

Algorithm and design improvements for indirect time of flight range imaging cameras

by

Benjamin Mark Moffat Drayton

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

Victoria
UNIVERSITY OF WELLINGTON

*Te Whare Wānanga
o te Ūpoko o te Ika a Māui*



2013

ABSTRACT

This thesis describes the development of a compact and modularised indirect time of flight range imaging camera. These cameras commonly use the Amplitude Modulated Continuous Wave (AMCW) technique. For this technique, an entire scene is illuminated with light modulated at a high frequency. An image sensor is also modulated and the phase shift introduced between the two modulation signals, due to the transit time of the light reflecting off objects in the scene and returning to the camera, is used to measure the distance.

The system constructed for this thesis is controlled by a Cyclone III FPGA and is capable of producing full field of view range images in real time with no additional computational resources. A PMD19K-2 sensor is used as the modulatable image sensor, and is capable of modulation frequencies up to 40 MHz.

One significant issue identified with this range imaging technology is that the precision of the range measurements are often dependent on the properties of the object being measured. The dynamic range of the camera is therefore very important when imaging high contrast scenes. Variable Frame Rate Imaging is a novel technique that is developed as part of this thesis and is shown to have promise for addressing this issue. Traditional theory for indirect time of flight cameras is expanded to describe this technique and is experimentally verified. A comparison is made between this technique and traditional High Dynamic Range Imaging. Furthermore, this technique is extended to provide a constant precision measurement of a scene, regardless of the properties of the objects in the scene.

It is shown that the replacement of the standard phase detection algorithm with a different algorithm can both reduce the linearity error in the phase measurements caused by harmonics in the correlation waveform and ameliorate axial motion error caused by relative motion of the camera and the object being measured. The new algorithm requires a trivial increase in computational power over the standard algorithm

and can be implemented without any significant changes to the standard hardware used in indirect time of flight cameras.

Finally, the complete system is evaluated in a number of real world scenarios. Applications in both 3D modelling and mobile robotics are demonstrated and tests are performed for a variety of scenarios including dynamic scenes using a Pioneer 2 robot.

ACKNOWLEDGEMENTS

I would like to extend my thanks to a number of people for their support and guidance. I would like to thank my primary supervisor Professor Dale Carnegie, who has patiently supported me and given me the freedom to pursue the topics I found most interesting, and my secondary supervisor, Dr. Adrian Dorrington, for providing me with in-depth knowledge on the subject matter. I would also like to thank the Chronoptics group at Waikato University for their input and advice, and the technicians at Victoria University of Wellington for efficiently procuring whatever I needed.

I would like to thank the Foundation for Research, Science and Technology for providing me with a Te Tipu Pūtaiao Fellowship and supporting me throughout my studies.

Finally, thank you to all my friends, family and loved ones who have provided me with support. Special thanks to my partner for supporting me through my studies and putting up with not seeing me for long stretches of time.

TABLE OF CONTENTS

ABSTRACT.....	III
ACKNOWLEDGEMENTS.....	V
TABLE OF CONTENTS	VII
LIST OF TABLES	XIII
LIST OF FIGURES	XV
CHAPTER 1 INTRODUCTION	1
1.1 OBJECTIVES.....	2
1.2 THESIS OUTLINE	4
CHAPTER 2 BACKGROUND THEORY.....	7
2.1 INTRODUCTION.....	7
2.2 CURRENT RANGE FINDING TECHNIQUES	7
2.2.1 <i>Triangulation</i>	7
2.2.1.1 Active.....	8
2.2.1.2 Passive.....	9
2.2.2 <i>Structured light</i>	10
2.2.3 <i>Time of flight</i>	10
2.2.3.1 Direct time of flight measurement.....	12
2.2.3.2 Shuttered Light Pulse	12
2.2.3.3 Indirect time of flight measurement	13
2.3 INDIRECT TIME OF FLIGHT SYSTEMS	18
2.3.1 <i>Elements of an Indirect ToF system</i>	18
2.3.2 <i>CMOS ToF sensors</i>	19
2.3.3 <i>Measurement Precision</i>	20
2.3.4 <i>Systematic Error Sources</i>	22
2.3.4.1 Harmonic Linearity error	23
2.3.4.2 Quantisation.....	24
2.3.4.3 Fixed Pattern Noise	24
2.3.4.4 Temporal variation in the modulation waveform	25
2.3.4.5 Multipath Error.....	25
2.3.4.6 Motion error.....	26
2.4 COMMERCIAL SYSTEMS.....	29
2.4.1 <i>Mesa Imaging SR-4000</i>	29

2.4.2	<i>PMDTech CamCube</i>	30
2.4.3	<i>Panasonic D-IMager</i>	30
2.4.4	<i>SoftKinetic DepthSense 311</i>	30
2.4.5	<i>Deficiencies of commercial cameras</i>	31
2.5	SUMMARY	31
CHAPTER 3	HARDWARE	33
3.1	INTRODUCTION	33
3.2	THE VICTORIA UNIVERSITY RANGE IMAGING SYSTEM	33
3.2.1	<i>The Prototype System</i>	33
3.2.2	<i>ToF Sensor</i>	33
3.2.3	<i>FPGA Development Kit</i>	35
3.2.4	<i>Illumination Board</i>	35
3.2.5	<i>Image Capture Board</i>	37
3.2.6	<i>External Interface Board</i>	38
3.2.7	<i>Problems with the prototype system</i>	39
3.3	DEVELOPMENT OF A COMPACT RANGE IMAGING SYSTEM.....	39
3.3.1	<i>Design Methodology</i>	40
3.3.2	<i>Image Capture Board</i>	43
3.3.3	<i>Illumination Board</i>	44
3.3.4	<i>FPGA Board</i>	46
3.3.5	<i>The External Interface Board</i>	48
3.3.6	<i>Firmware</i>	49
3.4	REVISION OF THE COMPACT RANGE IMAGING SYSTEM	51
3.4.1	<i>Changes To The Illumination Board</i>	51
3.4.2	<i>Changes To The Image Capture Board</i>	51
3.4.3	<i>Changes To The FPGA Board</i>	55
3.4.4	<i>Changes To The External Interface Board</i>	57
3.5	OPERATION OF THE SYSTEM	58
3.6	SETUP OF THE MOTION TESTING SYSTEM.....	59
3.6.1	<i>Linear Actuator</i>	60
3.6.2	<i>Control System</i>	61
3.6.3	<i>Interface</i>	64
3.7	SUMMARY	66
CHAPTER 4	DYNAMIC RANGE IMPROVEMENTS	67
4.1	INTRODUCTION	67
4.2	EXPERIMENTAL SETUP	67

4.3	CHARACTERISATION OF OUR SYSTEM	69
4.4	AUTOMATIC FRAME TIME ADJUSTMENT	76
4.5	HIGH DYNAMIC RANGE IMAGING	77
4.6	VARIABLE FRAME RATE USING PIXEL THRESHOLDING	81
4.6.1	<i>Introduction</i>	81
4.6.2	<i>Theory developed in this research</i>	84
4.6.3	<i>Dynamic range testing</i>	90
4.6.4	<i>Comparison with a commercial camera</i>	94
4.6.5	<i>Effect of axial motion</i>	99
4.6.6	<i>Improved thresholding metric</i>	101
4.7	SUMMARY	105
CHAPTER 5	PHASE DETECTION ALGORITHMS	107
5.1	INTRODUCTION	107
5.2	THEORETICAL ANALYSIS OF PHASE ALGORITHMS	109
5.3	STANDARD ALGORITHM (N-BUCKET)	116
5.3.1	<i>Background</i>	116
5.3.2	<i>Characteristic polynomial analysis</i>	117
5.3.3	<i>Linearity measurements</i>	117
5.3.4	<i>Motion Error Analysis</i>	119
5.4	FIVE FRAME STANDARD ALGORITHM	120
5.4.1	<i>Background</i>	120
5.4.2	<i>Characteristic polynomial analysis</i>	120
5.4.3	<i>Linearity Measurements</i>	121
5.4.4	<i>Motion Error Analysis</i>	125
5.5	ORDER CHANGED N-BUCKET	125
5.5.1	<i>Background</i>	125
5.5.2	<i>Characteristic polynomial</i>	127
5.5.3	<i>Linearity Measurements</i>	127
5.5.4	<i>Motion Error Analysis</i>	128
5.6	CARRÉ ALGORITHM	129
5.6.1	<i>Background</i>	129
5.6.2	<i>Characteristic polynomial analysis</i>	131
5.6.3	<i>Linearity measurements</i>	132
5.6.4	<i>Motion Error Analysis</i>	133
5.7	HARIHARAN ALGORITHM	134
5.7.1	<i>Background</i>	134
5.7.2	<i>Characteristic polynomial analysis</i>	135

5.7.3	Linearity measurements	136
5.7.4	Motion Error Analysis	136
5.8	N + 1 TYPE B ALGORITHM	137
5.8.1	Background.....	137
5.8.2	Characteristic polynomial analysis	140
5.8.3	Linearity measurements	140
5.8.4	Motion Error Analysis	142
5.9	NOVAK'S ALGORITHM.....	142
5.9.1	Background.....	142
5.9.2	Characteristic polynomial analysis	143
5.9.3	Linearity measurements	144
5.9.4	Motion Error Analysis	145
5.10	WDFT ALGORITHM.....	146
5.10.1	Background.....	146
5.10.2	Characteristic polynomial analysis.....	147
5.10.3	Linearity measurements	147
5.10.4	Motion Error Analysis	148
5.11	N + 3 ALGORITHM	149
5.11.1	Background.....	149
5.11.2	Characteristic polynomial analysis.....	151
5.11.3	Linearity measurements	151
5.11.4	Motion Error Analysis	152
5.12	COMPARISON OF ALGORITHMS.....	153
5.13	MULTIPLE VELOCITY MEASUREMENTS	157
5.14	CHANGE IN AMPLITUDE WITH DISTANCE	158
5.15	FPGA IMPLEMENTATION	159
5.16	COMPARISON OF PRECISION	160
5.17	APPLICATION SPECIFIC ALGORITHMS.....	162
5.18	SUMMARY	164
CHAPTER 6	REAL SYSTEM IMPLEMENTATION	165
6.1	INTRODUCTION	165
6.2	COMBINATION OF IMPROVEMENTS	165
6.3	COMPARISON OF RANGE IMAGES OF REAL SCENES.....	168
6.4	3D MODELLING.....	175
6.5	IMPLEMENTATION ON A MOBILE PLATFORM.....	176
6.5.1	Driving towards a wall.....	177
6.5.2	Detection of thin obstacles	183

6.5.3	<i>Evaluation of suitability for mobile robotics.....</i>	<i>184</i>
CHAPTER 7	CONCLUSIONS.....	187
7.1	THESIS SUMMARY	187
7.1.1	<i>Development of the compact Victoria University Range Imaging System</i>	<i>187</i>
7.1.2	<i>Implementation of the Windowed Discrete Fourier Transform.....</i>	<i>188</i>
7.1.3	<i>Implementation of VFRI.....</i>	<i>190</i>
7.1.4	<i>Combination into a real system.....</i>	<i>191</i>
7.2	PUBLICATIONS ARISING FROM THIS THESIS	191
7.3	FUTURE WORK	192
7.4	CONCLUSIONS	194
REFERENCES	195
APPENDIX A	PCB LAYOUTS	203
A.1	COMPACT SYSTEM.....	203
A.2	COMPACT SYSTEM REVISION 2.1	208

LIST OF TABLES

Table 3.1 Comparison of FPGA devices (Altera, 2011)	46
Table 3.2 Comparison of laser diode suppliers.....	51
Table 3.3 Voltage regulators for the compact range imaging system.....	57
Table 4.1 Average standard deviation of pixels using standard operating mode for different reflectivity regions	75
Table 4.2 Comparison of the standard deviation and frame rate of different reflectivity regions using a variable pixel rate	93
Table 5.1 List of the Relative Amplitudes of harmonics and their inverses in the Victoria University Range Imaging System.....	114
Table 5.2 Comparison of Motion Error and Linearity of phase detection algorithms	154
Table 5.3 Summary of phase detection algorithms.....	155
Table 6.1 Precision and linearity of different operating modes of the Victoria University Range Imaging System	167

LIST OF FIGURES

Figure 2.1 Diagram demonstrating the principle of triangulation.....	8
Figure 2.2 Diagram demonstrating the principle of Shuttered Light Pulse imaging	13
Figure 2.3 Comparison of the modulation signals for a close object (top) and a far object (bottom).....	14
Figure 2.4 Correlation waveform assuming sinusoidal modulation.....	16
Figure 2.5 Components of an indirect time of flight system	18
Figure 2.6 Diagram of an indirect time of flight pixel(Ringbeck & Hagebeuker, 2007).....	20
Figure 2.7 Theoretical error from axial motion for various speeds using 30 MHz modulation frequency, 25 ms integration time and 4 frames per measurement	28
Figure 2.8 Comparison of theoretical axial motion error versus distance between reference frames.....	29
Figure 3.1 Prototype Victoria University of Wellington Range Imaging System (McClymont et al., 2010)	34
Figure 3.2 Optical lens assembly photograph (left) and sensor IC (right)	34
Figure 3.3 Photograph of physical interface between the Illumination board and the sensor.....	36
Figure 3.4 Laser diode driver schematic	37
Figure 3.5 PMD19K-2 modulation driver circuit schematic.....	38
Figure 3.6 Schematic of modulation signal generation by a LVDS receiver	38
Figure 3.7 Photograph of the compact Victoria University Range Imaging System (McClymont et al., 2010)	40
Figure 3.8 Physical connections diagram of the compact range imaging system.....	41
Figure 3.9 Compact range imaging system board data flow diagram.....	43
Figure 3.10 Functional diagram of the Image Capture board	44
Figure 3.11 Functional diagram of the Illumination Board.....	45
Figure 3.12 Updated schematic of the laser diode driver	46
Figure 3.13 Functional diagram of the External Interface Board	48

Figure 3.14 AD5311 timing diagram (Analog Devices, 2010)	50
Figure 3.15 State machine diagram for setting the laser diode current	50
Figure 3.16 Image capture board temperature vs. modulation frequency.....	52
Figure 3.17 Image Capture Board temperature with and without fan (McClymont, 2010)	52
Figure 3.18 Comparison of modulation driver layouts between the original compact system (top) and the revised compact system (bottom). The drivers have been split into pairs and spread around the board	53
Figure 3.19 Schematic of the temperature control system implemented in the revised compact system	54
Figure 3.20 Functional diagram of revision 2.1 of the Image Capture Board	55
Figure 3.21 Functional diagram of revision 2.1 of the External Interface Board	58
Figure 3.22 Diagram of the standard imaging process implemented on the compact system FPGA	59
Figure 3.23 Photograph of the M6A linear actuator setup	60
Figure 3.24 NanoPro software interface	62
Figure 3.25 Distance vs. Time for the fastest stable motor profile (recorded using the encoder).....	63
Figure 3.26 Velocity vs. Time for the fastest stable motor profile (recorded using the encoder).....	63
Figure 3.27 T-slot sensor homing run	64
Figure 3.28 Linear Actuation control box	65
Figure 3.29 Motion Testing System interfaces	66
Figure 4.1 Photograph (top) and top view diagram (bottom) of apparatus for testing the dynamic range of indirect time of flight cameras.....	68
Figure 4.2 Layout of the three regions used for statistical analysis of the dynamic range	69
Figure 4.3 Standard Deviation of the measured phase versus Frame time for the bright region	70
Figure 4.4 Standard Deviation of the measured phase versus Frame time for the bright region (log-log).....	70

Figure 4.5 Measured Intensity of all four frames versus Frame Time for the bright region	71
Figure 4.6 Amplitude versus Frame Time for the bright region	71
Figure 4.7 Offset versus Frame Time for the bright region	72
Figure 4.8 Log-log plot of standard deviation versus amplitude for all three regions	73
Figure 4.9 \sqrt{B}/A versus Frame Time for the bright region comparison with measured precision.....	73
Figure 4.10 \sqrt{B}/A versus Frame Time for the bright region (log-log).....	74
Figure 4.11 Division of the standard deviation over the \sqrt{B}/A metric for different frame times.....	74
Figure 4.12 Typical range image using standard measurement procedure and 25 ms frame time	75
Figure 4.13 Typical range image using standard measurement procedure and 50 ms frame time	76
Figure 4.14 Typical range image using standard measurement procedure and 100 ms frame time	76
Figure 4.15 Operational diagram of the High Dynamic Range Imaging technique	78
Figure 4.16 Typical Short capture (top, 50 ms) and long capture (bottom, 200 ms) used for High Dynamic Range Imaging.....	80
Figure 4.17 High Dynamic Range Image (top) and selection matrix used (bottom) for a typical multi capture range image	81
Figure 4.18 Measured Phase versus frame time for the Victoria University Range Imaging System.....	82
Figure 4.19 Operational Diagram of Variable Frame Rate Imaging.....	83
Figure 4.20 Measured distribution of intensity values for the four frames used in phase calculation	84
Figure 4.21 Distribution of intensities for a well imaged area of a scene	85
Figure 4.22 Distribution of intensities for a poorly imaged area of a scene	85
Figure 4.23 Average intensity value versus number of measurements for the Victoria University Range Imaging System.....	87

Figure 4.24 Simulated amplitude versus number of measurements for different intensities.....	88
Figure 4.25 Measured amplitude versus number of measurements for different regions	88
Figure 4.26 Simulated standard deviation in phase when integrating over multiple measurements for a variety of amplitude values.....	89
Figure 4.27 Measured standard deviation versus number of frames integrated over for the dark region.....	90
Figure 4.28 Measured standard deviation versus number of frames integrated over for the sloped region	90
Figure 4.29 Measured distance (top) and frame rate (bottom) for a typical image using a frame time of 25 ms and a threshold value of 100	91
Figure 4.30 Measured distance (top) and frame rate (bottom) for a typical image using a frame time of 50 ms and a threshold value of 100	92
Figure 4.31 Standard deviation for different configurations using Variable Frame Rate Imaging.....	93
Figure 4.32 Mrasured intensity versus Frame Time for all four frames of the SR-4000 camera	95
Figure 4.33 Standard Deviation versus effective frame time comparing increasing frame time with integrating over multiple frames	95
Figure 4.34 Standard deviation versus number of measurements averaged over for all three regions using an SR-4000 commercial camera	96
Figure 4.35 Measured amplitude versus effective frame time for the sloped region using an SR-4000	97
Figure 4.36 Measured offset versus frame time for the sloped region using an SR-4000	97
Figure 4.37 Measured offset versus frame time for the dark region using an SR-4000 ...	98
Figure 4.38 Measured offset versus frame time for the bright region using an SR-4000.	98
Figure 4.39 Offset versus frame time comparison between normal background light and darkness in the sloped region using an SR-4000 camera	99
Figure 4.40 Theoretical Axial motion error versus phase when integrating over 1,2 and 3 measurements	101

Figure 4.41 Simulated standard deviation versus measured Amplitude for various correlation amplitudes.....	102
Figure 4.42 Simulated standard deviation versus \sqrt{B}/A for various correlation amplitudes	102
Figure 4.43 Standard Deviation versus threshold for the three regions using the SR400 camera with precision thresholding (top) and amplitude thresholding (bottom).....	104
Figure 4.44 Effective frame rate versus threshold for the three regions using the SR400 camera with precision thresholding (top) and amplitude thresholding (bottom).....	105
Figure 5.1 Simulated miscalibration error with and without the negative fundamental	112
Figure 5.2 Measured raw Intensity versus Sample Number using 64 steps per cycle	112
Figure 5.3 Relative Amplitude versus Frequency for the Victoria University Range Imaging System	113
Figure 5.4 Simulated linearity error of algorithms with different harmonic sensitivities	114
Figure 5.5 Comparison of motion error for different multiplicity of roots at $m = -1$	115
Figure 5.6 Comparison of the effect of motion on an algorithm with increasing multiplicity of roots at $m = 3$ and -3	115
Figure 5.7 Simulated linear miscalibration error for the standard algorithm using the measured harmonics for our system	116
Figure 5.8 Measured linearity error for the standard four frame algorithm	118
Figure 5.9 Comparison of measured linearity error between measurement techniques for the 4 step standard algorithm.....	118
Figure 5.10 Axial motion error versus distance for the standard four frame algorithm.	119
Figure 5.11 Linear miscalibration error for the four frame standard algorithm with a linear miscalibration of $\pi/40$	120
Figure 5.12 Measured linearity error for the standard five frame algorithm	121
Figure 5.13 Measured linearity error for the 64 frame N bucket algorithm.....	122
Figure 5.14 Comparison of Linearity Error between measurement techniques for the 5 step standard algorithm.....	123
Figure 5.15 Intensity versus Integration time for all four frames for the Victoria University Range Imaging System.....	124

Figure 5.16 Simulated error versus phase for non-linear pixels in an indirect time of flight camera	124
Figure 5.17 Axial motion error versus distance for the standard five frame algorithm..	125
Figure 5.18 Theoretical motion error versus distance for selected orderings of frame offset (Drayton et al., 2011).....	126
Figure 5.19 Comparison of motion error for different orderings of the standard algorithm with the third harmonic included	127
Figure 5.20 Linearity error for the order changed four frame algorithm.....	128
Figure 5.21 Axial motion error versus distance for the order changed four frame algorithm.....	129
Figure 5.22 Simulated relationship between measured phase and actual phase using Carré's algorithm	131
Figure 5.23 Theoretical linearity error for Carrés algorithm	131
Figure 5.24 Theoretical motion error versus distance for Carrés algorithm.....	132
Figure 5.25 Linearity error for Carrés algorithm.....	132
Figure 5.26 Axial motion error versus distance for Carré's Algorithm.....	133
Figure 5.27 Axial motion error versus distance for Carré's algorithm using phase unwrapping	134
Figure 5.28 Linearity error for Hariharan's algorithm	135
Figure 5.29 Axial motion error versus distance for Hariharan's algorithm	136
Figure 5.30 Linearity error for the four frame type B N+1 algorithm.....	141
Figure 5.31 Theoretical Linearity Error versus phase for the N+1 algorithm	141
Figure 5.32 Axial motion error versus distance for the four frame N+1 algorithm	142
Figure 5.33 Theoretical linearity error for Novak's algorithm	144
Figure 5.34 Theoretical motion error for Novak's algorithm	144
Figure 5.35 Linearity error for Novak's algorithm	145
Figure 5.36 Axial motion error versus distance for Novak's algorithm	145
Figure 5.37 Axial motion error versus distance for Novak's algorithm with phase unwrapped.....	146
Figure 5.38 Linearity error for the five frame WDFT algorithm	148
Figure 5.39 Theoretical Linearity Error versus phase for the WDFT algorithm.....	148

Figure 5.40 Axial motion error versus distance for the four frame WDFT algorithm	149
Figure 5.41 Linear miscalibration error for the five frame WDFT algorithm.....	150
Figure 5.42 Linearity error for the five frame N+3 algorithm.....	152
Figure 5.43 Theoretical Linearity Error versus phase for the five frame N+3 algorithm.	152
Figure 5.44 Axial motion error versus distance for the five frame N+3 algorithm	153
Figure 5.45 Comparison of axial motion error between all algorithms	156
Figure 5.46 Comparison of axial motion error versus distance for the four frame standard algorithm for multiple velocities.....	157
Figure 5.47 Comparison of axial motion error versus distance for the five frame WDFT algorithm for multiple velocities.....	158
Figure 5.48 Effect of inverse square decrease on the motion error for both the four frame standard and five frame WDFT algorithms	158
Figure 5.49 Comparison of axial motion error for the WDFT between Matlab and FPGA implementations.....	159
Figure 5.50 Distribution of standard deviations for the WDFT with processing done in Matlab	161
Figure 5.51 Distribution of standard deviations for the WDFT with processing done on the FPGA	161
Figure 5.52 Comparison of the standard deviation of range measurements versus distance for both the four frame standard algorithm and the five frame WDFT algorithm	162
Figure 5.53 Comparison of acceleration error between algorithms with different multiplicity of roots at $m = -1, 3$ and -3 for an acceleration of $\pi/80$ per frame	163
Figure 5.54 Comparison of linearity error using the phase stepping method	163
Figure 6.1 Linearity comparison of different operating modes implemented on the Victoria University Range Imaging System	166
Figure 6.2 Frame rate versus distance for both the VFRI and the WDFT + VFRI operating modes.....	168
Figure 6.3 Photo of Scene 1	169
Figure 6.4 Typical capture of Scene 1 using the standard system (left) and the WDFT + VFRI system (right)	170

Figure 6.5 Standard deviation of range measurements for Scene 1 using the standard system	171
Figure 6.6 Standard deviation of range measurements for Scene 1 using the WDFT + VFRI system	171
Figure 6.7 Photo of Scene 2	172
Figure 6.8 Typical capture of Scene 2 using the standard system (left) and the WDFT + VFRI system (right).....	172
Figure 6.9 Standard deviation of range measurements for Scene 2 using the standard system	173
Figure 6.10 Standard deviation of range measurements for Scene 2 using system with WDFT + VFRI.....	173
Figure 6.11 Typical measurement of Scene 2 using the WDFT + VFRI system at night ..	174
Figure 6.12 High angle view of Scene 2	174
Figure 6.13 Photograph of foam hand.....	175
Figure 6.14 Profile of model hand measured using the WDFT +VFRI system	176
Figure 6.15 Photograph of Pioneer 2 robot setup front view (left) and side view (right)	177
Figure 6.16 Photograph of the scene used for the straight line driving test.....	177
Figure 6.17 Four consecutive range images indicating lateral movement from first (top) to last (bottom).....	179
Figure 6.18 Initial range measurement for straight line driving test measurements using the standard system (left) and WDFT system (right).....	180
Figure 6.19 Comparison of range images for the straight line test using the standard algorithm (left) and the WDFT (right)	181
Figure 6.20 Comparison of the average measured distance of the seatback over four measurements between the standard operating mode and the WDFT mode	182
Figure 6.21 Model of the Robot and chair system	183
Figure 6.22 Photograph of the thin object detection scene.....	183
Figure 6.23 Initial partial detection of wire using the WDFT system	184
Figure 6.24 Comprehensive detection of wire using the WDFT system	185

Figure 7.1 Comparison between the prototype system (left) and newly developed system (right) (McClymont, 2010)	188
--	-----

Chapter 1 INTRODUCTION

Indirect time of flight cameras are increasingly being used in a wide variety of applications such as plant phenotyping (Alenya et al., 2011), profiling culturally significant objects (Chiabrando et al., 2011) and mobile robotics (Wiedemann et al., 2008). The field of mobile robotics is a particularly active field of research, with robots being used for applications such as urban search and rescue (Murphy, 2004), healthcare provision (Broadbent et al., 2009), the military (Voth, 2004) and space exploration (Katz & Some, 2003). A solution is required that can provide high precision, real time full field of view range images. Low latency real time measurements are crucial for mobile robotics as the environment is expected to change during the operation of the robot and navigation based on offline measurements is only possible in very limited and specific situations. In complex real world environments, obstacles such as thin objects, protruding objects and low or high objects require a full field of view range measurement to reliably detect and therefore safely navigate.

There are a number of technologies that have the potential to fulfil this requirement. One of the most promising is indirect time of flight cameras, which have recently become viable due to the integration of electronic modulation into CMOS imaging chips. As these range imaging systems are a relatively recent advancement, there is still a large amount of work required to ascertain the suitability of these cameras for mobile robotics and ameliorate deficiencies in these cameras.

Currently, the most common full field of view range imaging systems for mobile robotics are stereo vision systems and laser scanning systems. Stereo vision systems can provide high quality range measurements for a full field of view. However, for scenes with low texture the computational power required to match the objects in one camera's field of view with the corresponding object as viewed by the second camera is very large. A large number of mobile robots are relatively small and cannot provide this processing power. Laser scanning systems can take a considerable amount of time to scan the full

field of view and, due to the sequential method of acquiring data, moving objects may become segmented. Due to this segmentation it is preferable to acquire the entire range image simultaneously.

Commercial versions of indirect time of flight cameras are already available. However, the technology is still in its infancy and a number of deficiencies of these cameras, particularly for mobile robotics applications, are yet to be addressed.

1.1 Objectives

The objective of this thesis is to develop an indirect time of flight camera for the purposes of research into the suitability of this technology for applications such as mobile robotics. To achieve this, the system requires the following properties:

- Compactness – as the system is intended to be used in real world tests it is crucial that it can be mounted on a mobile robot. The requisite size is obviously robot dependent. For the purposes of this research, compactness indicates a suitability to be implemented on a Pioneer 2 robot (Adept MobileRobots, Amherst, NH, US), a common two wheeled robot with a maximum speed of 1.6 m/s and a payload capacity of 20 kg. Further miniaturisation is possible without compromising the quality of the measurements. This is in contrast with stereo vision systems where the measurement quality is dependent on the separation of the two cameras.
- A simple power supply structure – commonly mobile robots will be powered using a single battery and therefore the camera must be capable of operating from this type of power source. It should have a single, unregulated DC power input.
- Configurability – to be useful for research purposes it is essential that as much control as possible is afforded to the operator of the system. In particular, precise control of the modulation of both the sensor and the light source, and of the method used to calculate the phase value, is highly desirable. Parameters such as the integration time and the illumination power should also be adjustable in real time. The system should also be modular to allow for future improvements without having to design an entirely new system.

- Computational self sufficiency – many mobile robotic systems have limited computational power, mainly due to the limited weight bearing capacity and power available. This computational power must be shared by a number of control and sensor systems on the robot and therefore any solution for mobile robotics should ideally be computationally self sufficient, as the solution will then be usable by the largest number of mobile robotic platforms. For our system, this means all processing must be performed on an onboard FPGA.
- Future proofed – to the degree possible, future proofing should be provided. A prime example of this is in the resolution of the sensor, which can be reliably expected to increase. Computational power and memory should be selected based on this expectation.

One disadvantage of indirect time of flight cameras that has already been identified is the dependence of their precision on the properties and distance of the objects being measured. As diffused light is used to illuminate the scene, there is an inverse square decrease in intensity with distance. The received illumination can therefore change by several orders of magnitude over the scene, making the dynamic range of the camera of critical importance. It is desirable to measure the entire scene with high precision, therefore methods of increasing the dynamic range of these cameras is required.

A further problem of interest to mobile robotics is the effect of motion on the distance measurements. Moving objects in the scene, or motion of the camera, introduces systematic errors into the distance measurements. This error can be separated into axial motion error and lateral motion error. Current solutions have focused on the issue of lateral motion error. For these cameras to be used for mobile robotics applications, a method must be developed to ameliorate the systematic error due to axial motion.

The final objective for this thesis is to run a series of real world tests using the developed system. This will evaluate the suitability of the system for mobile robotics using a Pioneer 2 robot, as well as documenting the general performance of the system in a number of situations.

1.2 Thesis outline

The objectives of this thesis are addressed in 7 chapters:

Chapter 2 presents relevant background theory on distance measurement techniques, with an emphasis on range imaging. The error sources of indirect time of flight cameras, both systematic and random, are described and current techniques for ameliorating them are discussed. Finally, the constituent components of an indirect time of flight range imaging system are outlined and an overview of the current commercial cameras is provided.

Chapter 3 describes the development of the hardware for the compact Victoria University Range Imaging System. The initial prototype design and its advantages and disadvantages are first explained and then two revisions of a new compact system, completed as part of this thesis, are described. This compact system has been instrumental in allowing for real world testing of the range imaging system.

Chapter 4 presents methods to improve the dynamic range of indirect time of flight cameras. High Dynamic Range Imaging, a technique reported throughout the literature, is outlined and demonstrated using the Victoria University Range Imaging System. A new method is then proposed and compared to High Dynamic Range Imaging. A theoretical framework for the new method is developed and its efficacy is confirmed experimentally.

Chapter 5 presents improvements to the phase detection algorithm used in indirect time of flight cameras. A method for theoretically evaluating phase detection algorithms is introduced and, following a literature review, algorithms are evaluated both theoretically and experimentally. The best performing algorithm is tested further and shown to provide significant benefits over the standard algorithm.

Chapter 6 documents the combination of the work from Chapters 3, 4 and 5 into a real world system, and the experimental testing of that system. The system is evaluated in both static and dynamic environments and a number of potential applications are explored. Real world tests are performed to evaluate the system for use in mobile robotics.

Chapter 7 concludes this thesis. The major components of the system developed as part of this thesis are reviewed and future work is suggested to improve the system further. A list of publications arising from this thesis is provided and the novel contributions of this thesis are summarised.

Chapter 2 BACKGROUND THEORY

2.1 Introduction

This chapter provides an overview of range finding techniques currently reported in the literature, outlining their advantages and disadvantages. The operational methodology and theoretical framework of indirect time of flight systems is explored, including the recent advances in CMOS technology allowing for these systems to become compact and suitable for mobile robotic platforms. An analysis of the error sources involved in indirect time of flight measurements is then performed. Finally, commercial systems currently available are described.

2.2 Current range finding techniques

Numerous techniques have been developed to determine the distance to an object or objects. These are primarily based on triangulation, structured light or time of flight. This section will present an overview of the technologies currently available and their advantages and disadvantages, particularly in relation to mobile robotics applications.

2.2.1 TRIANGULATION

Triangulation is based on the geometric properties of right angled triangles. Measurement of the angle from two points to an object, as well as a known distance between the two points, is sufficient for measurement of the distance to that object. This configuration is shown in Figure 2.1.

The relationship between the base line lengths L_1 and L_2 , the angles to the object θ_1 and θ_2 and the perpendicular distance to the object D is given by:

$$L_1 = \frac{D}{\tan \theta_1} \text{ and } L_2 = \frac{D}{\tan \theta_2}. \quad (2.1)$$

Adding these equations together gives

$$L_1 + L_2 = \frac{D}{\tan \theta_1} + \frac{D}{\tan \theta_2}, \quad (2.2)$$

which can be rearranged to make D the subject giving:

$$D = \frac{L_1 + L_2}{\left(\frac{1}{\tan \theta_1} + \frac{1}{\tan \theta_2}\right)}. \quad (2.3)$$

Triangulation range finders can be either active, meaning they emit a signal and measure its reflection, or passive, which use background illumination.




Figure 2.1 Diagram demonstrating the principle of triangulation

2.2.1.1 ACTIVE

Active triangulation sensors generally measure the position of the returning light beam with a set angle, instead of measuring the angles θ_1 and θ_2 directly. These are often referred to as Position Sensitive Detectors (PSDs). A light emitter is set at a specific angle and the position of the reflected beam is captured by a linear photosensitive sensor. The further away the object being detected is, the further across the detector the returning signal will appear. Because this technique does not depend on the amplitude of the returning wave, it is not affected by the reflectivity of the object it is measuring, provided a minimum level of intensity is received by the PSD.

Active triangulation sensors can either be a single point sensor or, by using a cylindrical lens, the light source can be formed into a line detecting all objects along this line. In this

case a two dimensional detector is required. Often a charged coupled device (CCD) camera is used as they are inexpensive and readily available.

The range of a PSD is dependent on the angle of the light source and the width of the detector. If an object is too close, the beam will be reflected back onto the emitter or into the gap between the emitter and the detector, causing the reflected signal to not be detected. Therefore these sensors have a minimum measureable distance. If the object is too far away, the reflected light will miss the detector and therefore the distance cannot be measured.

Several of Victoria University's robotic vehicles use PSDs (GP2Y3A003K0F and GP2Y3A002K0F) from Sharp (Osaka, Japan) (McClymont, 2010). Both of these PSDs use 5 individual LEDs at 5° offsets to give a 25° vertical field of view and have ranges of 0.4 m to 3 m and 0.2 m to 1.5 m respectively. Both sensors retail for ~\$100.

While active triangulation sensors can be useful, a full field of view range measurement is often desired. Single point and line sensors are useful for detecting walls and some other objects, but since they do not provide a full field of view they are not ideal for navigating complex real world environments. It is possible to use multiple sensors, however this introduces problems with interference between the sensors if they are run simultaneously or latency if they are polled. The size and cost also increases proportional to the information received.

2.2.1.2 PASSIVE

Passive triangulation (more commonly known as stereo vision) uses two cameras set a specific distance apart. By using correspondence algorithms, which relate objects in one camera's field of view with the same objects in the other camera, the distance to objects within the field of view can be measured. This is effectively the same principle as used for depth perception by human eyes. The main difficulty for these range finding systems is the complexity of the correspondence problem, particularly for homogeneous or untextured scenes. Advanced algorithms can be used to improve the performance of these cameras, however, they are accompanied by an increase in computational effort.

Because of the image processing used, this method requires significant computational power, although real time systems are possible on a mid-range laptop.

Stereo vision systems do provide a full field of view range measurement. However, due to the correspondence problem, the accuracy of the data is dependent on the properties of the object being measured (such as homogeneity, contrast and the quality of light available) and on the computational power of the system.

2.2.2 STRUCTURED LIGHT

Structured light works by projecting a set pattern onto the scene being imaged and observing how the pattern deforms over the objects in the scene. Simple systems can be implemented using a single line of light, while more advanced systems are capable of imaging an entire field of view using an infrared projector.

A common structured light system is the Kinect sensor used on the popular gaming system, the Xbox 360. It has a field of view of 57° horizontally and 43° vertically and outputs video frames at 30 fps with a resolution of 640×480 pixels. Depth is calculated from measurements recorded as 10 bit disparity values. The depth resolution decreases with increasing distance, with 1 cm depth resolution at 2 m and 7 cm resolution at 5 m (Khoshelham & Elberink, 2012).

The main drawback for structured light systems is that the resolution of the system is dependent on the separation between the observing camera and the light source. For mobile robotics applications compactness is a highly desirable property so this requirement for a separation between emitter and receiver is a significant disadvantage.

2.2.3 TIME OF FLIGHT

Distance can be measured by emitting a signal of known velocity and measuring the time taken for it to reflect back off an object. This is known as the time of flight principle. Radar and Sonar are two widely known time of flight systems. For short distance range finding, historically only sound waves have been used, as the speed of sound is much slower than that of electromagnetic waves and therefore the timing electronics do not have to be as accurate. The main disadvantage of sound waves is that it is difficult to

maintain a narrow beam. Achieving a high spatial resolution is therefore difficult. The precision of the measurements is determined by the frequency of the sound waves. High frequency sound waves can be used to make high resolution measurements, however they are heavily attenuated by air. Low frequency sound waves have lower resolution, however they are less attenuated in air and so a larger measurement distance can be achieved. Indicatively, the X1 Ranging module Pro (SensComp, 2011), an ultrasonic sensor for mobile robotics, has a frequency of 50 kHz, an operating range of 0.15 m to 6.1 m and a beam width of 15°. With an aperture size of 4.3 cm the maximum angular resolution of this sensor is 11°. The attenuation of sound waves in air for a frequency of 50 kHz, using a temperature of 20 °C and a humidity of 60%, is 1.7 dB/m (Jakevicius & Demcenko, 2008). To achieve high resolution images, such as used in ultrasound imaging, frequencies of 1 MHz or higher are used. At this frequency the attenuation in air using the same parameters is 164 dB/m (Jakevicius & Demcenko, 2008) and the maximum angular resolution is 0.5° (assuming the same aperture size). Because of this high attenuation, high resolution ultrasound can only be used for close range imaging in mediums with lower attenuation, such as water and human tissue.

Common to both waveforms is the problem of multiple reflections, where the signal reflects more than once, altering the return time and potentially interfering with the current or future measurements.

For time of flight the distance to an object is measured as

$$d = \frac{vt}{2}, \quad (2.4)$$

where d is the distance, v is the velocity of the signal and t is the time taken for the signal to return to the sensor.

Increases in the speed of electronic devices have now made it possible to use electromagnetic waves for short distance time of flight applications. Visible or infrared light is desirable as it is significantly easier to focus than sound, with narrow laser beams possible over long distances. However, since it is travelling six orders of magnitude faster, the timing accuracy required to measure it directly is much higher. There are several approaches that can be used for time of flight measurement. The most common

are measuring the time directly for a light pulse, Shuttered Light Pulse and encoding the time into a phase shift between two modulation waves (often referred to as Indirect Time of Flight).

2.2.3.1 DIRECT TIME OF FLIGHT MEASUREMENT

Direct time of flight sensors use a diode or laser to emit a pulse of light then measure the time taken for the pulse to return. Indicatively, to measure an object to an accuracy of 1 mm, the timer must be able to resolve the timing of the pulse to within 7 ps. The pulse generator also has to be able to operate at a very high speed.

Direct time of flight systems are available in single point, line and full field varieties. A line sensor is normally made by dispersing the light in one direction similar to PSDs, however this increases the required complexity of the detection circuitry. Direct time of flight systems can be expanded to provide a full field range image by mechanically moving the emitter across the field of view (Besl, 1988) or directing the beam using a mirror. While this greatly reduces the complexity of obtaining multiple data points compared to using multiple sensors, it does introduce the problem of mechanical wear. The accuracy of the system is generally dependent on the accuracy of the mechanical stage, rather than the accuracy of the sensor itself, and it can take a considerable amount of time to scan the entire field of interest (Carnegie et al., 2005).

There is currently research into the use of a 2D array of single-photon avalanche diodes (SPADs), which have the potential to provide direct measurement of a full field of view (Niclass et al., 2008), but this technology is still currently in its infancy with only very low resolutions available.

2.2.3.2 SHUTTERED LIGHT PULSE

For the Shuttered Light Pulse technique, an illumination pulse is emitted and a gating mechanism allows only a portion of the reflected light to be observed by the sensor. This results in the intensity observed being dependent on distance. As the optical properties of objects in the scene is expected to be variable, a second calibration measurement is required. In this measurement the gate is kept open and therefore the full reflected

pulse from the scene is acquired. The distance the object can then be calculated using the ratio of the intensity values measured using the equation

$$d = \frac{cT_0(1 - I_1/I_2)}{2}, \quad (2.5)$$

where c is the speed of light, T_0 is the width of the light pulse and I_1 and I_2 are the intensity measurements of the range frame and calibration frame respectively (Christie et al., 1995). Figure 2.2 shows this operation.

Using an image intensifier and a CCD camera this method can provide a full field of view range measurement, however image intensifiers require high voltage power supplies to operate and are bulky and expensive devices that are not suitable for mobile applications.

2.2.3.3 INDIRECT TIME OF FLIGHT MEASUREMENT

Requiring extremely high speed electronics can be avoided by encoding the time information into another signal. The time information can be encoded into a phase shift by modulating both the illumination source and a shutter on the sensor at high frequency, generally on the order of 10 - 100 MHz. This is called Amplitude Modulated Continuous Wave imaging.




Figure 2.2 Diagram demonstrating the principle of Shuttered Light Pulse imaging

The time taken for the light to travel to the object and return introduces a phase shift between the illumination and shutter modulation waveforms. This phase shift as a fraction of an entire cycle, multiplied by the modulation period, gives the effective time for the signal to reflect from the object as

$$t = T_{mod} \frac{\varphi}{2\pi} = \frac{\varphi}{2\pi f_{mod}}, \quad (2.6)$$

where φ is the phase shift introduced by the time taken for the light to travel to the object and back, and f_{mod} is the modulation frequency. Substituting (2.6) into (2.4) and using the speed of light c as the velocity gives

$$d = \frac{c\varphi}{4\pi f_{mod}}. \quad (2.7)$$

A diagram demonstrating the phase shift is shown in Figure 2.3. For a closer object (top) the phase shift introduced in the reflected light is slight and most of the light is collected by the sensor. For an object further away (bottom) the shift is greater and less light is collected by the sensor.




Figure 2.3 Comparison of the modulation signals for a close object (top) and a far object (bottom)

Naively, the phase shift could be measured by the intensity of the light, as the phase shift between the sensor shutter and the light modulation correlates intensity with distance. However, intensity is also affected by the reflectivity of the object, the inverse square decrease with distance caused by waves spreading from a point source and background light levels. This problem can be solved by taking multiple frames, with introduced phase offsets, and measuring the phase of the resulting waveform. This waveform is commonly referred to as the correlation waveform, as it represents the correlation between the illumination and sensor modulation signals. The phase of this

waveform is the phase we wish to measure, as the first measurement point will have no additional phase added.

To measure distance using this technique, an equation has to be found for measuring the phase of the waveform resulting from the introduced phase offsets. For simplicity, the sensor and illumination waveforms are assumed to be sinusoidal, therefore the reflected illumination signal, $r(t)$, can be represented as (Jongenelen, 2010)

$$r(t) = R \cos(2\pi f_{mod} t - \varphi) + R_0, \quad (2.8)$$

where R is the average received illumination signal amplitude, accounting for the average transmitted amplitude, the inverse square relationship due to spreading waves and the reflectivity of the object being measured, t is the time delay associated with round trip travel and R_0 is a DC offset accounting for background lighting and imperfections in the illumination waveform.

Similarly the sensor modulation can be written as

$$s(t) = S \cos(2\pi f_{mod} t - \theta) + S_0, \quad (2.9)$$

where S is the average sensor gain, S_0 is the sensor's DC offset and θ is a phase offset introduced in the sensor modulation.

The returning modulated waveform is integrated over a large number of cycles. This results in a pixel intensity $I(\varphi, \theta)$ that is the multiplication of the reflected wave and sensor modulation integrated over the period T . This can be evaluated:

$$\begin{aligned} I(\varphi, \theta) &= \int_0^T r(t) \times s(t) dt \\ &= \int_0^T [R \cos(2\pi f_{mod} t - \varphi) + R_0] [S \cos(2\pi f_{mod} t - \theta) + S_0] dt \\ &= \int_0^T [RS \cos(2\pi f_{mod} t - \varphi) \cos(2\pi f_{mod} t - \theta) + RS_0 \cos(2\pi f_{mod} t - \varphi) \\ &\quad + R_0 S \cos(2\pi f_{mod} t - \theta) + R_0 S_0] dt \\ &= \int_0^T \left[\frac{RS}{2} \cos(\theta - \varphi) + \frac{RS}{2} \cos(4\pi f_{mod} t - \theta - \varphi) + RS_0 \cos(2\pi f_{mod} t - \varphi) \right. \\ &\quad \left. + R_0 S \cos(2\pi f_{mod} t - \theta) + R_0 S_0 \right] dt \\ &= \frac{RST}{2} \cos(\theta - \varphi) + \frac{RS}{8\pi f_{mod}} \sin(4\pi f_{mod} t - \theta - \varphi) \Big|_0^T + \frac{RS_0}{2\pi f_{mod}} \sin(2\pi f_{mod} t - \varphi) \Big|_0^T \\ &\quad + \frac{R_0 S}{2\pi f_{mod}} \sin(2\pi f_{mod} t - \theta) \Big|_0^T + R_0 S_0 T \end{aligned} \quad (2.10)$$

The integration time, T , is much larger than the modulation period ($1/f_{mod}$), therefore the 2nd, 3rd and 4th terms in (2.10) are insignificant compared to the 5th term, which is constant with respect to φ and θ . Therefore (2.10) can be simplified to:

$$\begin{aligned} I(\varphi, \theta) &= \frac{RST}{2} \cos(\theta - \varphi) + R_0 S_0 T \\ &= A \cos(\theta - \varphi) + B. \end{aligned} \quad (2.11)$$

For each range measurement N frames are recorded. The sensor modulation phase offset θ is stepped by $2\pi/N$ radians each frame. Therefore the intensity at frame n ($n = 0 \dots N-1$) is

$$I_n = A \cos(n\delta - \varphi) + B, \quad (2.12)$$

where δ is the phase step. A diagram of the correlation waveform, assuming sinusoidal modulation, is shown in Figure 2.4. The offset B can be separated into offset due to ambient light, A_{amb} , and offset due to active illumination, A_{sig} .

From the Discrete Fourier Transform of the correlation waveform, the parameters A , B and φ can be calculated as shown below

$$\varphi = \tan^{-1} \left(\frac{\sum_{n=0}^{N-1} I_n \sin\left(\frac{2\pi n}{N}\right)}{\sum_{n=0}^{N-1} I_n \cos\left(\frac{2\pi n}{N}\right)} \right) \quad (2.13)$$




Figure 2.4 Correlation waveform assuming sinusoidal modulation

$$A = \frac{2}{N} \sqrt{\left[\sum_{n=0}^{N-1} I_n \cos\left(\frac{2\pi n}{N}\right) \right]^2 + \left[\sum_{n=0}^{N-1} I_n \sin\left(\frac{2\pi n}{N}\right) \right]^2} \quad (2.14)$$

$$B = \frac{1}{N} \sum_{n=0}^{N-1} I_n. \quad (2.15)$$

It is common in many systems to use four frames for each range measurement (Lange & Seitz, 2001) (Blanc et al., 2004) (Hussmann et al., 2011) as this greatly simplifies the phase calculation to

$$\varphi = \tan^{-1} \left(\frac{I_1 - I_3}{I_0 - I_2} \right) \quad (2.16)$$

and the amplitude calculation to

$$A = \frac{1}{2} \sqrt{(I_0 - I_2)^2 + (I_1 - I_3)^2}. \quad (2.17)$$

This technique has the desired result of removing the amplitude dependence of the phase calculation. There is still an ambiguity problem however. If an object is at a distance such that it takes greater than 2π radians of phase for the light to return, it will be indistinguishable from a closer object. The distance is therefore more accurately characterised by the equation

$$d = \frac{c}{2f} \left(\frac{\varphi}{2\pi} + k \right), \quad (2.18)$$

where k is an integer. It has been shown that the unambiguous measurement distance can be extended significantly by acquiring measurements with multiple modulation frequencies (Payne et al., 2009), however long range measurements are also limited by eye safety concerns for many applications.

Indirect time of flight systems are possible using all the methods outlined for direct time of flight (point and line) but with the added possibility of making a full field of view system by illuminating the entire scene with modulated light and using a two dimensional detector with a fast shutter, instead of requiring an array of SPADs.

Indirect time of flight systems are capable of providing a full field of view range measurement with high frame rates. These cameras have small physical dimensions, have no moving parts and require little computational power. Therefore these cameras

are a potential solution to the sensor requirements for autonomous mobile robots. This thesis focuses on overcoming some of the remaining problems with these cameras, which are discussed later in this chapter.

2.3 Indirect Time of Flight systems

2.3.1 ELEMENTS OF AN INDIRECT TOF SYSTEM

The core components required for an indirect time of flight camera are:

- A light emitter – This is used to provide modulated light to illuminate the scene. The light source must be capable of switching at high frequencies.
- A light detector – An intensity based detector is required to measure the returning light. A sensor array and imaging optics are normally used to provide a full field of view.
- A high speed shuttering system – A method of modulating the detector on the order of 10 MHz – 100 MHz is required to generate a modulation phase change that can be measured.

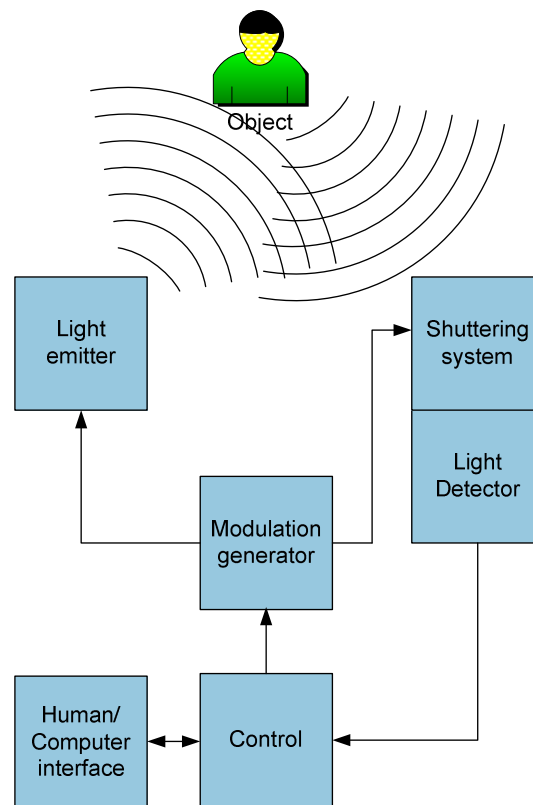


Figure 2.5 Components of an indirect time of flight system

- A modulation source – Multiphase high frequency modulation signals are required to make distance measurements that are independent of factors external to the camera.
- A control component – A control component is required to coordinate the phase of the modulation signals and perform the phase calculation.
- A human / computer interface – An interface for outputting the range measurements to either a computer or a human for interpretation is required for the system to be implemented in real world applications

Figure 2.5 demonstrates a system with these basic components.

2.3.2 CMOS ToF SENSORS

Until recently, image intensifiers were commonly used to provide a shutter capable of operating at 10 MHz or more (Cree, et al. 2006). The output of the image intensifier was then viewed with a standard CCD camera to provide a full field of view. Image intensifiers are not suitable for mobile robotics as they require high voltage power supplies (on the order of 600 V), have a high unit cost and have a large physical size. The development of Complementary Metal Oxide Semiconductor (CMOS) based sensors has decreased the cost, size and power requirements to a level where it is practical for these cameras to be used on mobile robots. These devices provide custom intensity pixels that can be shuttered electronically using low voltage signals, however, they have low resolution by modern camera standards. Current commercial models offer resolutions up to 200×200 pixels, however 1 MegaPixel sensors are expected in the near future.

Each sensor pixel is a 5 terminal device with two modulation input gates and two readout gates. A diagram of a typical pixel is shown in Figure 2.6. If no voltage is applied to the modulation gates, the charge carriers stimulated on the sensor surface will drift equally to both readout gates. However, if a complimentary modulation voltage is applied to the modulation inputs a potential is developed between the two readout gates. If a 50% duty cycle signal is used, in the presence of background light, accumulated charge will be equally shared across the two readout gates, as 50% of the time charge will be directed towards one gate and 50% of the time it will be directed

towards the other. In the presence of light modulated at the same frequency as the modulation inputs, charge will be shared between the readout gates based on the relative phase of the pixel modulation signal and the illumination modulation signal. After each integration period, the output voltage of both gates is subtracted. This cancels the background illumination and provides a voltage proportional to the relative phase of the two modulation signals.




Figure 2.6 Diagram of an indirect time of flight pixel (Ringbeck & Hagebeuker, 2007)

2.3.3 MEASUREMENT PRECISION

The precision of indirect time of flight cameras, as a 1-sigma standard deviation, can be expressed as (Jongenelen, 2010)

$$\sigma_\phi = \frac{\sqrt{B}}{\sqrt{2}c_d A_{sig}}, \quad (2.19)$$

where c_d is the demodulation contrast of the sensor. The demodulation contrast is the ratio of the amplitude of the correlation waveform to the offset of the correlation waveform due to the illumination (Payne et al., 2011), or

$$c_d = \frac{A}{A_{sig}}. \quad (2.20)$$

Therefore (2.19) can be rewritten as

$$\sigma_\phi = \frac{\sqrt{B}}{\sqrt{2}A}. \quad (2.21)$$

There are a variety of options for improving the precision of the phase measurements. Generally this will not include improving the demodulation contrast, as the demodulation contrast is primarily determined by the sensor architecture.

A number of hardware based methods can be used to minimise this ratio. Optical filters can be used to decrease A_{amb} , and hence improve the precision. The optical power of the illumination source can be increased, however if the system is to be used in public areas this is limited by eye safety concerns. When operating the camera, the most readily available method for improving the precision of the measurements is to increase the frame time. Both B (Jongenelen, 2010) and A (Falie & Buzuloiu, 2007) are expected to be proportional to the frame time and therefore increasing the frame time should provide an inverse square root decrease in the standard deviation of the phase measurements.

The standard deviation of the distance measurements is related to the standard deviation of the phase measurements by the equation

$$\sigma_d = \frac{c\sigma_\phi}{4\pi f_{mod}}. \quad (2.22)$$

Substituting (2.19) into (2.22) gives the equation for the standard deviation of the distance measurements as

$$\sigma_d = \frac{c}{4\pi f_{mod}} \frac{\sqrt{B}}{\sqrt{2}c_d A_{sig}}. \quad (2.23)$$

This equation shows that the precision of the range measurements can be improved by increasing the modulation frequency. However, this comes with the trade off of decreasing the maximum unambiguous measurement distance. Furthermore, as the demodulation contrast will decrease with increasing modulation frequency, there is an optimal operating modulation frequency that maximizes the distance precision.

Currently commercial camera manufacturers, such as Mesa Imaging and Canesta, use optimisation of the frame time as a method for improving the precision of measurements taken by their cameras. The Mesa Imaging SR-4000 uses an optimisation based on a histogram of the intensity values in the cameras field of view (Mesa Imaging, 2011). It takes four input values, the minimum acceptable frame time, the maximum

acceptable frame time, percentOverPos and desiredPos. PercentOverPos is a value between 0 and 100 indicating the percentage of values in the histogram desired to be above the intensity value desiredPos (0-255). The frame time is adapted slowly to shift the intensity value PercentOverPos values are above until it matches desiredPos. The recommended settings for the SR-4000 are 1,150,5,70 (Mesa Imaging, 2011).

Another implementation, used in cameras manufactured by Canesta (Sunnyvale, CA, USA), assigns pixels a quantized value representing the intensity of the pixel. The pixels are then categorised into three bins, pixels with too much light (saturated), pixels with an acceptable amount of light and pixels without enough light (Gokturk & Rafii, 2008). If the number of pixels with an acceptable amount of light is over a threshold then no action is taken. Otherwise, one or more system parameters (frame time, common mode resets, video gain and potentially others) are adjusted, depending on whether there are more pixels with too much or too little light.

As well as these commercial implementations, a number of methods have been reported in the literature. These include maintaining the mean amplitude value of the scene at a set value (Wiedemann et al., 2008), estimating and trying to maximize the mean accuracy (May et al., 2006) and calculating optimal integration times for a visual servo control task based on the mean amplitude (Pomares et al., 2010).

Optimisation of the frame time does not improve the fundamental dynamic range of the camera for a particular configuration, it simply optimises the camera configuration for the scene. Ensuring bright objects in the scene do not saturate can deteriorate the quality of the range image in darker areas. There are diminishing returns for increasing the frame time when the area is already reasonably well imaged. Therefore, optimising the intensity using this method can cause the frame rate for bright objects to be lower than is required to acquire a good image of them, as the frame time is increased to try to get a high quality image in darker areas of the scene.

2.3.4 SYSTEMATIC ERROR SOURCES

This section will provide an analysis of the sources of systematic error in indirect time of flight measurements. A thorough understanding of these errors is required to undertake

an improvement in the quality of the data from these cameras. This section is split into several sub-sections each addressing a particular source of error.

2.3.4.1 HARMONIC LINEARITY ERROR

In section 2.2.3.2 an equation for the phase was derived assuming a sinusoidal correlation signal, which requires at least one of the modulation signals to be sinusoidal. In reality, due to the ease with which they can be generated digitally, square wave modulation signals are normally used. Furthermore, the non-linear transfer function of the illumination source and sensor pixels also has an effect on the correlation waveform. These factors introduce harmonics into the system, which are expressed as a sinusoidal linearity error with distance. For the $\pm m^{\text{th}}$ harmonic a sinusoidal error with $m \mp 1$ cycles is observed within the unambiguous measurement distance. It is known that a four cycle error is observed for current systems (Jongenelen et al., 2009), this is due to the negative 3rd and positive 5th harmonics.

There are several factors that can impact the harmonic content of the correlation waveform. Slight differences in pixels due to the CMOS fabrication process cause slight differences in harmonic response between different pixels on the same sensor. There may also be systematic variations in the harmonic content due to the propagation of the modulation signals across the sensor (Drayton et al., 2012 b). The choice of modulation frequency will also influence the harmonic content of the correlation waveform, as the harmonics are attenuated due to the bandwidth limitations of components in the system. Having a variable modulation frequency can be advantageous due to the trade off between the precision and the unambiguous measurement distance described in section 2.3.3. The internal temperature of the components in the system can also affect the harmonic content of the correlation waveform.

A number of different attempts have been made to calibrate this error using sinusoids (Chiabrando et al., 2009), 6th order polynomials (Kim et al., 2008), b-spline fitting (Lindner & Kolb, 2006) (Fuchs & Hirzinger, 2008), and look-up tables (Kahlmann et al., 2006). Some of these techniques were reasonably successful, particularly b-splines and look-up tables. However, as described above, this error is dependent on temperature,

the particular pixel used, and the modulation frequency. Generating a comprehensive calibration taking into account all of these variables is therefore very difficult. Generally a Fixed Pattern Noise calibration is combined with a distance calibration to try to alleviate the spatial variation, which is discussed further in Section 2.3.4.3.

A number of other approaches have been tried to mitigate this error. Instead of treating the modulation signal as sinusoidal, it was treated as a triangular wave, and various intermediaries between the two (Lindner et al., 2008). While this can provide improved linearity, the correct correlation waveform to use is system dependent, as the harmonic content will differ between hardware implementations as well as with the factors described previously. Harmonic cancelation using frame encoding (Payne et al., 2010) and the use of a heterodyne operating mode (Conroy et al., 2009) both provide a solution that removes harmonic linearity error and is not dependent on the exact harmonic content of the system. However, both of these methods complicate the relationship between the measured phase and the actual phase when the object being measured is moving, potentially complicating solutions to this error source. This motion error is discussed further in Section 2.3.4.6.

2.3.4.2 *QUANTISATION*

The conversion of an analog voltage into a digital number introduces a quantisation error into the phase calculation. It has been shown that the quantisation of the phase values is influenced not only by the resolution of the ADC, but also by the amplitude of the correlation waveform (Frank et al., 2009). For low amplitude returns, the set of possible phases that can be calculated becomes sparse resulting in a large quantisation error.

2.3.4.3 *FIXED PATTERN NOISE*

Due to the construction of the sensor, indirect time of flight systems are susceptible to fixed pattern noise (May et al., 2009). There are two sources of this error. Firstly the construction of CMOS devices means that homogeneity between pixels construction is not guaranteed so the response characteristics of individual pixels on the same sensor are not identical. The second cause is propagation delay across the sensor. Since the

modulation signals generally enter the sensor from one point and then propagate across the sensor, a phase shift is introduced depending on the position of the pixel on the sensor (Payne et al., 2009). This error can generally be calibrated out and some manufacturer's include calibration values for this purpose. However, there is also some temperature dependence of this error and therefore the calibration may not always be valid.

2.3.4.4 TEMPORAL VARIATION IN THE MODULATION WAVEFORM

Another potential source of error is temporal changes in the modulation waveform (Godbaz et al., 2011). Depending on the hardware and measurement setup, the shape of the modulation waveform, particularly for the illumination modulation, may change during the measurement due to the discharge of capacitors and temperature build-up in the illumination source during the measurement. This is particularly noticeable when the camera is operating in a triggered acquisition mode, rather than recording frames continuously. As it is hardware specific, the exact form of error observed due to this effect is difficult to determine. Careful hardware design should be undertaken to minimise this error.

2.3.4.5 MULTIPATH ERROR

As well as the returning light coming directly from the imaged object, normally referred to as the primary return, there are also returning light paths that have reflected off two or more objects. This interfering light causes errors in the phase calculation. While this type of error is hard to quantify, as it is highly scene dependent, it is particularly apparent at acute corners of objects, where the two surfaces blend into each other in a curve rather than forming a corner between two flat planes (Guomundsson et al., 2007), and where there are particularly bright objects in the scene. Methods for eliminating multipath effects by recording measurements using multiple frequencies have been shown to work in some circumstances (Dorrington et al., 2011).

The finite spatial resolution of indirect time of flight cameras causes another multipath problem at the edges of objects. At these edges, there are pixels that represent modulated light returning from more than one object, and therefore distance. The

returning light from these objects interfere with each other and the resulting calculated phase represents neither pixel. This is called the “mixed pixel” effect.

Finally, multipath error can also occur due to scattering internal to the camera lens system. Light from an object, that does not necessarily have to be within the camera’s field of view, can reflect internally in the lens system and interfere with the measurement of another object.

2.3.4.6 MOTION ERROR

Because each range measurement requires a number of successive intensity measurements to be acquired, object motion introduces errors. These can be classified as lateral motion errors, from movement across the field of view, and axial motion errors, from movement along the viewing axis.

If an object moves laterally within the ranger’s field of view, the pixels at the edges of the object experience error. In this case the edge is not a single row of pixels around the edge of the object, as it is for mixed pixels, but is all the pixels that are on both surfaces for at least some of the measurement time. This can be arbitrarily large depending on the velocity of the object and the measurement time used. These pixels experience a step change in phase. Lottner (Lottner et al., 2007) measured the effect this had for a simulated motion of 250 mm/s, and showed that the introduction of lateral motion causes significant error compared to the static case, particularly around the edge of the object where the difference between the correct distance and the measured distance reached 100 mm.

There is currently research exploring the use of a 2D camera for edge detection and correction (Lottner et al., 2007) and optical flow algorithms (Lindner & Kolb, 2009) to ameliorate lateral motion error. The use of a 2D camera was successful for identifying edge pixels in the 3D measurements, however the methods for correcting these pixels presented in the paper are not completely satisfactory. The addition of a second camera is also not ideal. Optical flow algorithms were shown to significantly improve the measured data in a dynamic scene, however they required additional computational power and decreased the maximum frame rate.

Axial motion has a more subtle effect on range imaging than lateral motion, as the change in distance for successive frames is generally smaller. By including axial motion (2.12) becomes

$$I_n = A \cos(n\delta - \varphi_n) + B, \quad (2.24)$$

where φ_n is the actual phase at frame n (Lindner, 2010). This can be related to the phase of the previous frame using the equation

$$\begin{aligned} \varphi_n &= \varphi_{n-1} + \frac{4\pi f_{mod} v t_f}{c} \\ &= \varphi_0 + \frac{4\pi n f_{mod} v t_f}{c} \\ &= \varphi_0 + n v \alpha, \end{aligned} \quad (2.25)$$

where t_f is the time taken to measure one frame and v is the axial velocity of the object. α is introduced for convenience as

$$\alpha = \frac{4\pi f_{mod} t_f}{c}. \quad (2.26)$$

This assumes the velocity is linear over the short frame time. It has also been assumed that the amplitude is constant over the measurement time. In reality, due to the inverse square decrease in illumination with distance, the amplitude will change with distance. This will be addressed in Chapter 5.

Substituting (2.24) into (2.16) gives

$$\begin{aligned} \varphi_m &= \tan^{-1} \left(\frac{I_1 - I_3}{I_0 - I_2} \right) \\ &= \tan^{-1} \left(\frac{\cos\left(\frac{\pi}{2} - \varphi_1\right) - \cos\left(\frac{3\pi}{2} - \varphi_3\right)}{\cos(-\varphi_0) - \cos(\pi - \varphi_2)} \right) \\ &= \tan^{-1} \left(\frac{\cos\left(\frac{\pi}{2} - \varphi_0 - v\alpha\right) - \cos\left(\frac{3\pi}{2} - \varphi_0 - 3v\alpha\right)}{\cos(-\varphi_0) - \cos(\pi - \varphi_0 - 2v\alpha)} \right), \end{aligned} \quad (2.27)$$

where φ_m is the measured phase. This can then be simplified to

$$\varphi_m = \tan^{-1} \left(\frac{\sin(\varphi_0 + v\alpha) + \sin(\varphi_0 + 3v\alpha)}{\cos(\varphi_0) + \cos(\varphi_0 + 2v\alpha)} \right). \quad (2.28)$$

The theoretical error vs. distance has been plotted in Figure 2.7 for various velocities. Motion error has a non-linear relationship with distance. This error contains two components, a sinusoidal error with two cycles in the unambiguous measurement

distance, the amplitude of which is proportional to the velocity, and an offset that is proportional to velocity.

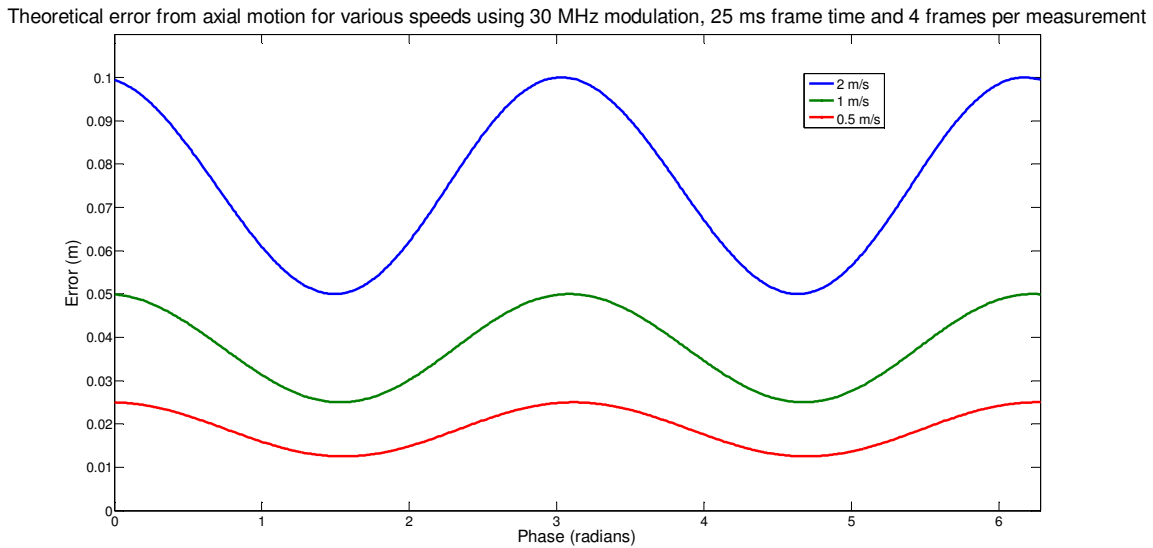


Figure 2.7 Theoretical error from axial motion for various speeds using 30 MHz modulation frequency, 25 ms integration time and 4 frames per measurement

As an object moves axially it will shrink or grow from the camera's point of view. Because of this, the edges of the object will experience similar problems to objects moving laterally.

The offset error can be shown to be caused by the selection of reference frame (Drayton et al., 2012 a). Traditionally the first frame has been used as the reference frame for the actual phase, as shown in the previous analysis; however, there is no reason to have a preferential reference frame. The measured motion error using each of the four available frames as the reference frame is shown in Figure 2.8. This demonstrates that the offset observed is due to the selection of reference frame and the actual time at which the phase would have zero offset is within the measurement.

This analysis has assumed a sinusoidal correlation waveform. Similar to linearity, harmonics have an effect on the axial motion error experienced by indirect time of flight cameras. Traditional theory from indirect time of flight literature is not sufficient to analyse axial motion error with the inclusion of harmonics. However, theory can be

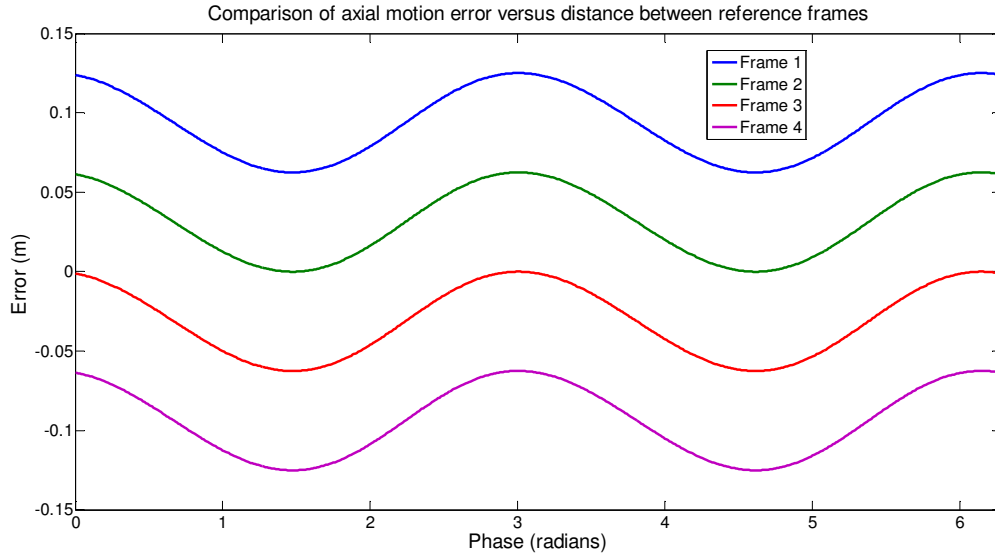


Figure 2.8 Comparison of theoretical axial motion error versus distance between reference frames

adapted from Phase Shifting Interferometry for this purpose, which is covered in detail in Chapter 5.

2.4 Commercial Systems

There are a number of commercial cameras available that can provide high quality range images. However, they have a high cost and lack the configurability required for some of the research contained in this thesis. Because of this, it was decided to develop a highly configurable custom range finding system suitable for mobile applications research, described in Chapter 3. In this section the leading commercial ranging cameras are profiled, particularly in reference to their accuracy and precision.

2.4.1 MESA IMAGING SR-4000

Mesa Imaging AG released the SwissRanger™ SR-4000 in 2008. The SR-4000 has a resolution of 176 (h) × 144 (v) pixels and comes in two models, one with a range of 0.8 m to 5 m and an accuracy of ± 10 mm and one with a range of 0.8 to 8 m and an accuracy of ± 15 mm. Accuracy is defined as being over the calibrated distance, for a target with 99% reflectivity and for the 11×11 block of centre pixels of the camera. Both models are available with a field of view of 43.6° (h) × 34.6° (v) or 69° (h) × 56° (v). A 12 VDC power supply is required. The SR-4000 costs ~\$4,300 USD (Acroname Robotics, 2012). A unique

feature of this camera is the inclusion of an optical feedback loop, which provides self-calibration of a number of parameters of the camera.

2.4.2 PMDTECH CAMCUBE

PMDTech released the PMD[vision][®] CamCube 3.0 (PMD, 2010) in June 2010. It has a range of 0.3 to 7 m, 200×200 pixels of resolution at 40 frames per second and a $40^\circ \times 40^\circ$ field of view. In its standard configuration it has physical dimensions of $60 \text{ mm} \times 194 \text{ mm} \times 60 \text{ mm}$ and is run off a single 12 VDC power supply. The PMDTech camera can be used outdoors due to its Suppression of Background Illumination (SBI) technology. The CamCube 3.0 is no longer available for sale as PMDTech has shifted focus to producing only sensors.

2.4.3 PANASONIC D-IMAGER

In February 2012 Panasonic released two new models of their D-IMager camera, a high precision model (EKL3105) and a high illumination model for outdoor environments (EKL3106). Both models have a resolution of $160 \text{ (h)} \times 120 \text{ (v)}$, a field of view of $60^\circ \text{ (v)} \times 44^\circ \text{ (h)}$ and a range of 1.2 m to 5 m under low ambient lighting conditions. They are capable of frame rates up to 30 fps. The high precision model has a precision of 2 cm under low illumination and 5 cm with an illumination of 20,000 lx, for a target with 90% reflectance, at a distance of 2 m using a central pixel. The high illumination model has a precision of 3 cm under low illumination and 14 cm with an illumination of 100,000 lx. The high precision and high illumination models cost \$1,950 USD and \$2,730 USD respectively (Digi-key, 2012). A 24 VDC power supply is required.

2.4.4 SOFTKINETIC DEPTHSense 311

In December 2011 SoftKinetic released the DS311 camera. It has a resolution of $160 \text{ (h)} \times 120 \text{ (v)}$ pixels, a range of 1.5 m to 4.5 m and a field of view of $57.3^\circ \text{ (h)} \times 42.0^\circ \text{ (v)}$. The system also has an integrated standard colour camera with a resolution of 640×480 pixels and a field of view of $50^\circ \text{ (h)} \times 40^\circ \text{ (v)}$. The camera has a quoted depth resolution of $< 3 \text{ cm}$ at 3 m, the parameters for this measurement are not

provided. Frame rates between 25 fps and 60 fps are achievable with this camera. The DS311 costs \$299 USD (Softkinetic, 2012).

2.4.5 DEFICIENCIES OF COMMERCIAL CAMERAS

While the cameras outlined in this section are capable of providing high quality range measurements, they have some common disadvantages that must be addressed. For a number of research topics having a highly configurable system is required, as being restricted to implementing high level techniques to post-process the range data outputted by these cameras rules out a number of promising methods. In particular, techniques that operate by altering the phase calculation or modulation waveform (including the phase steps, waveform shape and the modulation frequency or frequencies) have shown promise in addressing both systematic and random errors and therefore control of these parameters is highly desirable. These commercial camera manufacturers are not currently willing to provide this control to researchers or end users.

The research focus of the commercial cameras discussed in this chapter has been to provide highly linear and stable distance measurements. Some of the systematic errors outlined in this chapter have not been addressed. In particular, all of these cameras are susceptible to systematic error due to relative motion between the camera and the object being observed, as described in section 2.3.4.6. The precision of measurements using these cameras is dependent on both the pixel being used and the properties of the objects in the scene as described in section 2.3.3, meaning acquiring high quality range data over an entire field of view may not be possible.

2.5 Summary

In this chapter different range measurement techniques have been described. Indirect time of flight cameras were identified as a promising technology for providing high quality full field of view range measurements in real time. An overview was performed of error sources in indirect time of flight cameras, with theoretical explanations for observed errors. Current attempts to eliminate these errors reported in the literature were outlined.

Chapter 3 HARDWARE

3.1 Introduction

This chapter first describes the Victoria University Range Imaging System that existed at the beginning of this research. The limitations of this system are discussed, and a new revised system, developed and tested as part of this research, is illustrated. The development of an apparatus used to provide repeatable and controllable motion of a target object for experimental purposes is also presented.

3.2 The Victoria University Range Imaging System

3.2.1 *THE PROTOTYPE SYSTEM*

At the beginning of this research, a prototype indirect time of flight range imaging system existed at Victoria University of Wellington. A photograph of this system is shown in Figure 3.1. This system was mounted on an optical table measuring 600×300 mm and required a bench top power supply. The system was based around an Altera Stratix III FPGA Development board (Altera, San Jose, CA, USA) accompanied by a number of custom PCBs developed by researchers at Victoria University of Wellington and Waikato University (Jongenelen, 2010), namely;

- The Illumination Board - This board illuminates the scene being measured using two banks of laser diodes (developed by the Chronoptics group at Waikato University).
- The Image Capture Board – This board provides an interface between the camera sensor and the FPGA.
- The External Interface Board – This board provides an Ethernet and VGA interface for outputting data to external systems for long term storage or display.

3.2.2 *TOF SENSOR*

The sensor used in this system was a PMD19K-2 video sensor array from PMDTechnologies. This sensor provides 160×120 pixels of image resolution and allows




Figure 3.1 Prototype Victoria University of Wellington Range Imaging System (McClymont et al., 2010)

for external electronic modulation. The pixels are grouped into four independently modulated blocks of 40×120 pixels, each of which have two complimentary modulation inputs.

The PMD19K-2 sensor is mounted on a small daughter board, and is fitted with a Goyo Optical GM31614MCN lens (Goyo Optical, Hamasaki, Japan) with a 16 mm focal length,




Figure 3.2 Optical lens assembly photograph (left) and sensor IC (right)

manual focus and iris, and a 27 mm filter thread. A LP645 filter from Midwest Optical Systems is placed in the filter thread to suppress background illumination. This filter passes over 90% of red and NIR light while filtering shorter wavelengths. A photograph of the daughter board, with and without the lens assembly, is shown in Figure 3.2. The daughter board is connected to the Image Capture Board by two DF17C 20 way headers (Hirose, Tokyo, Japan). Bolt holes are provided to facilitate a strong physical connection between the two boards.

3.2.3 *FPGA DEVELOPMENT KIT*

An Altera EP3SL150F1152 Stratix III FPGA performed general purpose processing for the system. An FPGA has advantages over a microcontroller or DSP as it has the ability to perform many different digital logic functions, along with the required image processing, in a single package. The parallel processing capability of the FPGA allows these different systems to be implemented without compromising the timing of the sensor readout.

The Stratix family of FPGAs was chosen as it provides easily customisable phase locked loops (PLLs) which are used to generate the accurate multi-phase modulation signals. A single PLL is used to produce the modulation signals for both the illumination and the sensor. The phase of the sensor modulation signal can be stepped in real time by the FPGA with a step size of 1.50° when using the recommended multiplier setting for our frequency range. Stratix FPGAs also have a large amount of on chip memory, which is required for storing range images while they are being processed. A development kit was used to avoid the difficulty of undertaking a complicated custom FPGA PCB design. Using the development board has the drawback of a large physical size, as the development board incorporates a number of features not necessary for the ranging system. Connections to the External Interface Board and Image Capture Board were made using high speed mezzanine connectors.

3.2.4 *ILLUMINATION BOARD*

The Illumination Board provides the modulated light source required for the indirect time of flight method. Illumination is provided by 8 IR (980 nm) L980P100 (Thor Labs, Newton, NJ, USA) and 8 red (660 nm) ML101J27 (Mitsubishi, Cypress, CA, USA) laser

diodes arranged in a concentric circle around a hole in the board that accommodates the camera lens, as shown in Figure 3.3. Diffusers are used to provide homogeneous illumination over the scene with FWHM angles of $40^\circ \times 60^\circ$ (Luminit, Torrance, CA, USA).




Figure 3.3 Photograph of physical interface between the Illumination board and the sensor

The laser diodes are driven by iC-HK laser diode switches (iC Haus GmbH, Bodenheim, Germany) in a controlled current configuration providing a continuous output power of up to 50 mW per diode using a 50% duty cycle. A schematic of the laser diode switch circuit is shown in Figure 3.4. The drivers have two channels the current of which is set by the resistors R5 and R6 depending on the current required for the laser diode. For the ML101J27 diodes a $10\ \Omega$ resistor and a $5\ \Omega$ resistor are used to provide 200 mA of current at 2.7 V input on the CI pin (iC-Haus GmbH, 2011). The L980P100 diodes use resistors of $10\ \Omega$ and $3\ \Omega$ to provide the required 240 mA. The CI input is set using a PIC16F684 microcontroller (Microchip, Chandler, AZ, USA) and an AD5311 DAC (Analog Devices, Norwood, MA, USA). Because the relationship between the current through the laser diodes and the voltage across them is temperature dependent, when the system is first turned on the current through the diodes is gradually increased. This ensures that the maximum power rating of the diodes is never exceeded.

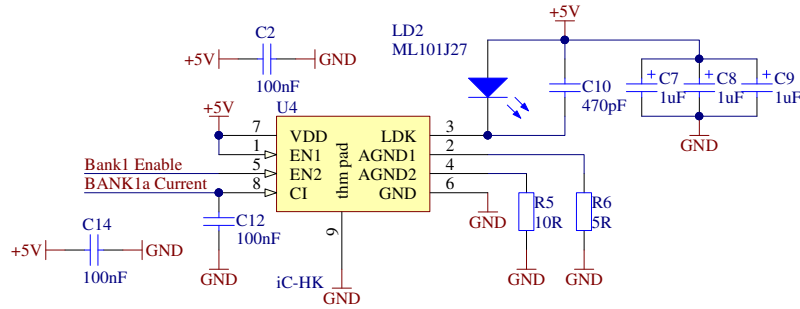


Figure 3.4 Laser diode driver schematic

3.2.5 IMAGE CAPTURE BOARD

The image capture board provides an interface to the PMD daughter board that contains the PMD 19K-2 sensor. It also performs modulation of the sensor using two EL7158 Ultra-High Current Pin Drivers (Intersil, Milpitas, CA, USA) shown in Figure 3.5, each connected to four of the modulation inputs. The voltage is switched between 0 V and a high voltage set using an AD8531 as a voltage follower. The voltage is set to 2 V, as recommended in the PMD19K-2 sensor datasheet (PMD Technologies, 2008), using a voltage divider formed by two resistors (R4 and R7). The output of the switches is connected to the modulation inputs via the $3\ \Omega$ resistors R1 and R2. Each modulation input has a capacitance of 250 pF meaning each switch is driving 1 nF. The drivers have an internal on-resistance of $0.5\ \Omega$. Combined with the $3\ \Omega$ resistors, this forms a low pass filter with a cut off frequency of 45 MHz. It is desirable to limit the bandwidth of the system to the frequencies that are intended to be used, as this helps to mitigate the linearity errors due to harmonics discussed in Chapter 2.

The LVDS modulation signal from the FPGA board is connected to two inputs on a LVDS receiver, one with reverse polarity, as shown in Figure 3.6. This generates the two complimentary modulation signals MODA_I and MODB_I.

The Image Capture Board also has an AD9826 16-bit two channel ADC (Analog Devices, Norwood, MA, USA) that converts the analogue video stream from both readout gates of the PMD19K-2 sensor into digital video frames which are then sent to the FPGA for processing. The gain and offset of the ADC can be adjusted by the FPGA via a three wire SPI interface. The Image Capture Board is connected to the FPGA development board via a High Speed Mezzanine connection.

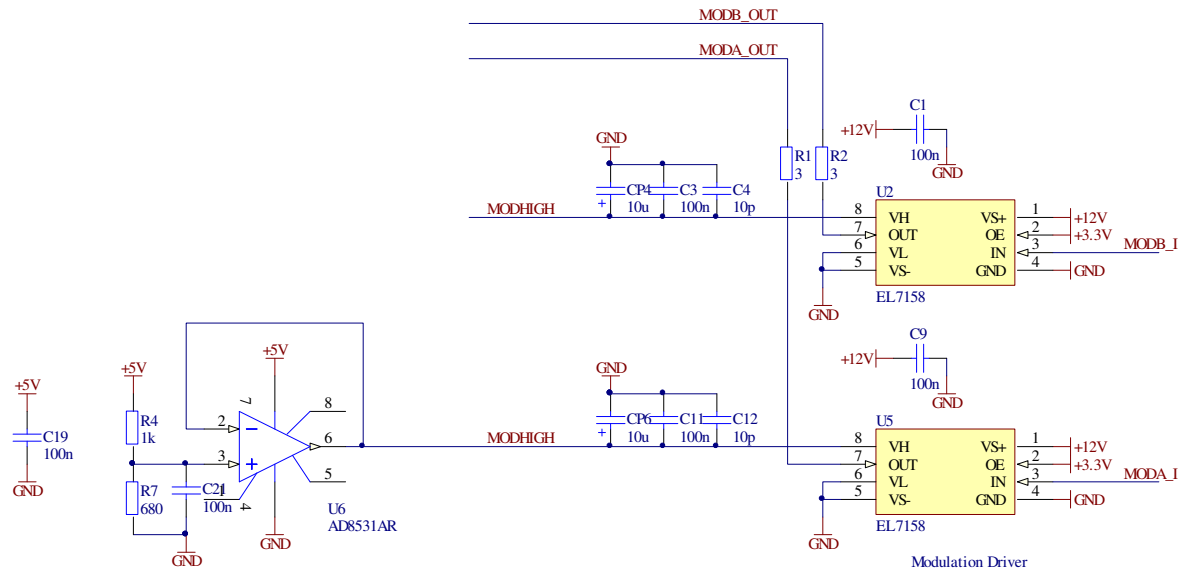


Figure 3.5 PMD19K-2 modulation driver circuit schematic

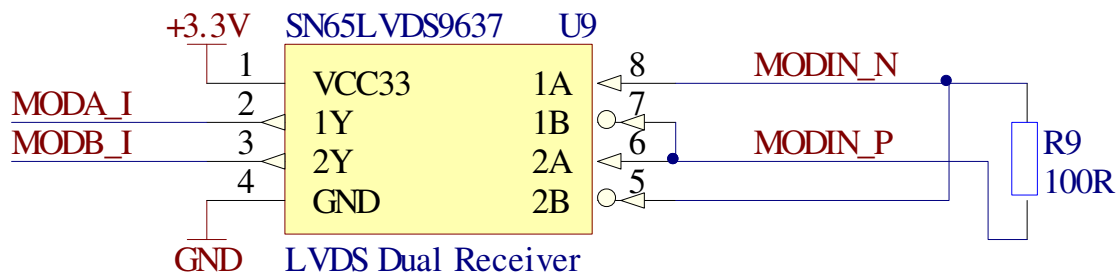


Figure 3.6 Schematic of modulation signal generation by a LVDS receiver

3.2.6 EXTERNAL INTERFACE BOARD

The External Interface Board allows for observation and long term storage of the data collected by the system. A VGA connection, driven by an ADV7123 Triple 10-Bit High Speed Video DAC (Analog Devices, Norwood, MA, USA) is used to output data to a 640×480 VGA monitor. The VGA interface is controlled by the FPGA. Four different data streams can be displayed on the monitor at any time. These data are configurable, but are usually the raw data values from the two readout gates of the sensor, the amplitude of the correlation waveform and the distance data. These are the data outputs that are likely to be useful for robotic navigation or measurement purposes.

A DM9000 Ethernet Controller (Davicom, Hsinchu, Taiwan) is used to provide an interface for recording data onto a computer for long term storage. A NIOS II processor

implemented on the FPGA provides a driver for this controller. The NIOS II processor also controls the dataflow of frames to the Ethernet interface, allowing the frame data to be collected and saved on an external computer using a Java application.

3.2.7 PROBLEMS WITH THE PROTOTYPE SYSTEM

Because it is based around an FPGA development board, the physical dimensions of this system are large ($400 \times 200 \times 250$ mm excluding the optical table it is mounted on). These dimensions preclude this system from being implemented on most mobile robotic platforms. The FPGA development board also has a high replacement cost if the system is damaged (\$3500). Two separate voltage supplies are required for the system, a 12 V regulated supply for the Stratix III development board and a 10 V regulated supply for the Illumination board. This was implemented using a laboratory bench top power supply.

The EL7158 modulation drivers used on the Image Capture Board limit the modulation frequency to a maximum of 40 MHz. Each modulation block represents a capacitive load of 250 pF that should be driven to 2 V to provide maximum contrast. For a frequency of 40 MHz modulating each block requires 125 mA, calculated using a perfect capacitor model. Each EL7158 is driving four modulation blocks and has a continuous current limit of 500 mA. Therefore, the number of modulation drivers should be increased to allow higher modulation frequencies and allow some safety headroom for the current modulation frequency range.

3.3 Development of a compact Range Imaging System

Due to the problems with the existing Victoria University of Wellington Range Imaging System, a new compact system was designed. The work on this system was completed concurrently with Johnny McClymont as part of his Master's thesis (McClymont, 2010). The hardware portion of this thesis focused on the PCB design of the Image Capture Board, routing the DDR2 memory on the FPGA board, and testing, debugging and further modification of the system. It also included alteration of the FPGA firmware. Overall design decisions were made as a team. A photograph of the final system is shown Figure 3.7.




Figure 3.7 Photograph of the compact Victoria University Range Imaging System (McClymont et al., 2010)

3.3.1 DESIGN METHODOLOGY

It is highly desirable for the new system to be easily configurable in order for it to serve as a useful research platform. Therefore a modular design was chosen using the same basic board structure as the prototype system to split the design into separate subsystems. The boards used in this system are an FPGA board, an Image Capture Board, an Illumination Board and an External Interface board. Figure 3.8 shows the physical connections between each of the boards.

The FPGA board is connected to the Image Capture and External Interface boards using rigid 172 pin High Speed Mezzanine board to board connectors. These connectors provide low voltage differential signal connections capable of transmissions up to 8.5 GHz. They also form a stable physical connection between the boards, in combination with spacers in the corners of the boards.

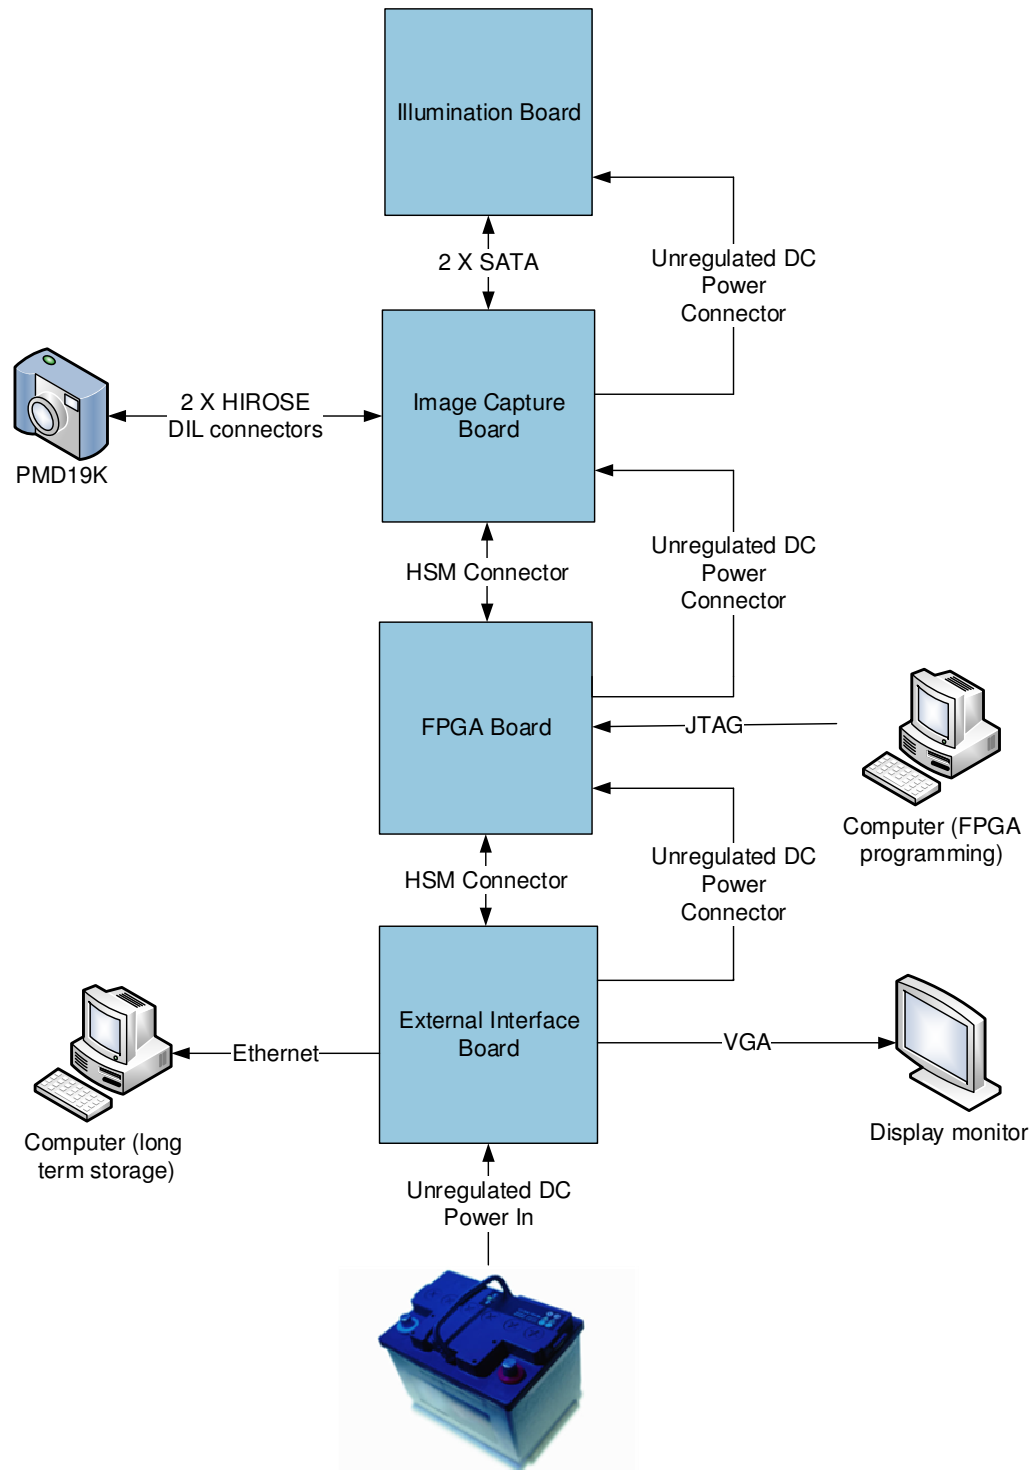


Figure 3.8 Physical connections diagram of the compact range imaging system

Modulation and current setting signals are sent from the Image Capture Board to the Illumination board using two SATA connections. SATA connections are used as the flexibility of SATA cables allows the distance between these two boards to be readily adjusted to accommodate different optical lenses.

Power distribution layout is an important factor in this design. As shown in Figure 3.8, a single unregulated DC power input is provided on the External Interface Board. From the External Interface board power is distributed to the other boards in the system via unregulated DC power plugs. It is then regulated individually on each board. This treats each board as an individual sub system, providing versatile modularity. Having a single external power connection simplifies interfacing the system with a mobile robot. Implementing individual regulators on each board also helps to distribute power dissipation and heat build up, and reduces cross-talk between subsystems.

The flow of data between the boards is shown in Figure 3.9. The FPGA board provides the control for the flow of data. It sends both modulation signals and signals to set the laser diode current through the Image Capture board to the Illumination board. Due to the stacked nature of the boards, a direct connection from the FPGA board to the Illumination board would require an additional, and unnecessary, cable. The Image Capture board also takes PMD control signals from the FPGA board to pass through to the PMD19K-2 sensor on the daughter board and returns the video signal from the sensor, after converting it to digital frames. The FPGA board then sends the processed data to external systems via the External Interface Board. FPGA configuration data is transmitted via JTAG to the FPGA board. JTAG is also used to interface with the NIOS II processor on the FPGA. The functions of these boards are discussed in more detail later in this chapter.

To reduce the size of the system, making it more suitable for mobile applications, the Stratix III FPGA development board used for the initial system is replaced with a custom FPGA board, as described in section 3.3.4.

While existing sensors have relatively low resolutions, it is expected that 1 MegaPixel sensors will be available in the near future. To allow for future-proofing, the system design, particularly memory buffer sizes and bus speeds, must be capable of handling this increased resolution. This requires the implementation of several external memory modules described in section 3.3.4.




Figure 3.9 Compact range imaging system board data flow diagram

3.3.2 IMAGE CAPTURE BOARD

Figure 3.10 shows a functional diagram of the Image Capture board. It was initially envisaged that the limitation on the modulation frequency described in section 3.2.7 could be overcome by increasing the number of EL7158 Ultra-High current drivers from 2 to 8, one for each modulation pin on the PMD 19K-2 sensor. However, this limitation is also due to the input bandwidth of the drivers, as well as the output drive current, so

the desired improvement was not realised. The sensor itself also has bandwidth limitations that limit any potential increase in modulation frequency.

An AD9826 ADC is used to convert the two analog video streams from the sensor's two output gates into digital frames for the FPGA.

Low Voltage Differential Signalling (LVDS) drivers and receivers are used to interface between the boards due to their high noise immunity, low power consumption and high speed. The LVDS repeaters are used to transfer signals from the FPGA board to the Illumination board. From the prototype version, these repeaters have been extended to include the current setting signals from the FPGA board.

The PCB layout for the Image Capture Board is shown in Appendix A.1.

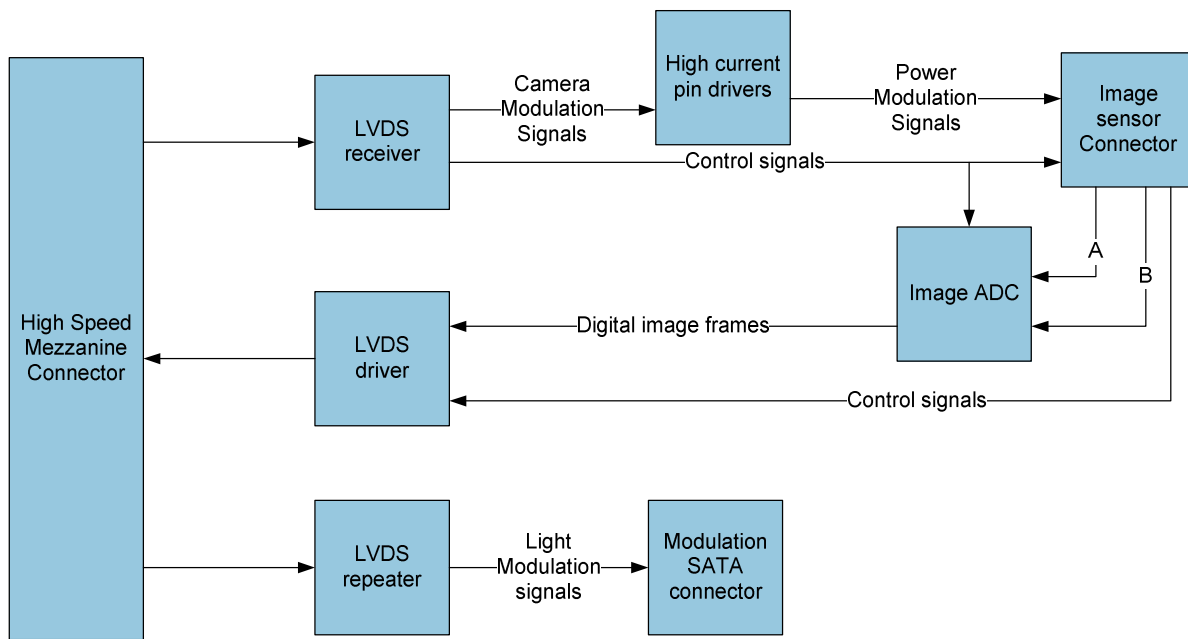


Figure 3.10 Functional diagram of the Image Capture board

3.3.3 ILLUMINATION BOARD

The microcontroller on the Illumination board was removed and the current of the laser diodes is instead set directly using the FPGA. This increases the flexibility of the system, as the light intensity can be adjusted in real time by the control systems implemented on the FPGA, which has information on the intensity of objects in the field of view that the

microcontroller did not. The signals to set the ADC value are sent via a second SATA connection, as each SATA connection only contains two differential signal wire pairs.

For this version, it was decided to replace the IR laser diodes with red diodes. The different diodes were used in the prototype to allow research into the effect that different light wavelengths had on the measurements. This feature is no longer necessary. Red laser diodes are preferable for safety as they are visible to humans and therefore activate blink reflexes, lowering the chance of eye damage. The PMD19K sensor also has higher contrast for red light than IR (PMD Technologies, 2008). The resistor values have been adjusted to operate using the new laser diodes and to operate at a maximum current control (CI) voltage of 5V. A functional diagram of the Illumination board is shown in Figure 3.11 and an updated schematic of the pin driver is shown in Figure 3.12.

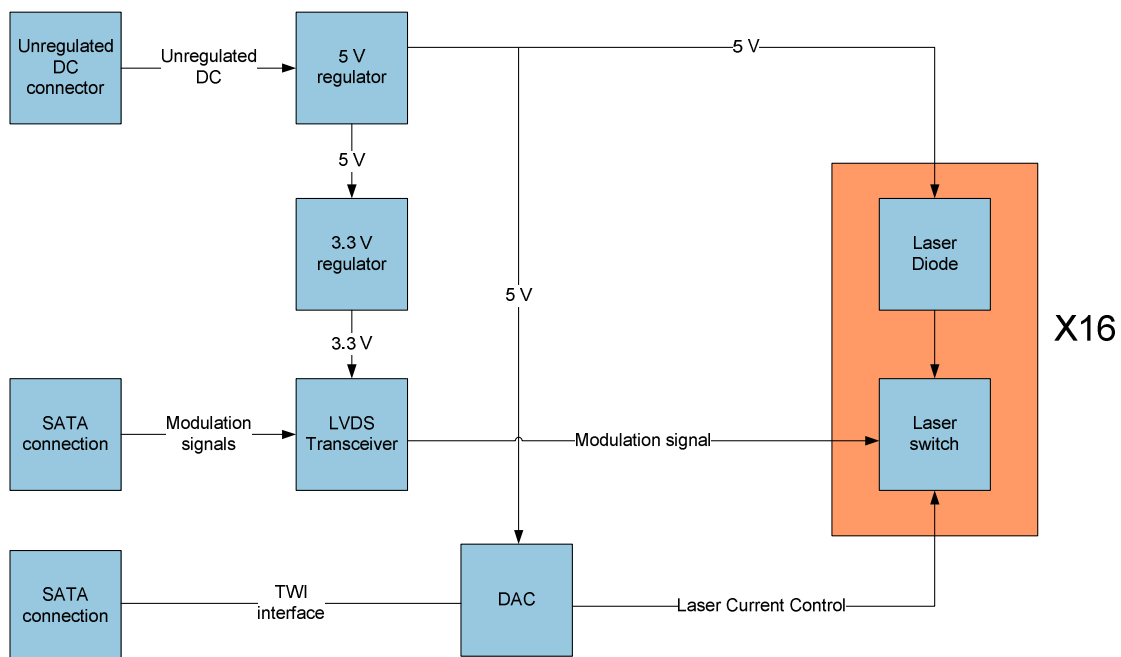


Figure 3.11 Functional diagram of the Illumination Board

BANK1_control is an added safety feature, required as the current is now being set on the FPGA board. It is generated by low pass filtering the modulation signal and comparing this to a set voltage. This means that if the duty cycle drops below a threshold level (for example if the SATA cable becomes unplugged) the laser diodes will shut down.

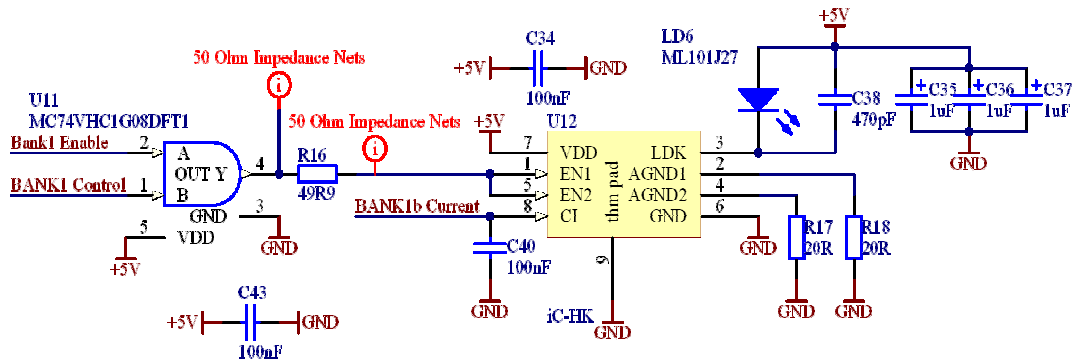


Figure 3.12 Updated schematic of the laser diode driver

3.3.4 FPGA BOARD

With the development of a custom FPGA board it was decided to replace the Stratix III FPGA with a Cyclone III FPGA, as the Cyclone is an order of magnitude less expensive (see Table 3.1 for a comparison of relevant FPGA devices). Only two PLLs are required for the range imaging system, therefore the lower number of PLLs in the Cyclone is not a limiting factor. The use of a custom board also means we can implement external memory systems suitable for our application, so the larger number of internal memory elements provided by the Stratix III are also no longer necessary.

Table 3.1 Comparison of FPGA devices (Altera, 2011)

FPGA	# of logic elements	# of RAM bits	# of PLLs	Cost (USD)
Stratix III EP3SL150F1152C2N	142,500	6,543,360	12	\$3777.00
Cyclone III EP3C120F780C7	119,088	3,981,312	4	\$502.00
Cyclone III EP3C40F780C6	39,600	1,161,216	4	\$166.50

The new FPGA board is compatible with both the EP3C40 and EP3C120 Cyclone III FPGAs. The latter is more expensive, but has approximately three times as many logic and memory elements as the EP3C40 (see Table 3.1). The EP3C40 FPGA was chosen for this version as it provides the lowest cost option.

To future proof the system, external memory chips are included on the FPGA board to allow the system to work with sensors with resolutions up to 1 MegaPixel. Each pixel is stored as a 16 bit number. This means a 1 MegaPixel camera will require 16 Mb per frame. Six frames are stored simultaneously during standard operation for processing, these are the sum of the imaginary components and the sum of the real components (used for the calculation of the phase), the current readout gate a image, and the current readout gate b image, the calculated phase image and the calculated amplitude image.

DDR2 memory was chosen for these memory buffers. This was primarily due to its high bandwidth and low cost per megabyte of storage. The Cyclone III also provides predesigned DDR2 memory interfaces simplifying the task of implementing the memory in firmware. Four DDR2 memory buffers are required for this system:

- Accumulator – this memory bank is used to store the images during processing. It stores both semi-processed and fully processed data used to determine the phase measurements. This memory bank is implemented using a 512 Mb MT47H64M8 DDR2 memory chip (Micron, Boise, ID, USA).
- Output Buffer – this memory bank stores finished images that will be accessed by the NIOS II processor, the VGA output and the Ethernet output. This bank is also implemented on a 512 Mb MT47H64M8 DDR2 memory chip (Micron, Boise, ID, USA).
- NIOS Buffer – this memory bank stores the NIOS II processor firmware. The NIOS II processor is used to control the VGA and Ethernet interfaces, as well as providing an interface to change the modulation and frame capture options. This memory bank is implemented using a 512 Mb MT47H32M16BT DDR2 memory chip (Micron, Boise, ID, USA). A 16 bit memory chip is used for this memory bank as the NIOS processor requires very high bandwidth.
- Ethernet Buffer – this memory bank stores the image frames that will be transferred via the Ethernet interface to a computer for long term storage. It buffers against the latency of the Ethernet connection. This memory bank is

implemented on two 4 Gb MT47H512M8 DDR2 memory chips (Micron, Boise, ID, USA).

The FPGA board is an eight layer board consisting of four data layers and four power plane layers. An eight layer board was used primarily due to the high density of pads under the FPGA requiring multiple breakout layers. The four data layers of the FPGA board are shown in Appendix A.1.

Because of the transfer speeds of the memory chips, the data and address PCB tracks for the DDR2 memory chips have to be length tuned to ensure signals arrive coincidentally.

As well as the DDR2 memory, a RC28F256P30BF parallel Flash memory chip (Micron, Boise, ID, USA) is included on the FPGA board. This chip contains the configuration for the FPGA. On start up the FPGA configuration is transferred from the Flash chip to the FPGA.

3.3.5 THE EXTERNAL INTERFACE BOARD

An ATMEGA32U4 microcontroller was added to the external interface board to provide a USB interface. This interface can be used to record data directly to a computer for long

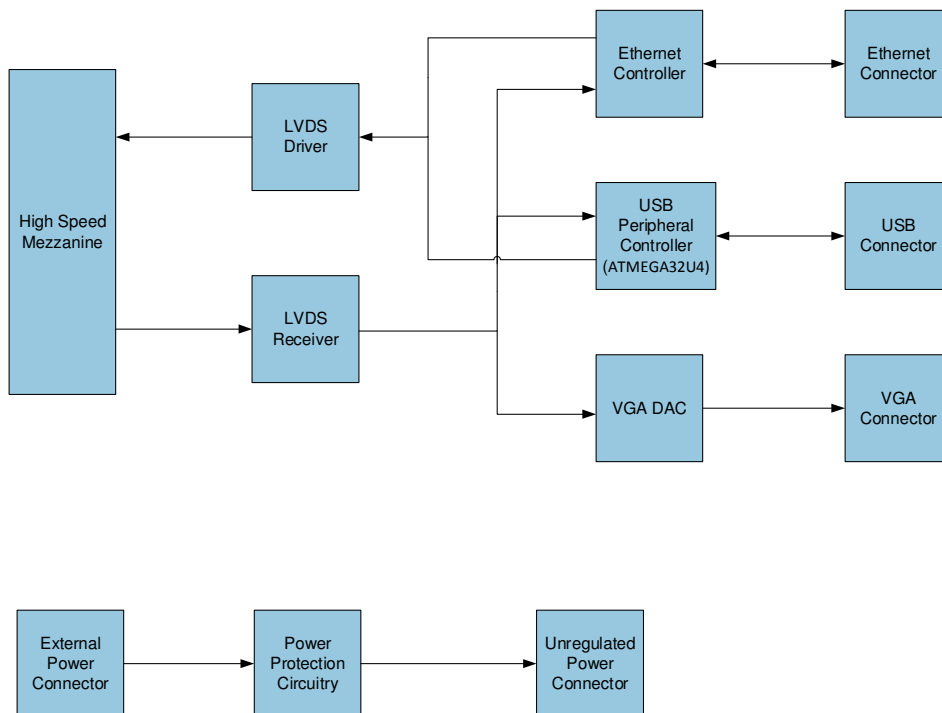


Figure 3.13 Functional diagram of the External Interface Board

term storage, providing a second interface option along with the Ethernet connection. The VGA connection remains unchanged.

In the compact system, the External interface board is used to connect to a single unregulated power supply that is used to power the entire system. A 48CTQ060SPBF rectifier diode is used to ensure the polarity of the power is correct and an LC low pass filter is used to remove high frequency noise from the input. This power is then connected to an unregulated DC power connector, daisy-chaining the system power to the other boards. A functional diagram of the External Interface Board is shown in Figure 3.13.

3.3.6 *FIRMWARE*

Because the Stratix III FPGA was replaced with a Cyclone III FPGA, significant firmware changes were required, particularly in how memory is interfaced. As previously discussed, the lower number of internal memory elements means that external memory must be implemented. DDR2 memory was selected (as discussed in section 3.3.4), which differs from the FPGA internal memory as data is transferred in blocks rather than having single line access. For the NIOS II processor the NIOS II system itself handles the transfer of data blocks to and from external memory without the need for additional code.

To port the previous FPGA code onto the new board a new pin map was required, in order to map the internal signals to the external peripherals. As discussed in section 3.2.4, in the previous system a PIC16F684 microcontroller (Microchip, Chandler, AZ, USA) on the illumination board was used to set the current of the laser diodes. In the compact version this microcontroller was removed and the current is set directly by the FPGA. This is done using a two wire interface (TWI) that sets the voltage of an AD5311 DAC. The timing diagram for this interface is shown in Figure 3.14. TWI was chosen as it provides a low cost way of setting the current, requiring only two conductors that can be implemented using a common SATA cable. The data rate is not important as the current is updated infrequently in a manner that is not time critical. To start and stop a write




Figure 3.14 AD5311 timing diagram (Analog Devices, 2010)

cycle, the SDA line must be pulled low and high respectively when SCL is high. All other transitions must occur while SCL is low.

A state machine is used to program the AD5311 DAC, the state diagram of which is shown in Figure 3.15. Count is a counter running off the 50 MHz clock and place is an index keeping track of how many data bits have been sent. SCL is generated by using a 10 bit counter to divide the 50 MHz FPGA clock down to 48.8 kHz. Transitions on the SDA line must occur while SCL is high and must remain stable during the high period of SCL. The state machine runs at twice the frequency of SCL and transitions occur at 90 and 270 degrees out of phase with SCL, allowing transitions to occur between positive levels of SCL. From start up the current is slowly increased to a maximum through successive writes to the DAC. This is required to ensure safe operation of the laser diodes as explained in section 3.2.4. The current starts at 10% of the final value and is increased to full current over a period of four minutes.




Figure 3.15 State machine diagram for setting the laser diode current

3.4 Revision of the Compact Range Imaging System

A number of issues were identified with the initial design of the compact range imaging system. Most of these were small layout errors, however there were also significant design changes required. During testing it became evident that several additional features were required, so it was decided to develop a second revision of the compact system. This second revision allowed all the layout issues to be fixed, design changes to be implemented, and additional features to be added. This section will detail the second revision of the compact Victoria University of Wellington Range Imaging System board by board, describing issues with each board and changes made in the second revision.

3.4.1 CHANGES TO THE ILLUMINATION BOARD

On the Illumination board some decoupling capacitors for the laser diodes were placed too close to the laser diodes. This meant that the metal casings of the laser diodes could not sit flush with the board without creating a short circuit. To fix this, the capacitors were spaced out. Due to supply issues and cost, as shown in Table 3.2, it was decided to replace the ML101J27 laser diodes with HL6545MG (Opnext, Fremont, CA, USA) laser diodes. These have the same wavelength of 660 nm and a maximum continuous output power of 120 mW. The footprint was changed accordingly.

Table 3.2 Comparison of laser diode suppliers

Diode	Supplier	Cost	Minimum order quantity
ML101J27	Thorlabs	\$120 USD	1
ML101J27	Orthotoronto	\$65 USD	50
HL6545MG	Thorlabs	\$65 USD	1

3.4.2 CHANGES TO THE IMAGE CAPTURE BOARD

There was a significant issue with heat dissipation on the Image Capture board. The high current drivers produce large amounts of heat (proportional to the modulation frequency) as shown in Figure 3.16. These data were taken using a Fluke 566 IR thermometer (Fluke, Everett, WA, USA) measuring the temperature of the 2.5 V regulator on the Image Capture Board. The ambient temperature was 21.1 °C.

The overheating problem was so severe that when modulated at 40 MHz the 2.5 V regulator would periodically go into thermal shutdown, as the temperature exceeded 85 °C. The short term solution to this was to provide active cooling via a 12 V computer fan attached to the bottom of the system. A K type thermocouple, attached to one of the high current driver chips using thermal paste, was used to measure the temperature over time with and without the fan to show its effect. Without the fan the temperature quickly reached the thermal shutdown level. However, with the fan the temperature stabilises at ~30 °C, well below this temperature. These data are shown in Figure 3.17.




Figure 3.16 Image capture board temperature vs. modulation frequency




Figure 3.17 Image Capture Board temperature with and without fan (McClymont, 2010)

In the first revision, the Ultra-High current pin drivers were all located close together. On the new board revision, the Ultra-High current pin drivers were spread out over the board to provide greater thermal dissipation. They were also moved from the FPGA Board facing side of the Image Capture Board to the side facing the Illumination Board. This side has greater vertical clearance due to the space required for the optical lens, making it easier to attach heat sinks to the pin drivers and improving natural air flow around the devices. A comparison between the layouts is shown in Figure 3.18.

For high frequency modulation applications, an LM64 fan controller (National Semiconductor, Santa Clara, CA, USA) with a MMBT3904 diode connected transistor (Fairchild Semiconductor, San Jose, CA, USA) as a temperature sensor was implemented to provide active fan control. Fan speed is set by a PWM output driving a MMBT3904 transistor, if the PWM is not enabled the transistor gate is pulled high by R24, making the fan run at full power. A two wire interface allows the FPGA to set the fan speed




Figure 3.18 Comparison of modulation driver layouts between the original compact system (top) and the revised compact system (bottom). The drivers have been split into pairs and spread around the board

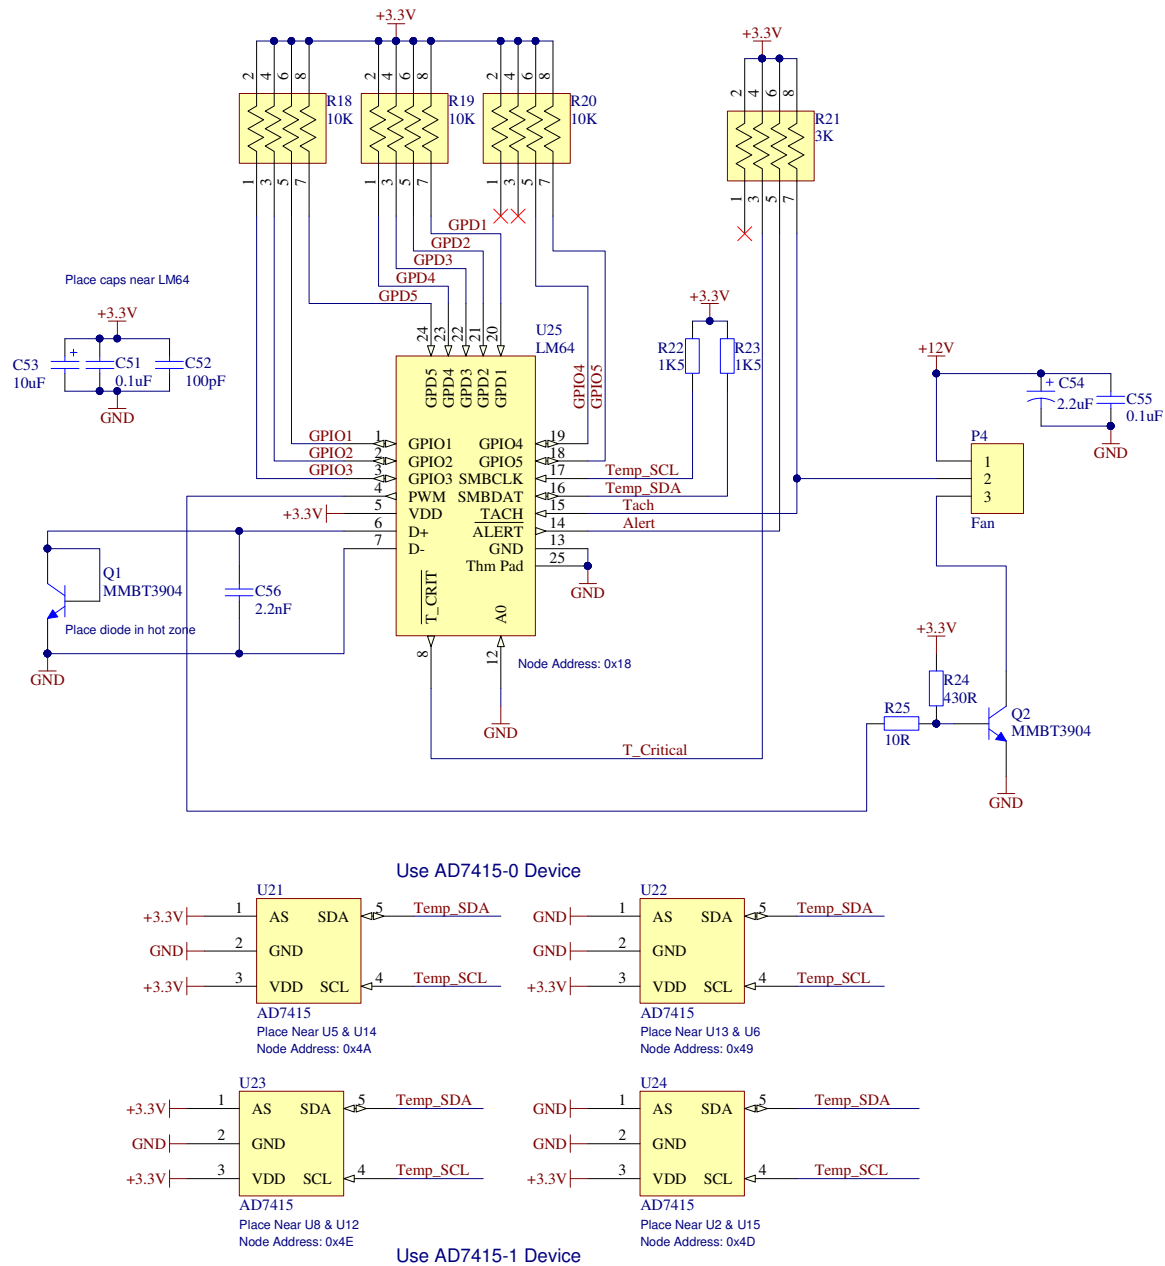


Figure 3.19 Schematic of the temperature control system implemented in the revised compact system directly if required. Four AD7415 band gap digital temperature sensors (Analog Devices, Norwood, MA, USA) spaced around the board near the pin drivers allow the FPGA to monitor the temperature of the board (U21-24 in Appendix A.2). This circuit is shown in Figure 3.19.

Development of the Linear Actuation System described in section 3.6 was undertaken in parallel with the development of the compact range imaging system. To allow these two systems to interact, an additional interface was required on the range imaging system. A 4 pin generic IO connector was added to the image capture board to allow connection to

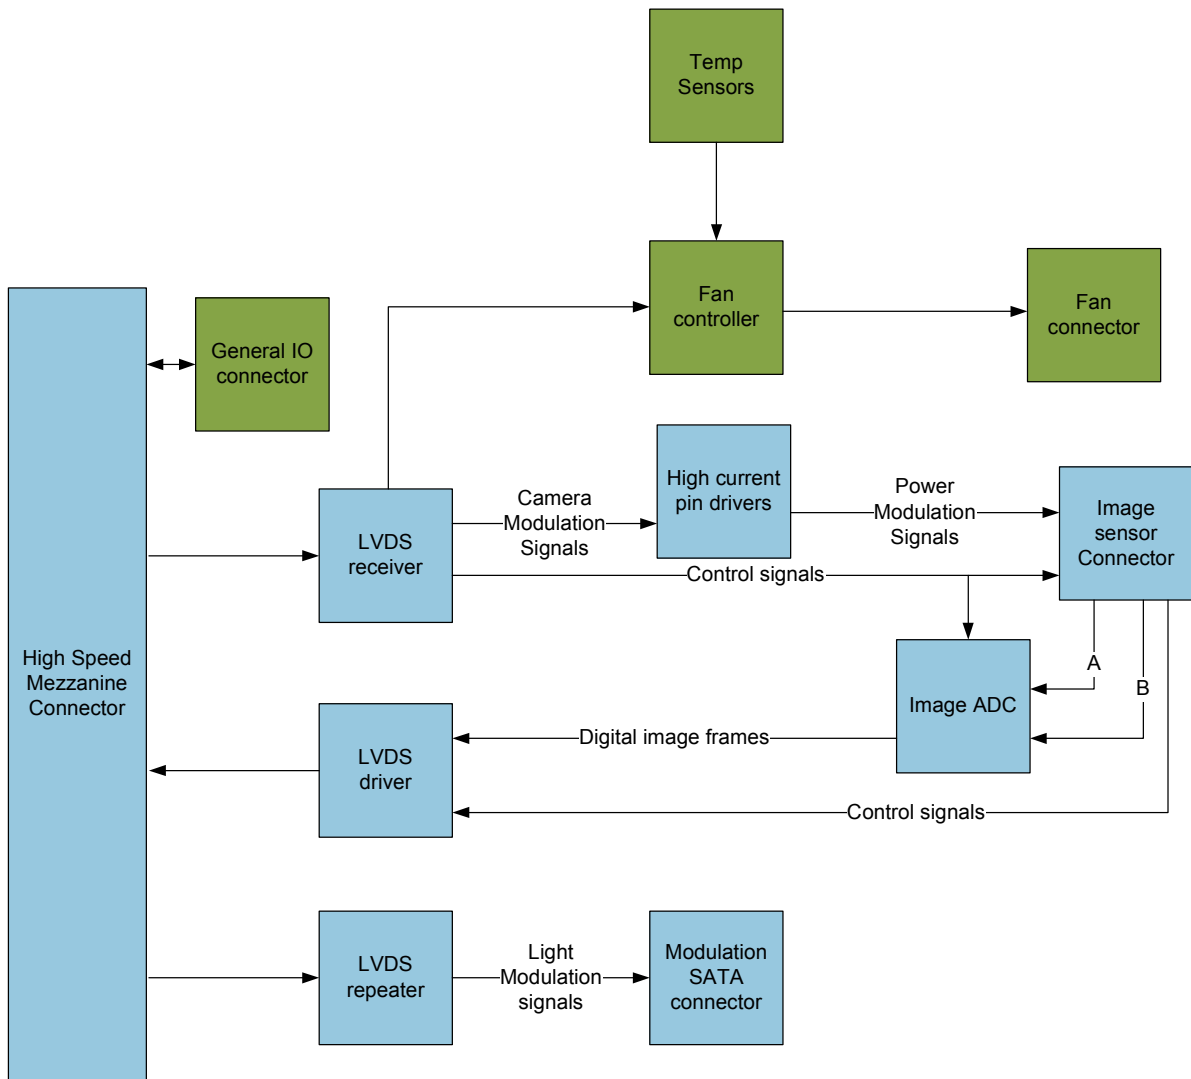


Figure 3.20 Functional diagram of revision 2.1 of the Image Capture Board

the Linear Actuation System and to provide the potential to connect to other devices in the future. An updated functional diagram is shown in Figure 3.20.

3.4.3 CHANGES TO THE FPGA BOARD

For our sensor's current resolution of 160×120 pixels, each frame requires 300 kb of RAM. Six frames are stored in buffers at any time hence requiring 1800 kb of RAM. It is therefore possible to implement this design using the internal memory of the Cyclone III EP3C120F780C7, which has 3888 kb of RAM. This simplifies the initial design and testing process significantly and therefore for the second revision of this system the EP3C120 was chosen over the EP3C40.

External memory is required for the NIOS II processor and is still desirable to allow higher resolution sensors to be used in the future. The memory banks on the FPGA board had several issues. The 16-bit DDR2 memory banks (Ethernet buffer and NIOS buffer) were not connected correctly, as the FPGA DDR2 interface requires the data lines be connected to specific high speed pins. This was fixed by changing the pin assignment and rerouting the board.

The Flash memory did not function correctly. This was because it requires connection to pins in I/O banks 1,6,7 and 8. Banks 7 and 8 are required by the Cyclone III architecture for DDR2 SDRAM and therefore operate off 1.8 V while banks 1 and 6 are used for the mezzanine connectors and require 3.3 V. Because of these specific pin requirements, successful implementation of the parallel flash would require a complete redesign of the FPGA board. As the use of a parallel flash does not provide any significant benefits, it was decided to replace the parallel flash memory chip with a serial flash memory chip, EPCS64SI16N (Altera, San Jose, CA, USA), so a redesign of the entire board was not required.

Table 3.3 shows all the voltage regulators used in the compact range imaging system. The acceptable input voltage range of the first revision was between 14 V and 20 V. This was because the LTM4601 regulators used on the FPGA and Illumination boards have a maximum voltage of 20 V and the LT3481 regulator used to provide a 12 V supply on the FPGA board has a minimum voltage of 14 V. To improve this so the system could be run from a single lead acid battery, as is common in mobile robotics, the LT3481 regulator was replaced with a LT3757 regulator, which has an input range of 5.5 V to 36 V. This makes the allowable input voltage range 10 V to 20 V (the minimum voltage is due to the LTM4601 5 V regulator on the Illumination board). The voltage range could be improved further by replacing the other regulators with SEPIC regulators, however, SEPIC regulators require more components than the other regulators used in this design and the current voltage range is suitable for most applications.

Table 3.3 Voltage regulators for the compact range imaging system

Regulator	Board	Regulator Type	Input voltage	Current	Output voltage	Input Range
		Step down switch				
LT3481	FPGA (rev 2)	mode regulator	UNREG DC	2 A	12 V	14 V - 34 V
LT3757	FPGA (rev 2.1)	SEPIC	UNREG DC	2 A	12 V	5.5 V - 36 V
LT1761	FPGA	Linear regulator	12 V	100 mA	5 V	NA
		Step down switch				
LTM4601	FPGA	mode regulator	UNREG DC	12 A	3.3 V	4.5 V - 20 V
LTC3026	FPGA	Linear regulator	3.3 V	1.5 A	2.5 V	NA
		Step down switch				
LTC3418	FPGA	mode regulator	3.3 V	8 A	1.8 V	NA
		Step down switch				
LTC3418	FPGA	mode regulator	3.3 V	8 A	1.2 V	NA
LTC3026	FPGA	Linear regulator	3.3 V	1.5 A	1.2 V	NA
TPS51100	FPGA	Linear Regulator	5 V	3 A	0.9 V	NA
		Step down switch				
LT3481	Image Capture	mode regulator	UNREG DC	2 A	5 V	6 V - 34 V
TPS79625	Image Capture	Linear regulator	5 V	1 A	2.5 V	NA
		Step down switch				
LTM4601	Illumination	mode regulator	UNREGDC	12 A	5 V	10 V - 20 V
ADP3300ART-3.3	Illumination	Linear regulator	5 V	50 mA	3.3 V	NA
	External	Step down switch				
LT3481	Interface	mode regulator	UNREGDC	2 A	5 V	6 V - 34 V
	External					
ADP1715ARMZ-3.3	Interface	Linear regulator	5 V	500 mA	3.3 V	NA
	External					
TPS79625	Interface	Linear regulator	5 V	1 A	2.5 V	A

3.4.4 CHANGES TO THE EXTERNAL INTERFACE BOARD

Only minor changes were needed to the External Interface board. The VGA interface IC was changed to operate from 5 V instead of 3.3 V, as this made it more stable and compatible with a larger variety of monitor models. A power switch was added that allows the control of power throughout the system. This utilised the unused pin on the unregulated power connector, sending a signal that can be used to enable and disable the voltage regulators on the other boards. If this feature is not desired a jumper can be used to enable the boards individually. This allows us to power or unpower individual boards as desired, which is useful for testing purposes.

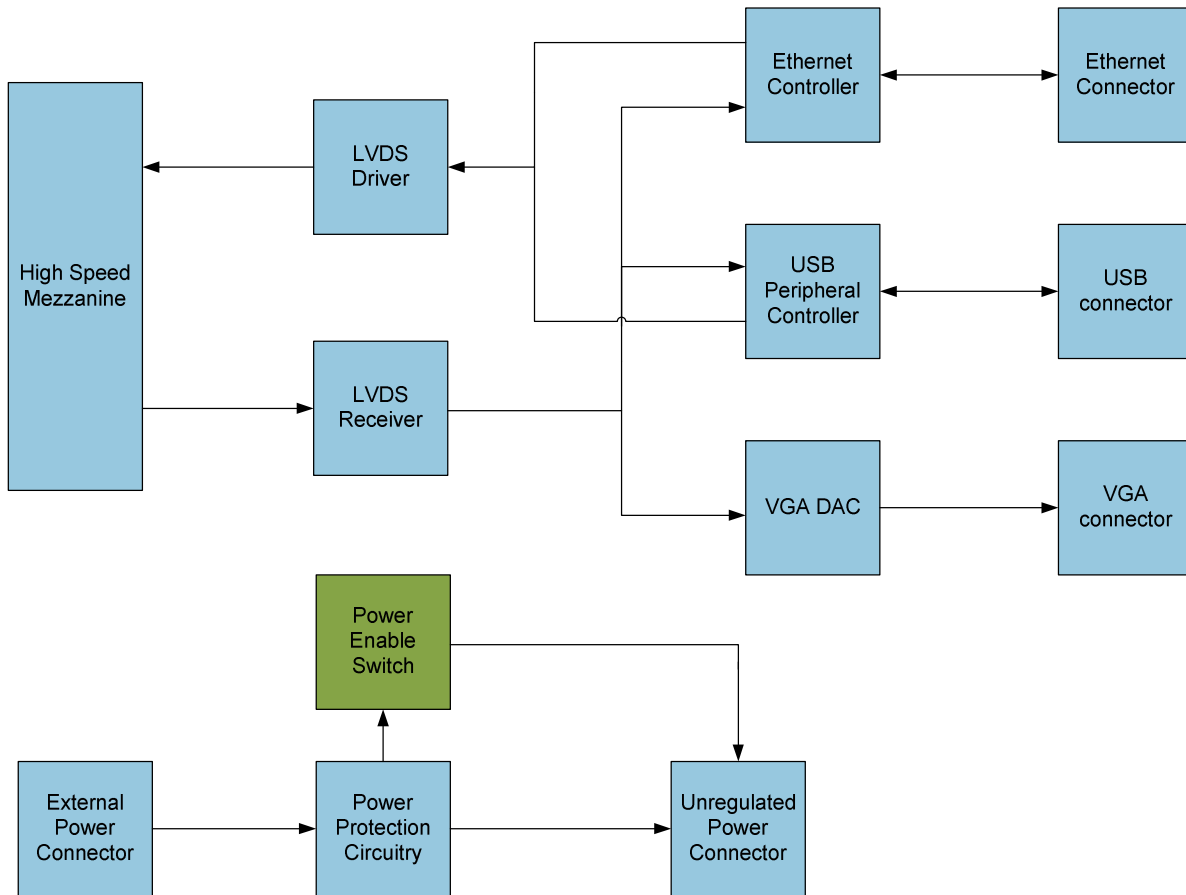


Figure 3.21 Functional diagram of revision 2.1 of the External Interface Board

Due to a supply shortage it was not possible to acquire the ATMEGA32U4 IC (Atmel, San Jose, CA, USA) intended to provide a USB interface in the first revision. This resulted in the first revision not having a USB interface. For the second revision the ATMEGA32U4 was replaced with an ATMEGA16U4 (Atmel, San Jose, CA, USA). This provides the same functionality but is in stock with providers. The footprint was changed accordingly. An updated functional diagram of the External Interface Board is shown in Figure 3.21.

3.5 Operation of the system

A diagram demonstrating the operation of the system is shown in Figure 3.22. Firstly the sensor is integrated for the set integration time t_i and is then read out. This measured intensity value is then multiplied by sine and cosine values calculated in accordance with (2.13). For the four frame algorithm these values simplify as shown in section 2.2.3.2, however, to allow the system to have maximal configurability, the more general solution is implemented. These values are then either added to the value stored in the associated accumulator or, if it is the first measurement ($n = 0$), remain unaltered and are stored in

the associated accumulator. Once all the measurements have been acquired, the phase and amplitude values are calculated and written to output buffers.

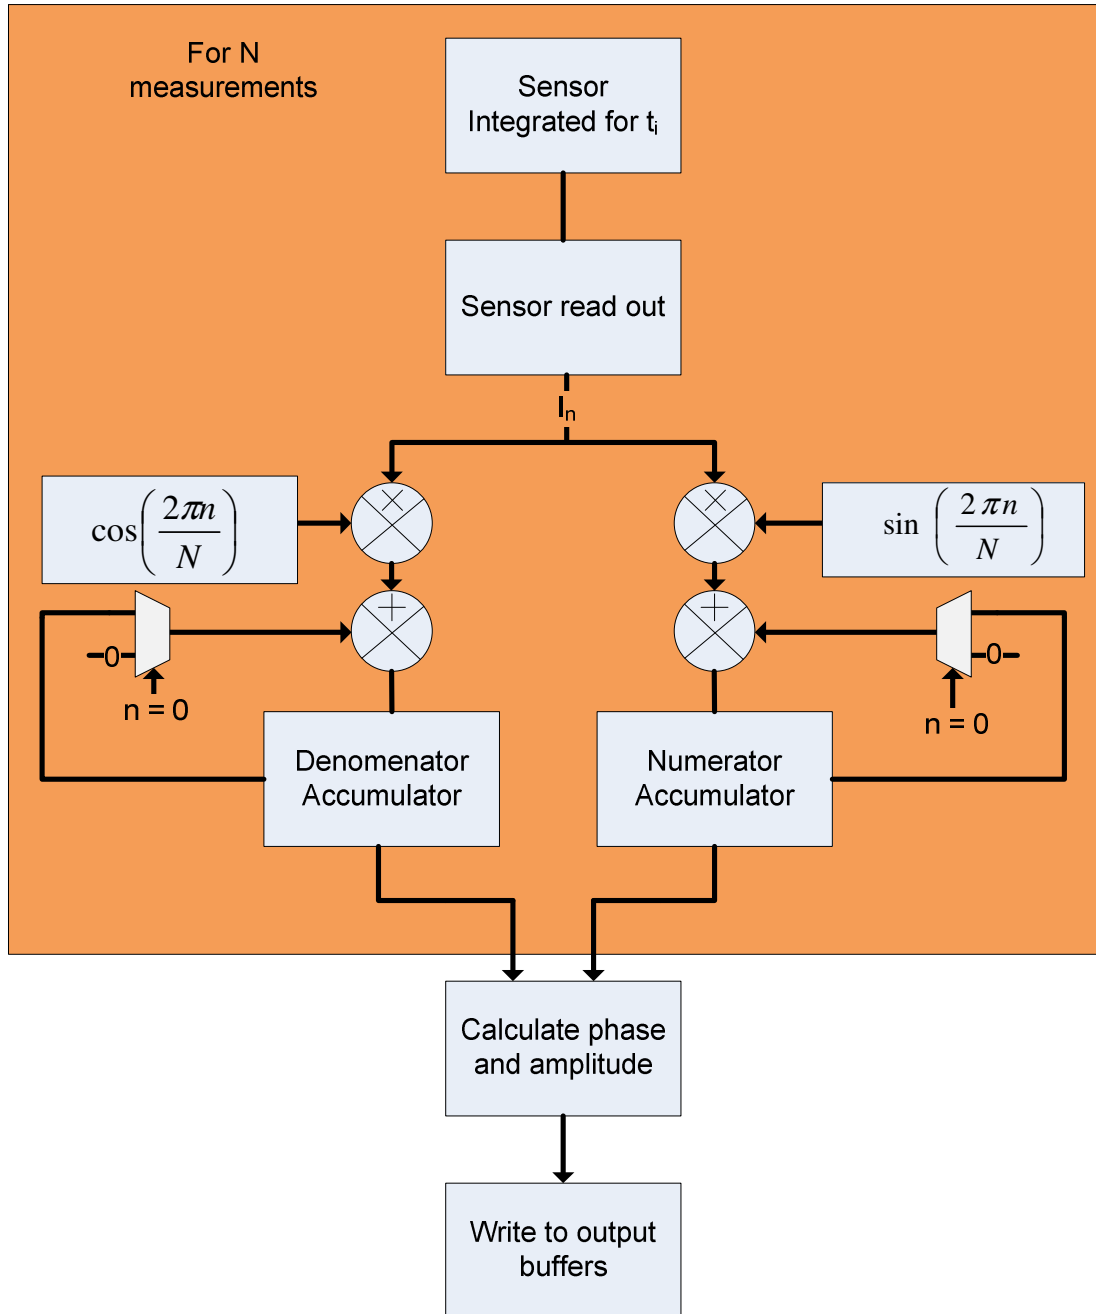


Figure 3.22 Diagram of the standard imaging process implemented on the compact system FPGA

3.6 Setup of the Motion Testing System

Central to this thesis is the study of the effect of motion on indirect time of flight range imaging cameras. To be able to measure such effects, accurate and repeatable motions must be performed. Victoria University did not have this capability at the start of this

research. This section outlines the design and implementation of a linear motion system for testing purposes.

3.6.1 LINEAR ACTUATOR

The linear actuator chosen to provide repeatable motion for the experimental section of this research is a MSA-M6S (Macron Dynamics Inc., 2011). This actuator was chosen as it has a high repeatability of ± 0.025 mm and is available in a travel of 4.2 m, which is sufficient for testing most of the measurement range of these cameras and able to fit in the laboratory space available. It also provides a T-slot in the actuator for easy mounting of sensors. The MSA-M6S is a belt driven actuator built for accelerations up to 5 g. The belt system has a feed rate of 150 mm/rev. To provide stability, the linear table is placed on four sets of supports with feet separated by 150 mm. A photograph of the linear actuator setup is shown in Figure 3.23. The surrounds of the actuator are covered in black cloth to minimise potential multipath errors. The first set of supports visible to the camera have been covered for the same reason.




Figure 3.23 Photograph of the M6A linear actuator setup

The belt is driven by a ST5909 high torque stepper motor from Nanotec (Nanotec GmbH, Munich, Germany). This motor has a step angle of 0.9° , resulting in a linear step size of 0.375 mm. A 48 V power rail is provided by a Topward 6303D bench top power supply configured in series mode generating ± 24 V rails.

To provide accurate position information, an HEDS-5540 optical encoder (Avago Technologies, San Jose, CA, USA) is attached to the stepper motor shaft. This provides 500 pulses per revolution and requires a 5 V power supply. With the given feed rate of 150 mm/rev this gives a linear resolution of ± 0.3 mm.

3.6.2 CONTROL SYSTEM

The ST5909 stepper motor is controlled using an SMCI33 stepper motor controller also from Nanotec. This controller can operate from a voltage supply between 12 V and 48 V, and includes one encoder and six optocoupler inputs.

The SMCI33 controller can be configured using Nanotec's NanoPro software, shown in Figure 3.24. A MATLAB interface is also provided that is used when automating tasks. The SMCI33 controller can be set to perform any programmed motor profile, consisting of a set speed, distance and acceleration and deceleration regimes described below, on an external trigger on one of its optocoupler inputs.

Using the NanoPro software, different motor drive profiles can be chosen. Three ramp profiles are available: Trapezoid, Sinus and Jerk Free. The Trapezoid ramp has a constant acceleration until the required speed is reached, the Sinus ramp minimises the jerk (the rate of change of acceleration) while reaching the required speed in the same amount of time as a Trapezoid ramp and the Jerk Free ramp sets a maximum level of acceptable jerk. For very low jerk the Jerk Free ramp is similar to the Sinus ramp, while for very large values of jerk it is similar to the Trapezoid ramp. The Jerk Free profile offers a compromise between fast acceleration and limiting the jerk experienced by the motor. The parameters available to adjust for this profile are the ramp (average acceleration), the jerk (maximum change in acceleration), the brake ramp and the brake jerk.

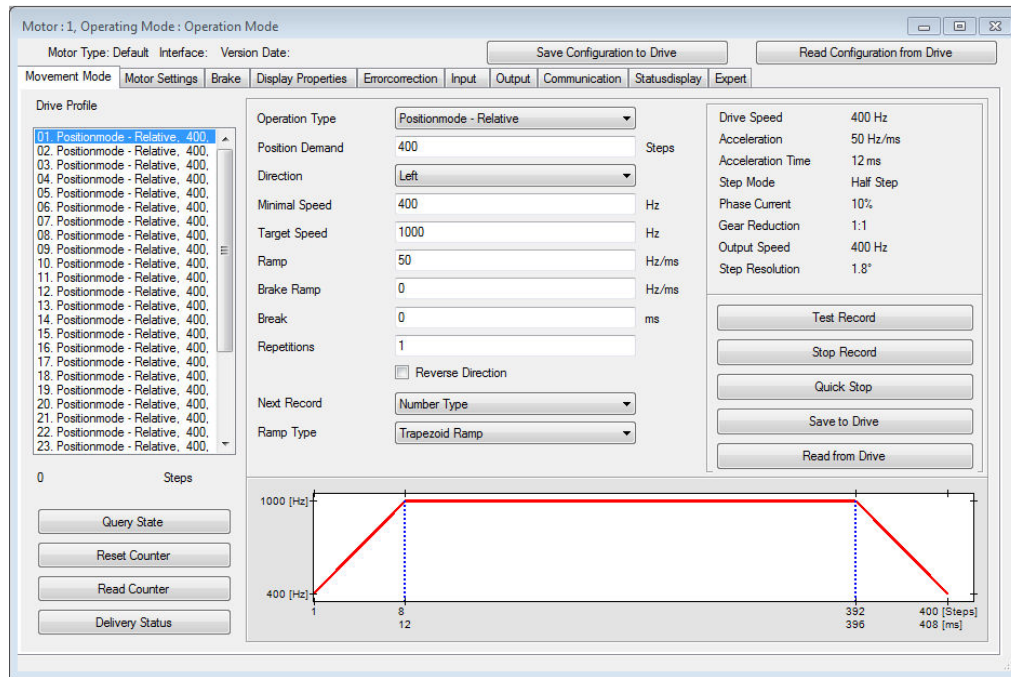


Figure 3.24 NanoPro software interface

The stepper motor may miss a phase step due to the speed, acceleration and loading on the motor causing it to stall, although this is not necessarily repeatable over multiple runs using the same settings. To test if a profile could be run reliably, the apparatus was set to run the profile both forward and in reverse 100 times. If the profile could complete this test without the motor stalling, it was considered reliable. The fastest reliable profile was capable of a speed of 2 m/s and could maintain this speed for 1 s before having to brake. This profile had a jerk of 1000 Hz/s^3 , a ramp of 10 Hz/ms , a break jerk of 5000 Hz/s^3 and a break ramp of 50 Hz/ms . The distance of the carriage versus time, recorded by the encoder, is shown in Figure 3.25 and the velocity is shown in Figure 3.26.

A reference position is provided by a 244-NPN-NC-06 T-Slot sensor (Macron Dynamics, Croydon, PA, USA). This is an inductive proximity sensor which is installed in the T-slot in the linear actuator. It detects the proximity of a metal tab on the carriage. This provides an absolute position to allow for continuity between tests. The NanoPro software has a homing run feature which aligns the carriage with the proximity sensor.




Figure 3.25 Distance vs. Time for the fastest stable motor profile (recorded using the encoder)




Figure 3.26 Velocity vs. Time for the fastest stable motor profile (recorded using the encoder)

To test the accuracy of this homing run, the carriage was first homed and then moved a random distance along the track by hand. The homing run was then run again. The encoder was used to detect if the carriage returned to the same position. Five tests were run and in all tests the carriage returned to the same position ± 1 encoder step. As discussed in section 3.6.1 the encoder has a linear resolution of 0.3 mm, meaning the homing is accurate to ± 0.3 mm. An example homing run recorded by the encoder is shown in Figure 3.27. The run initially overshoots the home sensor then returns over the sensor to find the exact location independent of the initial speed and braking time.

The repeatability of the linear table was tested independent of the encoder using a reference point marked on the bench below the linear actuator. The carriage was set at this point and moved forward then back 1875 mm (5000 steps). The settings used were:

- Ramp Profile - Jerk Free
- Target Speed – 1875 mm/s
- Ramp – 2 Hz/ms
- Jerk – 1000 Hz/s³
- Break Ramp – 2 Hz/ms
- Break Jerk – 1000 Hz/s³
- Break – 1 ms

The experiment was run three times with every run resulting in a final displacement of less than 0.5 mm (measured using a ruler). Since the motion errors being investigated using this apparatus are on the order of centimetres, this displacement is acceptable.




Figure 3.27 T-slot sensor homing run

3.6.3 INTERFACE

An RS422 connection is used to interface the SMCI33 motor controller to the FPGA controlling the range imaging system. This is implemented using two DS8921N controllers (National Semiconductor, Santa Clara, CA, USA). These are bi-directional

differential line drivers that are powered by 5 V. This interface is used to carry trigger signals from the FPGA to the linear motion system and return encoder data. Differential signals are chosen to eliminate ground loop noise, which was noticeable when the systems were connected directly.

The SMCI33 controller is placed in an ABS enclosure to improve stability. Four five pin sockets are available to connect to the range imaging camera, the encoder, the homing switch and the motor. Industrial 5 way series 678 circular sockets and plugs are used to ensure solid connections. A SCR-02 USB receptacle (Samtec, New Albany, IN, USA) is used to allow the connection of the SMCI33 controller to a computer via USB. A picture of the control box is shown in Figure 3.28.



Figure 3.28 Linear Actuation control box

LED indicators show the status of the trigger input from the FPGA and the homing sensor. A switch is available to disable triggering from the FPGA when it is not needed.

A diagram of the connection layout of the Motion Testing System is shown in Figure 3.29.

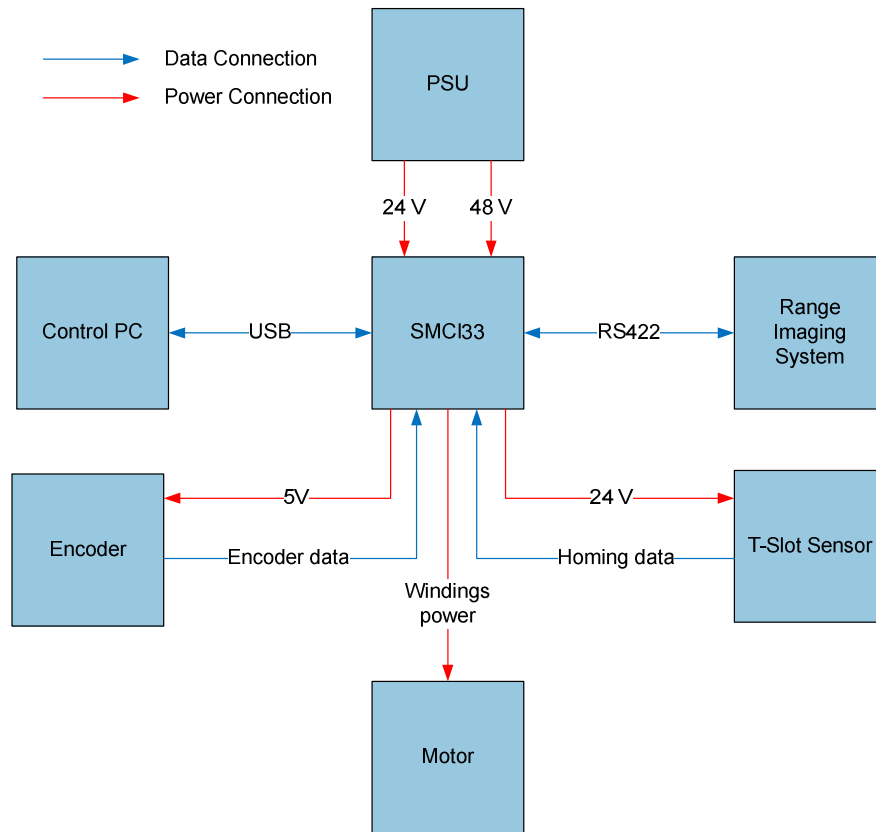


Figure 3.29 Motion Testing System interfaces

3.7 Summary

This chapter has detailed the hardware developments associated with this thesis. An indirect time of flight range imaging camera was designed specifically for use in mobile robotics research. It improves on a prototype system with the implementation of a custom made FPGA board to replace the bulky development board as well as introducing a number of improvements to auxiliary boards.

A linear actuation system was implemented to provide accurate repeatable motion. The system was set up to be triggerable by the range imaging system and will be used to provide experimental data on the range imaging system's response to moving objects.

Chapter 4 DYNAMIC RANGE IMPROVEMENTS

4.1 Introduction

As indirect time of flight range imaging cameras capture a full field of view, it is expected that objects in the scene will have different levels of reflectivity and will be at different distances. This results in the fraction of the illumination waveform reflected back to the sensor varying significantly across the scene. As the precision of the range measurements depends on the illumination received by the sensor, the dynamic range of the camera is crucial to acquiring high precision measurements across the entire scene. If the dynamic range of the camera is not adequate, some objects may saturate, meaning their range cannot be measured, or they may not return sufficient light to measure their distance with reasonable precision. This chapter will characterise the dynamic range and precision of the Victoria University Range Imaging System, present several methods for improving the dynamic range of indirect time of flight cameras in general and provide experimental data demonstrating the efficacy of each of these methods.

4.2 Experimental setup

An experimental apparatus containing three regions with different levels of reflectivity is used to provide quantitative measurements of the dynamic range of the Victoria University Range Imaging System. The three regions are a bright area perpendicular to the axis of the camera, a dark area perpendicular to the axis of the camera and a dark area at a 30° angle towards the camera. These were chosen to provide a large variation in returning light intensities. The apparatus consists of a flat wooden screen covered with black fabric. The bright area is a white piece of card attached to the screen. A triangular prism attached to the screen, with black card on the camera facing side, is used to provide the sloped area. A picture of the apparatus, and a top down diagram, is

shown in Figure 4.1. The bright, dark and sloped regions are labelled A, B and C respectively. The distance from the camera to the apparatus is 1.55 m.

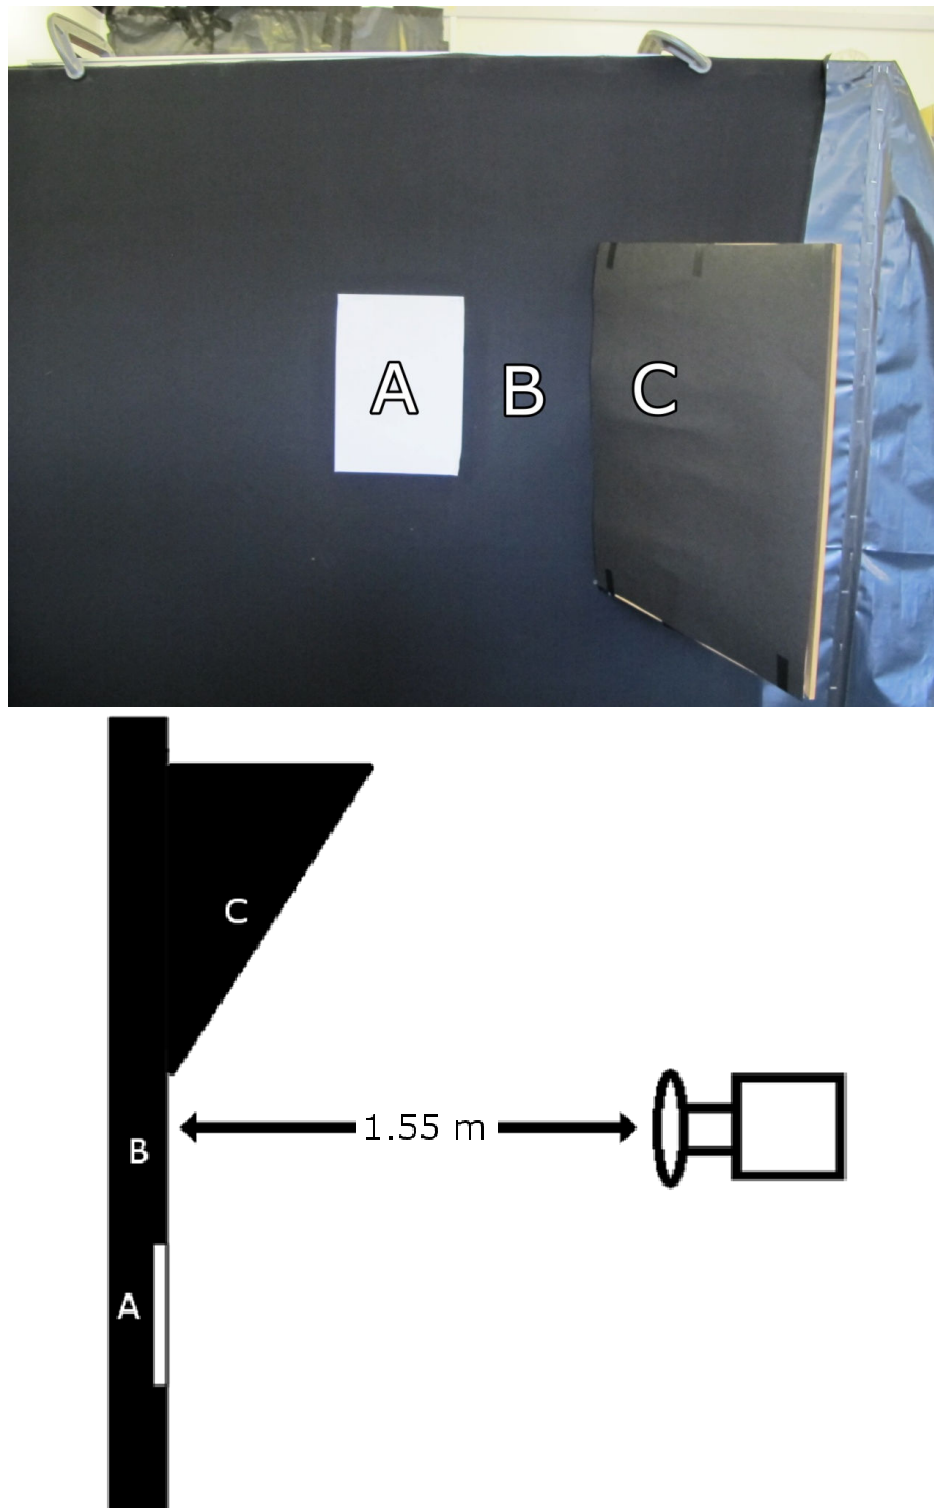


Figure 4.1 Photograph (top) and top view diagram (bottom) of apparatus for testing the dynamic range of indirect time of flight cameras

For each method described in this chapter three values for the measurement precision are calculated. These measurements are based on an area of 10 pixels by 60 pixels in the centre of each of the three regions. This ensures no border pixels are used, as each region is wider than 10 pixels, and that the regions are of the same size. The standard deviation of each pixel in the region is calculated and the result is then averaged over the region. The three regions are shown in Figure 4.2 on a typical range image. Colour is used in this image to represent the measured distance. All measurements in this chapter are acquired using the standard four frame phase algorithm. The maximum modulation frequency within the bandwidth limitations of the Victoria University Range Imaging System (40 MHz) was used for these measurements.




Figure 4.2 Layout of the three regions used for statistical analysis of the dynamic range

4.3 Characterisation of our system

In this section the Victoria University Range Imaging System will be characterised for a number of different phenomenon related to the precision and the response will be compared to theory. The dynamic range of the unaltered system, described in Chapter 3, will be measured to act as a control for the methods presented in this chapter.

Theory states that the standard deviation of the range measurements should have an inverse square root relationship with the frame time (Jongenelen, 2010). To test this, the bright region was imaged over 100 measurements and the standard deviation was

calculated for a range of frame times, with the results shown in Figure 4.3. This demonstrates that, as expected, there is an inverse relationship until the sensor saturates, at which point the standard deviation increases rapidly. To measure the exact relationship, a log-log plot was performed, and is shown in Figure 4.4 up until the point where the sensor saturates. A linear fit of this data has a slope of -1.21 ± 0.04 (Drayton et al., 2012 b). For an inverse square root relationship this slope is expected to be -0.5, therefore this is not the inverse square root relationship that was expected.




Figure 4.3 Standard Deviation of the measured phase versus Frame time for the bright region




Figure 4.4 Standard Deviation of the measured phase versus Frame time for the bright region (log-log)

To investigate the cause of this deviation from theory, 100 measurements were recorded of the individual intensity values for all four frames over a range of frame times. Again the bright region was used for these measurements. The average intensity values are shown in Figure 4.5. While it is expected that the intensity should increase linearly with frame time, until saturation is reached, the response of our camera is highly non-linear.




Figure 4.5 Measured Intensity of all four frames versus Frame Time for the bright region




Figure 4.6 Amplitude versus Frame Time for the bright region

Using the intensities from Figure 4.5, the amplitude and offset values for each frame time can be calculated. These values are expected to increase linearly. The amplitude is

shown in Figure 4.6 and the offset in Figure 4.7. These show that while the amplitude increases approximately linearly until the sensor starts saturating, the offset displays significantly non-linear behaviour well before saturation.




Figure 4.7 Offset versus Frame Time for the bright region

To test if the amplitude still has the inverse proportional relationship that is expected with changing frame time, the frame time was varied between 13 ms and 120 ms and both the amplitude and the phase were recorded over 100 measurements. A log-log plot of the data from all three regions is shown in Figure 4.8 with any saturated data removed. The slopes of linear fits to the log-log plots of the bright, dark and sloped regions are -0.83 ± 0.02 , -1.38 ± 0.03 and -1.56 ± 0.04 respectively. Therefore there is non-linearity in the amplitude of the measurements as well as the offset.

As the standard deviation should be proportional to \sqrt{B}/A (as shown in section 2.3.3) it is useful to also calculate this value with increasing frame time, which is shown in Figure 4.9 along with the precision, scaled appropriately, for comparison. A linear fit of a log-log plot of the ratio \sqrt{B}/A is shown in Figure 4.10 and has a slope of -1.16 ± 0.01 . This demonstrates that even though, for our system, the offset is not increasing linearly as desired, the ratio \sqrt{B}/A is still proportional to the standard deviation. Further confirmation of this is shown in Figure 4.11, where the two have been divided, giving a reasonably constant value. While the value is expected to be $1/\sqrt{2}$, because the amplitude and offset values are digital representations of the correlation waveform

some gain may be included. Therefore a deviation from the $1/\sqrt{2}$ value is not unexpected.




Figure 4.8 Log-log plot of standard deviation versus amplitude for all three regions




Figure 4.9 \sqrt{B}/A versus Frame Time for the bright region comparison with measured precision

To provide comparison values for the methods of improving the dynamic range of indirect time of flight cameras presented in this chapter, measurements were taken using the unmodified operating mode of the camera discussed in Chapter 3. This operating mode was the only mode available on the prototype range imaging system discussed in Section 3.2. A summary of the results are shown in Table 4.1. Due to the

cyclic nature of the phase measurement, which restricts the calculated phase values to the range 0 to 2π , the standard deviation flattens off at $\pi/\sqrt{3}$ radians, which is the standard deviation for randomly distributed values with this range. For a modulation frequency of 40 MHz this is equivalent to 1083 mm.




Figure 4.10 \sqrt{B}/A versus Frame Time for the bright region (log-log)





Figure 4.11 Division of the standard deviation over the \sqrt{B}/A metric for different frame times

Figure 4.12 shows a typical range image using a frame time of 25 ms. Using this frame time, the bright region produces a quality image with a standard deviation of

Table 4.1 Average standard deviation of pixels using standard operating mode for different reflectivity regions

Frame Time	Standard Deviation (mm)		
	Bright	Dark	Sloped
25 ms	5.7 ± 0.7	240 ± 70	700 ± 200
50 ms	3.1 ± 0.2	70 ± 20	180 ± 70
100 ms	2.2 ± 0.3	21 ± 2	41 ± 7

**Figure 4.12 Typical range image using standard measurement procedure and 25 ms frame time**

5.7 ± 0.7 mm, however both the dark and sloped regions have poor precision with standard deviations of 240 ± 70 mm and 700 ± 200 mm respectively.

Figure 4.13 shows a typical range image using a 50 ms frame time. Using this frame time decreases the standard deviation of the bright area to 3.1 ± 0.2 mm and improves the dark region to 70 ± 20 mm, which, while not ideal, is a significant improvement. The sloped region is still providing poor precision with a standard deviation of 180 ± 70 mm.

Figure 4.14 shows a typical range image using a 100 ms frame time. This frame time further improved the bright and dark areas and made the sloped area well imaged, with a standard deviation of 41 ± 7 mm. Frame times of higher than 100 ms caused the pixels in the bright region to saturate. Some curvature is visible in the measured distance across the sensor in the perpendicular region. This is caused by the propagation of the modulation signals through the sensor leading to spatial errors. Generally a Fixed Pattern Noise calibration is performed to correct this, as described in Chapter 2,




Figure 4.13 Typical range image using standard measurement procedure and 50 ms frame time




Figure 4.14 Typical range image using standard measurement procedure and 100 ms frame time

however, as it is not important for the dynamic range measurements, this has not been performed for these data.

4.4 Automatic frame time adjustment

This technique was discussed in Chapter 2. There is an element of subjectivity in how “good” a range image of a particular scene is, depending on the desired trade off between precision and frame rate for the application and the desire to eliminate dark pixels, saturated pixels or a combination of the two. As there are many different methods for finding the optimal frame time, depending on the individual

implementation, the optimal frame time for imaging the measurement apparatus could be a variety of values. While there is no clear “best” implementation that can be used for quantitative analysis of this method, the data acquired for the frame time of 100 ms is a reasonably good measure, as this is the highest frame time that could be used without pixels saturating, which is approximately the value these methods are generally trying to find.

4.5 High Dynamic Range Imaging

A method for improving dynamic range, which is commonly used in traditional photography, is combining two or more measurements using different frame times into a single range measurement. This is known as High Dynamic Range Imaging and has been demonstrated for indirect time of flight cameras (Gokturk et al., 2004) (Gokturk & Rafii, 2008). High Dynamic Range Imaging can combine an arbitrary number of frames, however a thorough analysis of the benefit of using more than two frames has not been performed. Figure 4.3 indicates that, unless the pixel saturates, when attempting to maximise precision it is always preferable to select a measurement using a longer frame time. Even implementing the simple approach of using only two different frame times can greatly increase the dynamic range of the camera, provided there is a significant difference between them (Gokturk & Rafii, 2008).

For simplicity, our implementation uses two range images with the longer time being four times the duration of the shorter time. The amplitudes of the two frames are compared on a pixel-by-pixel basis and the higher amplitude pixel is used. The amplitude of fully saturated pixels is essentially zero, since all four intensity values in saturated pixels are very similar, meaning non saturated pixels will always be chosen over fully saturated ones. It is possible for only one of the four samples to saturate while the other samples have not yet saturated. This will introduce some error into the measured phase and amplitude, but is still likely to have higher amplitude than a lower frame time measurement. A rigorous implementation of this technique will require a method for detecting these pixels, however that is outside the scope of this research.

While amplitude is not as good an indicator of the precision as \sqrt{B}/A , the precision of the two measurements is not expected to be similar for any individual pixel and

therefore the amplitude should be an acceptable metric to use. The FPGA implementation of the cordic arctangent function used in the Victoria University Range Imaging System automatically calculates the amplitude and it is therefore significantly easier to use as a metric and does not require additional computational power or memory, unlike the calculation of \sqrt{B}/A .

A diagram demonstrating this method is shown in Figure 4.15. TM is an integer multiplier on the frame time to provide the different exposure measurements. For our implementation this has two values, 1 and 4. The standard method is used to generate

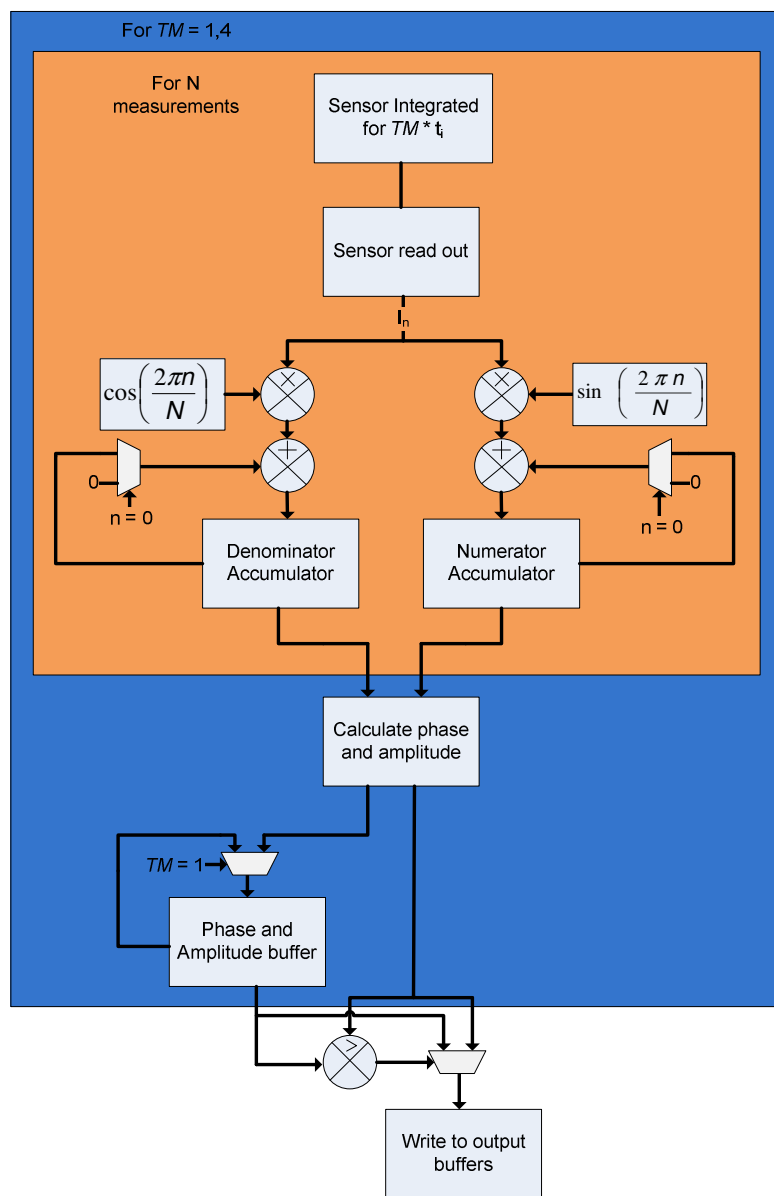


Figure 4.15 Operational diagram of the High Dynamic Range Imaging technique

the phase and amplitude measurements as described in Chapter 2. On one value of TM the values are stored in a buffer, on the other the two amplitude values are compared and the higher amplitude, and the corresponding phase value, are written to the output buffers.

In terms of precision, it will always be advantageous to choose a frame using a longer frame time, unless the pixel has saturated. However, this is not necessarily true for accuracy. There has been some research into using a fusing algorithm to combine data from multiple measurements using different frame times (Hahne & Alexa, 2011). This demonstrated that the overall accuracy could be improved using a fusing algorithm to combine multiple range images. They presented several fusing algorithms that provide compromises between computational time and accuracy, although all of the algorithms have some impact on the frame rate. The amount of improvement using this technique is scene dependent.

There are two possible implementations of this method. The intensity values can be cleared after each measurement or the shorter measurement data can be included into the longer measurement data. The frame rate will be either the frame rate of the sum of the frame times used or the frame rate of the longest capture used, depending on the implementation. Both of these are likely to be significantly longer than required to get a good image in at least some parts of the scene, particularly bright areas.

Figure 4.16 shows two typical range images used for our implementation of High Dynamic Range Imaging. The short exposure frame time is 50 ms and the long exposure frame time is 200 ms. For the short exposure the dark and the sloped regions have low precision. For the long exposure the bright area has completely saturated and the data from this region is unusable.

Figure 4.17 shows the range image resulting from combining the long and short capture images using the method described above. The resulting image provides a quality measurement in the bright, dark and sloped areas with standard deviations of 3.3 ± 0.5 mm, 12 ± 2 mm and 15 ± 3 mm respectively. The selection matrix used to choose between the images is also shown. The black region indicates pixels where the short frame time has higher amplitude than the long frame time and therefore the short

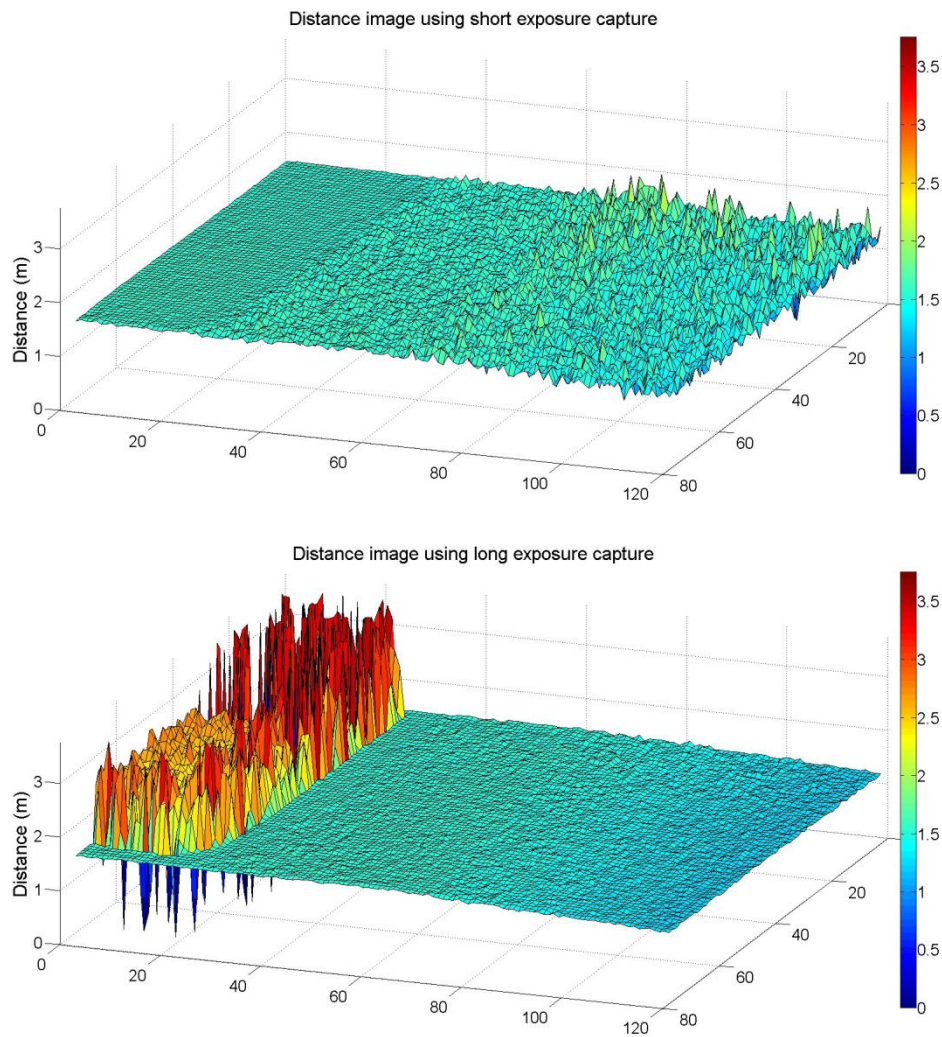


Figure 4.16 Typical Short capture (top, 50 ms) and long capture (bottom, 200 ms) used for High Dynamic Range Imaging

frame time phase value is used. Conversely the white area indicates the long frame time has higher amplitude so the long frame time phase value was used.

Close observation of the distance image in Figure 4.17 shows that there appears to be a slight discrepancy between the distance measured by the short capture and the distance measured for the long capture, which should be the same as it is outside of the sloped region of the scene. This can be explained by the non-linearity of our system. Using the data from Figure 4.5, the measured phase for the same object with changing frame time can be calculated. This is shown in Figure 4.18, confirming that our camera has a phase shift with changing frame time. Non-linearity of the phase with changing frame time has been observed in other PMD based cameras (Wiedemann et al., 2008) and also in

SwissRanger cameras (Kahlmann et al., 2006). Some camera calibration techniques have incorporated a calibration for frame time related error (Radmer et al., 2008) (Lindner & Kolb, 2007), which can mitigate this issue.

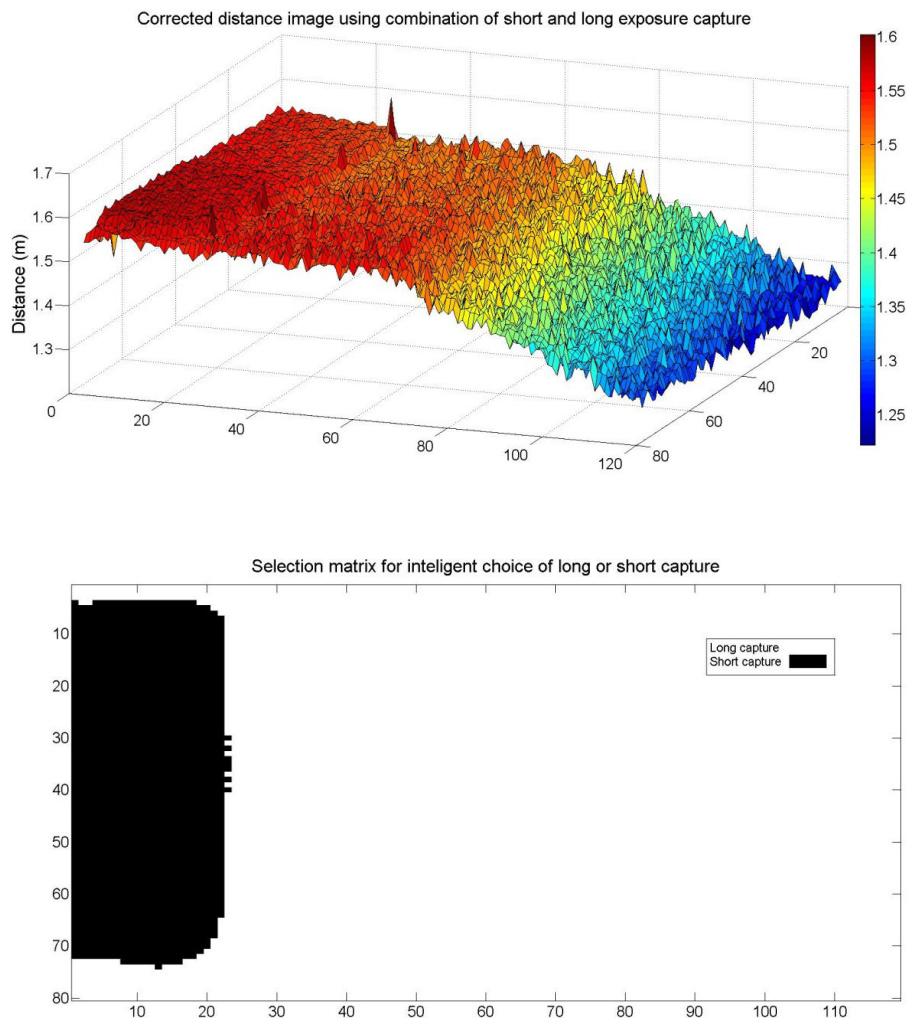


Figure 4.17 High Dynamic Range Image (top) and selection matrix used (bottom) for a typical multi capture range image

4.6 Variable Frame Rate Imaging

4.6.1 INTRODUCTION

In traditional indirect time of flight range imaging, the accumulators for the numerator and denominator of the arctangent are cleared after each measurement. The dynamic range of indirect time of flight cameras can be improved by selectively clearing or maintaining these accumulators on a pixel-by-pixel basis, depending on the quality of




Figure 4.18 Measured Phase versus frame time for the Victoria University Range Imaging System

the measurement (Drayton et al., 2013). Over numerous samples the noise will cancel and the remaining signal will add towards the correct phase. It is important that the individual numerator and denominator registers are averaged rather than averaging multiple phase measurements, as otherwise systematic errors can be introduced (Godbaz et al., 2011).

Due to the readout architecture of the PMD19K-2 sensor, it is not possible to individually vary the frame time for each pixel on the sensor. However, this effect can be replicated by choosing a relatively low frame time for the entire sensor and then summing the intensity frames over a number of measurements during processing. Other pixel designs have been demonstrated where on chip circuitry is used to stop integration based on a voltage threshold (Buttgen et al., 2005) and where the integration period is adapted on a pixel-by-pixel basis controlled by a comparison of the voltage output with a reference voltage on chip (Lehmann et al., 2005). However, these pixel designs increase the complexity of the pixel. These more complex pixel designs have not achieved wide spread use. The concept has only recently been explored in the processing stage, suitable for implementing on commercial cameras implementing standard pixel architectures (Drayton et al., 2013). An advantage of this method is that, with an increase in memory requirements, the standard dynamic range image can be calculated and output concurrently with the Variable Frame Rate image.

Two different implementations of this method have been developed. The first uses the amplitude as the metric for pixel quality, similar to the High Dynamic Range Imaging method in Section 4.5. The second method replaces this metric with the ratio \sqrt{B}/A .

The accumulators are not cleared until the threshold value is reached. When the threshold value is reached a flag is set that controls the source of the accumulators. The next time the pixel is processed zero will be added to the value to be stored in the accumulators, instead of its previous value. The flag also controls writing to the output buffer for that pixel so the phase value is only updated with the reset of the accumulators. Thresholds are set in arbitrary units used to compare with the amplitude produced by the cordic arctangent implemented on the FPGA.

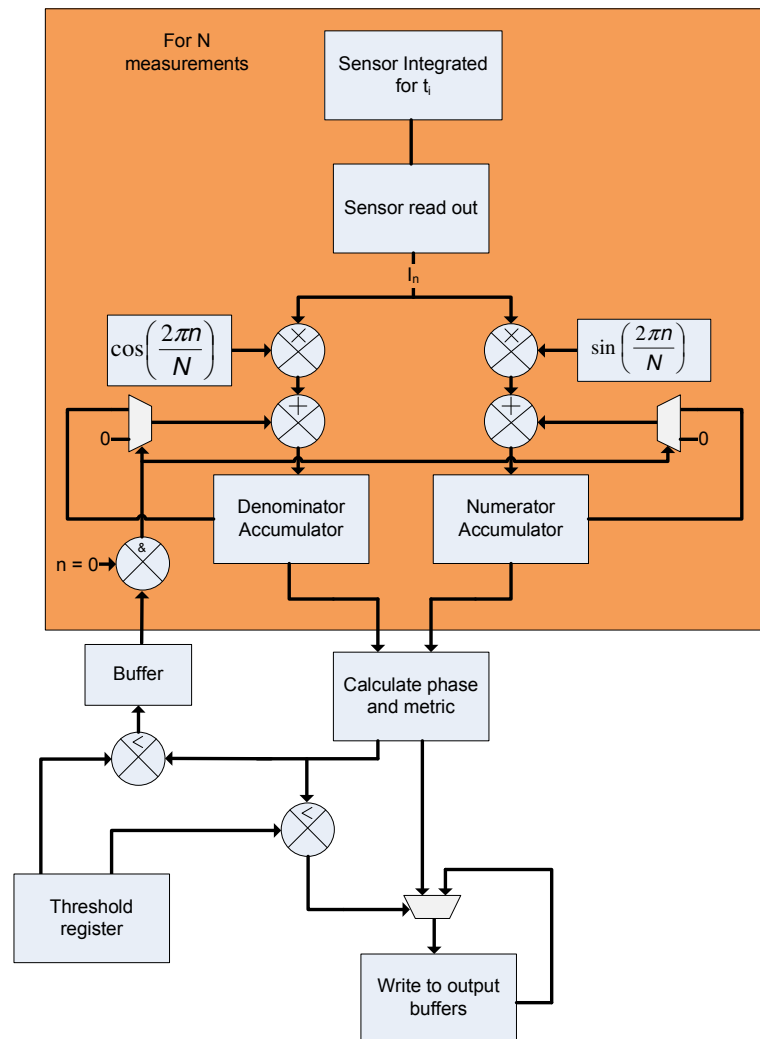


Figure 4.19 Operational Diagram of Variable Frame Rate Imaging

A diagram demonstrating this method is shown in Figure 4.19. After each measurement the phase and amplitude are calculated. The amplitude is then compared to a threshold value. The result of the comparison determines if the output buffers should be written to. It is also stored in a buffer so for the next measurement the real and imaginary accumulators are loaded with 0 instead of the previous value.

This method sacrifices the uniformity of the frame rate across the image in return for being able to acquire high quality images in dark areas without a detrimental impact on the frame rate of brighter areas of the scene.

4.6.2 THEORY DEVELOPED IN THIS RESEARCH

To provide theoretical analysis of Variable Frame Rate Imaging we need to understand the distribution of the errors in the intensity measurements. The distributions for the intensities of each frame are shown in Figure 4.20, measured by taking 1000 measurements of a static object. These indicate that the noise sources acting on the intensity values are approximately normally distributed. As the intensity values are calculated as the subtraction of two charge bins, shown in Figure 2.6, the sign therefore simply indicates which bin has accumulated more charge.

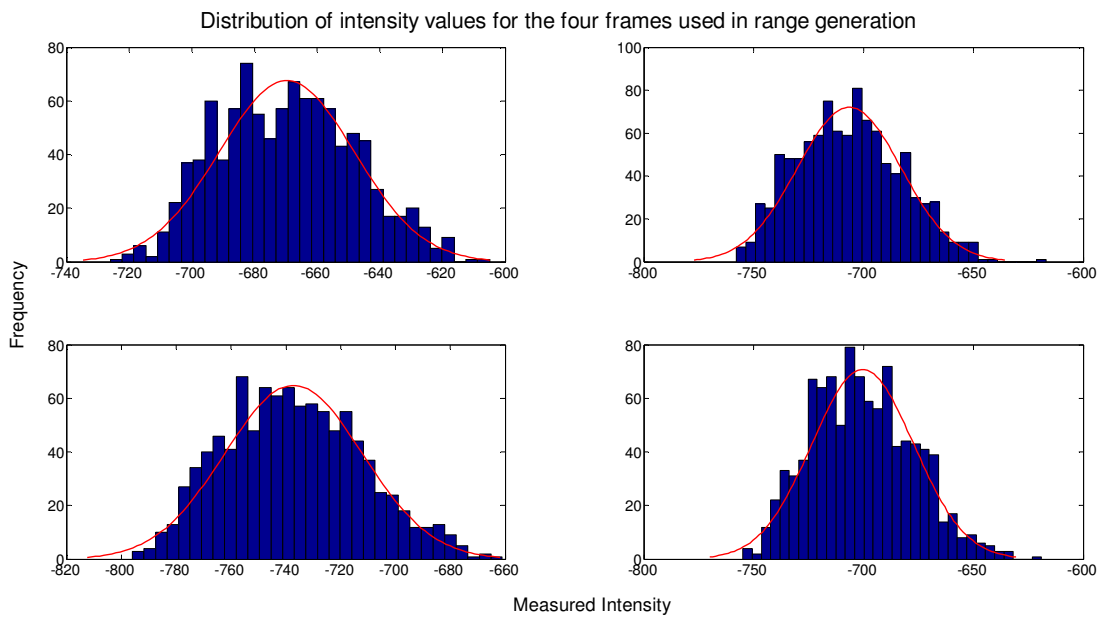


Figure 4.20 Measured distribution of intensity values for the four frames used in phase calculation

How this error changes in differently illuminated areas of the scene is also important if we want to simulate imaging a scene with high contrast. To ascertain this, 100 measurements were taken of a scene and two objects were identified with very different levels of precision. Histograms of the intensity values for the brighter object are shown in Figure 4.21 and histograms of the intensity values for the darker object are shown in Figure 4.22. This demonstrates that both the bright and dark objects intensity distributions have approximately the same standard deviations. However, the difference between the mean values of the intensity distributions (in other words the amplitude of the correlation waveform) and the offset, have changed significantly. Therefore, change




Figure 4.21 Distribution of intensities for a well imaged area of a scene




Figure 4.22 Distribution of intensities for a poorly imaged area of a scene

in the amplitude and offset of the noise free correlation waveform is the primary mechanism for the change in signal to noise ratio across the imaged scene, with the noise level being reasonably constant.

From Chapter 2 the equation for the amplitude is

$$A = \frac{2}{N} \sqrt{\left[\sum_{n=0}^{N-1} I_n \cos\left(\frac{2\pi n}{N}\right) \right]^2 + \left[\sum_{n=0}^{N-1} I_n \sin\left(\frac{2\pi n}{N}\right) \right]^2} \quad (4.1)$$

when measuring over M measurements and summing the real and imaginary terms this becomes

$$A = \frac{2}{N} \sqrt{\left[\sum_{m=0}^{M-1} \sum_{n=0}^{N-1} I_{n,m} \cos\left(\frac{2\pi n}{N}\right) \right]^2 + \left[\sum_{m=0}^{M-1} \sum_{n=0}^{N-1} I_{n,m} \sin\left(\frac{2\pi n}{N}\right) \right]^2}, \quad (4.2)$$

where m is the measurement number and $I_{n,m}$ is the intensity measured for the m^{th} measurement of the n^{th} frame. It is useful to express $I_{n,m}$ as

$$I_{n,m} = I_n + \varepsilon_{m,n}, \quad (4.3)$$

where I_n is the noise free intensity of frame n and $\varepsilon_{m,n}$ is distributed noise for the measurement m which has a mean of 0. From the measurements in Figure 4.20 we know that $\varepsilon_{m,n}$ can be assumed to be normally distributed. The amplitude is therefore

$$\begin{aligned} A &= \frac{2}{N} \sqrt{\left[\sum_{m=0}^{M-1} \sum_{n=0}^{N-1} (I_n + \varepsilon_{m,n}) \cos\left(\frac{2\pi n}{N}\right) \right]^2 + \left[\sum_{m=0}^{M-1} \sum_{n=0}^{N-1} (I_n + \varepsilon_{m,n}) \sin\left(\frac{2\pi n}{N}\right) \right]^2} \\ &= \frac{2}{N} \sqrt{\left[M \sum_{n=0}^{N-1} I_n \cos\left(\frac{2\pi n}{N}\right) + \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} \varepsilon_{m,n} \cos\left(\frac{2\pi n}{N}\right) \right]^2 + \left[M \sum_{n=0}^{N-1} I_n \sin\left(\frac{2\pi n}{N}\right) + \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} \varepsilon_{m,n} \sin\left(\frac{2\pi n}{N}\right) \right]^2}. \end{aligned} \quad (4.4)$$

Taking $N = 4$ and simplifying gives

$$A = \frac{M}{2} \sqrt{\left[(I_0 - I_2) + \frac{1}{M} \sum_{m=0}^{M-1} (\varepsilon_{m,0} - \varepsilon_{m,2}) \right]^2 + \left[(I_1 - I_3) + \frac{1}{M} \sum_{m=0}^{M-1} (\varepsilon_{m,1} - \varepsilon_{m,3}) \right]^2}. \quad (4.5)$$

For large values of M , the errors in the intensity measurements will become insignificant and the amplitude will increase proportional to M . For small values of M the individual errors will have a significant impact on the amplitude and the relationship between the

amplitude and the number of measurements M , while averaging to being linear, may have a very large spread.

It was stated that the noise $\epsilon_{m,n}$ had a mean of 0. This implies that the average value of I_n does not change with changing M . To validate this, the average intensity value for all four frames was measured for different values of M and is shown in Figure 4.23. The intensity values have been divided by M to make them consistent across the measurements. While all four intensity values show some dependence on the number of frames measured over, the strength of this relationship is very small. The largest effect is observed for the second frame and is a change of -0.04% between a single measurement and averaging over 100 measurements. This is sufficiently small that it can be ignored for our purposes.




Figure 4.23 Average intensity value versus number of measurements for the Victoria University Range Imaging System

Using (4.5) for the amplitude, and estimating an initial value for the amplitude (A_m), B and the standard deviation of $\epsilon_{m,n}$ from the measurements taken for Figure 4.20, the precision of the phase measurements with distance can be simulated. Several different amplitudes were tested centred around the amplitude calculated from the intensity measurements. For each experiment 1000 simulations were run. The change in amplitude versus number of measurements for different amplitudes is shown in Figure 4.24. This shows that the amplitude increases linearly with M , the slope of which is determined by the amplitude of the noise-free correlation waveform.




Figure 4.24 Simulated amplitude versus number of measurements for different intensities

This relationship can be experimentally verified by summing the output of the camera over a large number of measurements and recording the amplitude for each value of M . The data for both the dark and sloped regions are shown in Figure 4.25. Data for the bright region is not included as the amplitude was increasing too rapidly to acquire a reasonable number of frames. This figure demonstrates that, as expected, there is a linear relationship between amplitude and the number of frames integrated over, with the slope being determined by the amplitude of the noise-free correlation waveform.




Figure 4.25 Measured amplitude versus number of measurements for different regions

The same process used to find (4.5) can be used to determine the effect on the phase, which yields

$$\varphi = \tan^{-1} \left(\frac{(I_1 - I_3) + \frac{1}{M} \sum_{m=0}^{M-1} (\varepsilon_{m,1} - \varepsilon_{m,3})}{(I_0 - I_2) + \frac{1}{M} \sum_{m=0}^{M-1} (\varepsilon_{m,0} - \varepsilon_{m,2})} \right). \quad (4.6)$$

The relationship between the standard deviation and the number of measurements can be found using simulations. Again 1000 simulations were performed and the results are shown in Figure 4.26. As expected, the standard deviation decreases with an inverse square root relationship with the number of measurements. For low amplitude measurements, there is a flattening of the curve as the standard deviation of the measurements is limited by the fact the phase is within the range 0 to 2π . For these measurements the phase was set to π to avoid phase wrapping issues as the standard deviation reaches its limit. If other phases are used, the results for simulations with reasonable SNR or for large M are the same, with an inverse square root decrease. The results for low SNR measurements are highly noisy, as measurements crossing the phase wrapping point cause large perturbations in the standard deviation.




Figure 4.26 Simulated standard deviation in phase when integrating over multiple measurements for a variety of amplitude values

The precision of the system with changing M was measured directly by taking 100 measurements for each value of M and calculating the standard deviation. The results are shown in Figure 4.27 for the dark region and in Figure 4.28 for the sloped region. For the dark region the response has a constant slope in the log-log plot of -0.53 ± 0.03 ,

which is consistent with an object with a reasonable SNR. For the sloped region there are two sections to the response. For the first 11 measurements the standard deviation is improving slowly and is at approximately the maximum standard deviation value of 1083 mm. After the 11th measurement a slope of -0.6 ± 0.1 is observed, which agrees with the expected slope of -0.5 within uncertainty. This is consistent with the simulations for a pixel with a low amplitude to noise ratio.




Figure 4.27 Measured standard deviation versus number of frames integrated over for the dark region




Figure 4.28 Measured standard deviation versus number of frames integrated over for the sloped region

4.6.3 DYNAMIC RANGE TESTING

Using the scene described in Section 4.1, data were recorded for different frame times and thresholds using the amplitude as the metric. Similarly to the standard deviation, to

calculate the frame rate for each region the average frame rate of each pixel is calculated and the result is averaged over the region.

Figure 4.29 shows that using a 25 ms frame time with a threshold of 100 can recover the distance information for the dark area of the scene (compared with the non-thresholded capture shown in Figure 4.12), however the sloped area is still not measured correctly. The frame rate in the dark area has decreased to 0.6 ± 0.1 FPS while the maximum frame rate of 10 FPS is maintained in the bright area.

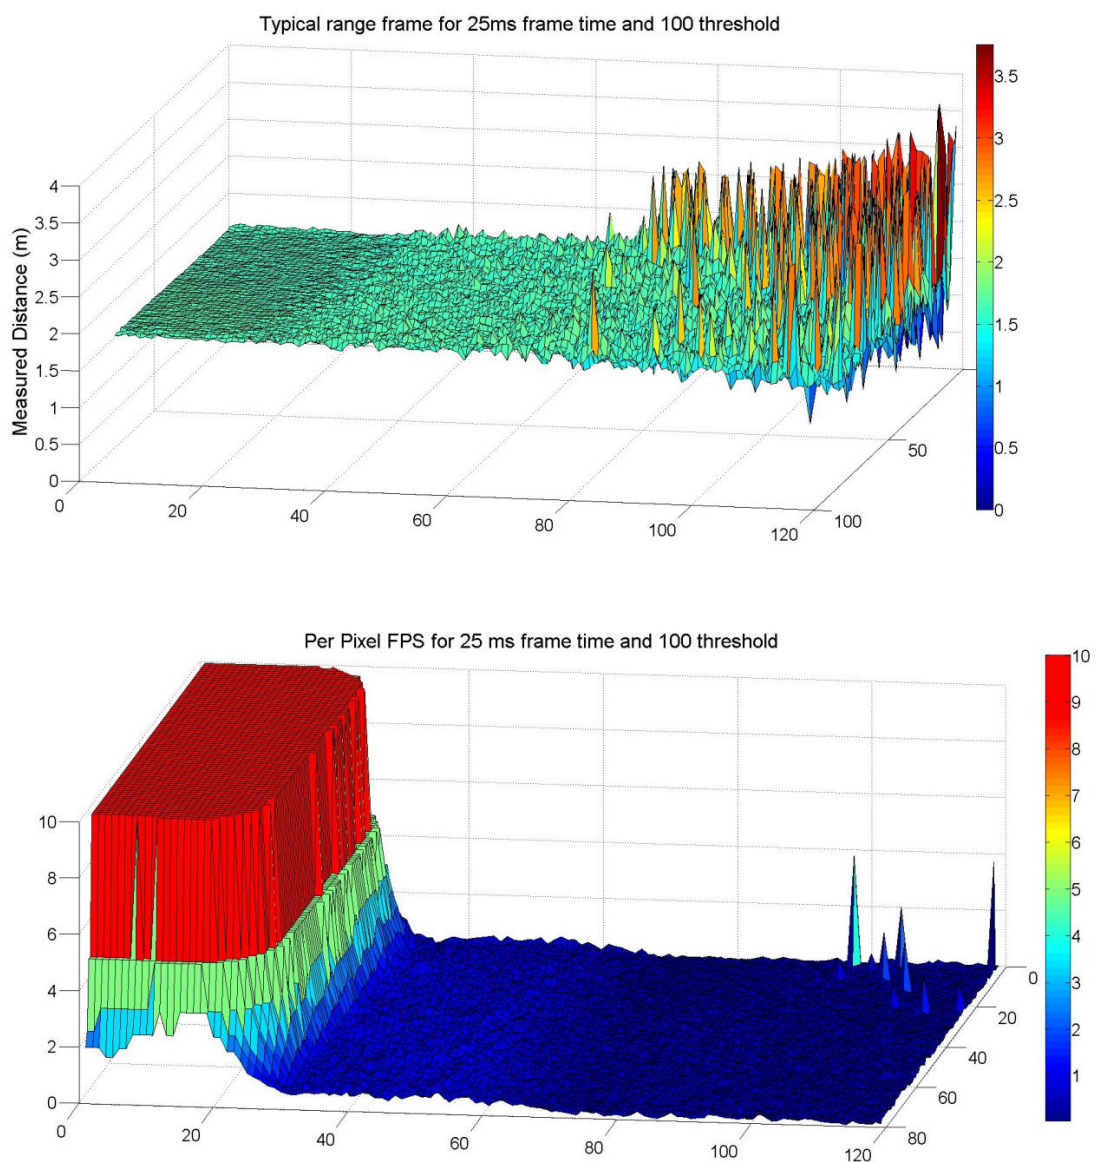


Figure 4.29 Measured distance (top) and frame rate (bottom) for a typical image using a frame time of 25 ms and a threshold value of 100

With a 50 ms frame time, shown in Figure 4.30, where previously the bright and dark regions were imaged reasonably well, as shown in Figure 4.13, the sloped area has now been recovered using a threshold of 100. The bright area has the maximum frame rate of 5 FPS while the dark area has a frame rate of 1.1 ± 0.1 FPS. The average frame rate of the sloped area is 0.38 ± 0.09 FPS.

Table 4.2 shows a comparison of both the standard deviation and frame rate for different regions using different combinations of both frame time and threshold value. The results are also plotted in Figure 4.31. The standard mode (threshold = 0) is included for comparison. Both increasing the threshold and increasing the frame time improves

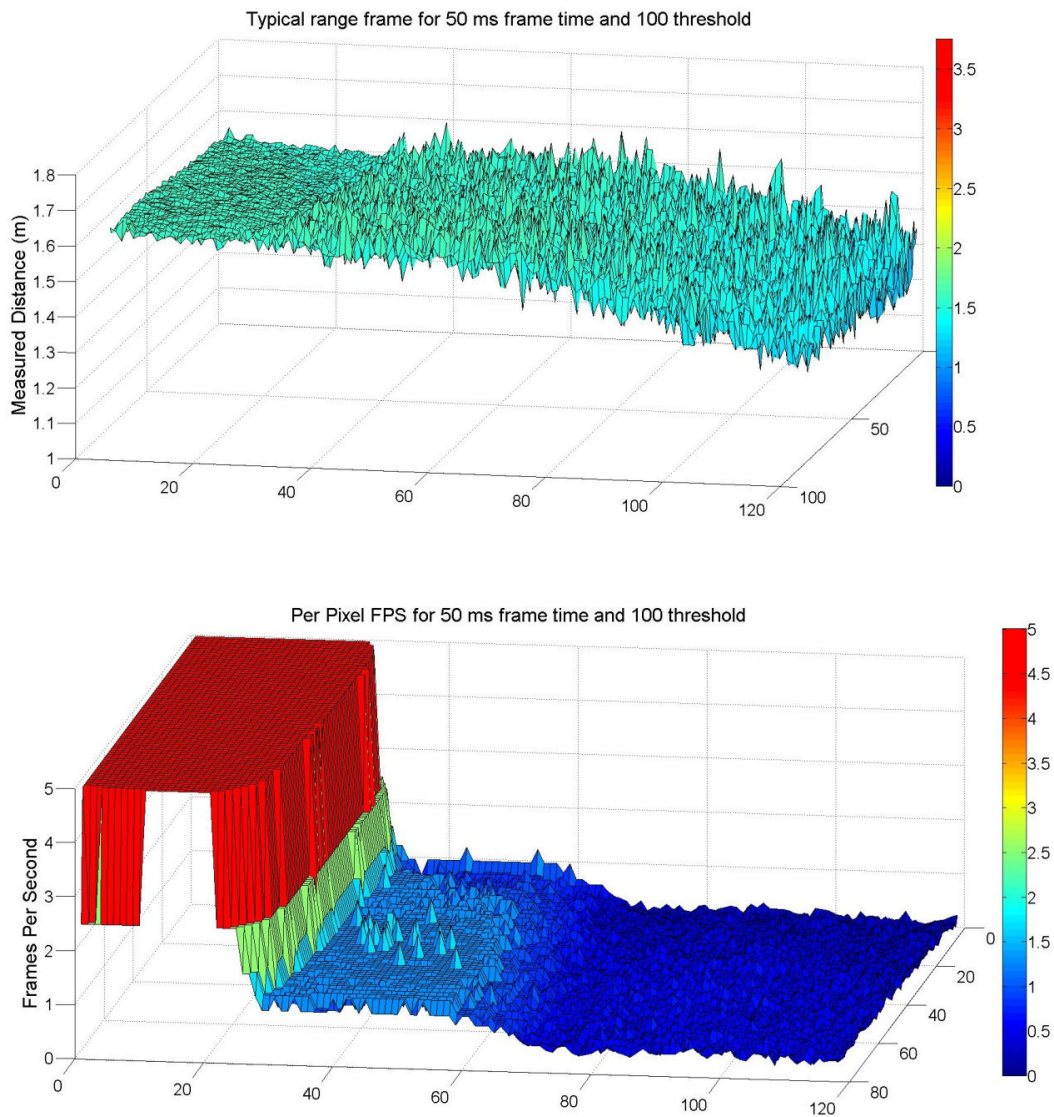


Figure 4.30 Measured distance (top) and frame rate (bottom) for a typical image using a frame time of 50 ms and a threshold value of 100

the precision of the system, however only a limited number of data points are available. Although the use of the amplitude as the threshold has successfully targeted areas of the scene that need longer frame times to provide a good measurement, it falls short of providing a constant precision across the field of view of the camera. While the dark and sloped regions were improved using thresholding, the precision of the dark region was always significantly worse than the bright region and the sloped region was always significantly worse than the dark region.

Table 4.2 Comparison of the standard deviation and frame rate of different reflectivity regions using a variable pixel rate

Frame time	Threshold	Standard Deviation (mm)			Frame rate (frames per second)		
		Bright	Dark	Sloped	Bright	Dark	Sloped
25 ms	0	5.7 ± 0.7	240 ± 70	700 ± 200	10 ± 0	10 ± 0	10 ± 0
	50	5.6 ± 0.8	70 ± 30	250 ± 80	10 ± 0	1.2 ± 0.3	0.3 ± 0.2
	100	5.5 ± 0.8	50 ± 60	200 ± 200	10 ± 0	0.6 ± 0.1	0.12 ± 0.05
50 ms	0	3.1 ± 0.2	70 ± 20	180 ± 70	5 ± 0	5 ± 0	5 ± 0
	50	3.0 ± 0.3	42 ± 5	61 ± 8	5 ± 0	2.2 ± 0.3	0.8 ± 0.02
	100	3.0 ± 0.3	28 ± 4	40 ± 30	5 ± 0	1.1 ± 0.1	0.38 ± 0.09
100 ms	0	2.2 ± 0.3	21 ± 2	41 ± 7	2.5 ± 0	2.5 ± 0	2.5 ± 0
	50	2.1 ± 0.2	20 ± 2	31 ± 5	2.5 ± 0	2.5 ± 0	1.4 ± 0.2
	100	2.2 ± 0.2	16 ± 2	23 ± 4	2.5 ± 0	1.6 ± 0.5	0.82 ± 0.08

Standard Deviation for different configurations of the pixel thresholding method for the bright (blue), dark (green) and sloped (red) regions




Figure 4.31 Standard deviation for different configurations using Variable Frame Rate Imaging

Using this method, a high frame rate for objects that are expected to be relatively bright (for example foreground objects) is achievable without having to sacrifice acquiring quality range images of the background and other dark objects, provided a drop in frame rate for these objects is acceptable.

The thresholded approach produces a significantly decreased frame rate to achieve a similar improvement in standard deviation to that from simply increasing the frame time. As an example of this, using a 25 ms frame time and a threshold of 50 gives a standard deviation in the dark area of 70 ± 30 mm compared to using 50 ms frame time with no thresholding, which gives a standard deviation of 70 ± 20 mm. The frame rates of these two measurements are 1.2 ± 0.3 FPS and 5 FPS respectively. However, there is advantage for VFRI in a higher frame rate in the bright region with comparable precision.

The requirement for part of the frame time to be used for readout of the sensor is a potential cause for degradation of the frame rate. However, it is not sufficient to explain this disparity. As shown in Section 4.3, the response of our system with changing frame time is non-linear and this results in the precision of the system improving not with the inverse square root improvement expected but significantly faster. This explains the improvement observed by recording a single long capture over increasing the number of measurements integrated over, as the latter follows the expected inverse square root improvement in precision.

4.6.4 COMPARISON WITH A COMMERCIAL CAMERA

To determine if this non-linearity is a problem with our particular sensor, or if it is common in indirect time of flight range imaging cameras, the measurements were repeated using a SR-4000 commercial camera from Mesa Imaging (Mesa Imaging, Zurich, Switzerland). These measurements were performed using the camera's "raw" mode to record uncalibrated raw intensity values. The results are shown in Figure 4.32. This demonstrates that the SR-4000 camera has a significantly more linear response than the Victoria University Range Imaging System.




Figure 4.32 Measured intensity versus Frame Time for all four frames of the SR-4000 camera

The measurements for both increasing the frame time and increasing the number of frames integrated over were repeated using the SR-4000 camera. As the intensity values for this sensor appear much more linear than for the Victoria University Range Imaging System, no relative degradation due to non-linearity is expected. Measurements were performed for all three regions and a comparison is made between increasing the Frame Time (T) and increasing the number of measurements (M) to provide the same effective measurement time. The results are shown in Figure 4.33. This demonstrates that there is still a disadvantage to integrating over multiple measurements compared to increasing the frame time.




Figure 4.33 Standard Deviation versus effective frame time comparing increasing frame time with integrating over multiple frames

The number of measurements integrated over can be extended significantly beyond the number used for comparison with increasing the frame time, as shown in Figure 4.34. This shows a consistent improvement with a log-log slope of -0.5 for all three regions. Similar to the Victoria University Range Imaging System, the SR-4000 shows better than expected improvement in precision with increasing frame time, therefore the change in the amplitude and offset of the correlation waveform with increasing integration time must be investigated more closely for the SR-4000 camera.




Figure 4.34 Standard deviation versus number of measurements averaged over for all three regions using an SR-4000 commercial camera

As it has the largest deviation between the two methods, the sloped region was used for this investigation. First the amplitude was measured with changing effective frame time and is shown in Figure 4.35. The effective frame time is the frame time multiplied by the number of frames integrated over. There is no significant difference in the amplitude of the correlation waveform between increasing the frame time and integrating over multiple measurements.

Measurement of the offset is not as simple as the measurement of the amplitude. The four intensity values produced by the SR-4000 camera are not direct representations of the measured correlation waveform, instead a digital offset is added to approximately centre them on the midway point of the 16-bit output data range (32768). However, using exactly this value provides negative offset values which increase with increasing frame time and therefore are not correct. To attempt to measure the zero offset point,

the camera was set to the minimum frame time possible (0.3 ms) and an image was recorded of an object with very low reflectivity in a darkened room. The resulting offset value was 31869.5 and subtracting this value from the measured offset provided reasonable answers. For the sloped region the offset for both methods is shown in Figure 4.36. This demonstrates that, similar to the Victoria University Range Imaging System, the offset is not increasing linearly with the frame time for the sloped region.




Figure 4.35 Measured amplitude versus effective frame time for the sloped region using an SR-4000




Figure 4.36 Measured offset versus frame time for the sloped region using an SR-4000

For comparison the dark and bright regions are shown in Figure 4.37 and Figure 4.38 respectively. Again the offset is not increasing at the same linear rate when increasing

the frame time compared to integrating over multiple frames. However, as the object being measured becomes brighter, this gap becomes smaller. For high frame times some pixels have saturated in both the bright and dark regions causing a sudden increase in the offset.

The cause of the non-linearity in the offset of the sensor for low light regions is not clear. One potential cause is the background illumination. The offset is a combination of both the offset caused by the ambient light and the offset due to the correlation waveform. To measure how much of this is contributed by the background illumination, the




Figure 4.37 Measured offset versus frame time for the dark region using an SR-4000




Figure 4.38 Measured offset versus frame time for the bright region using an SR-4000

measurements for the sloped region were repeated with the room darkened. A comparison between the measurements taken with normal background illumination and the measurements taken in the dark is shown in Figure 4.39. This demonstrates that the ambient light has added a constant offset, however, it is not responsible for the non-linearity observed.

Measurements using the SR-4000 camera have shown that the improved precision from increasing the frame time over integrating over multiple frames is not an effect limited to our particular indirect time of flight camera. In both cameras this effect appears to be caused by non-linearity in the offset of the correlation waveform with increasing frame time, however a theoretical explanation for this has not yet been formulated.




Figure 4.39 Offset versus frame time comparison between normal background light and darkness in the sloped region using an SR-4000 camera

4.6.5 EFFECT OF AXIAL MOTION

So far in the analysis of the VFRI technique the scene being imaged has been assumed to be static. As these cameras are likely to be used in dynamic environments it is interesting to investigate the effect of motion on this technique. Equation (2.24) for the intensity frame I_n can be extended to represent the intensity of frame I_n after m measurements as follows

$$\begin{aligned}
I_n &= \sum_{m=0}^{M-1} \cos(n\delta - \varphi_{n,m}) \\
I_n &= \sum_{m=0}^{M-1} \cos(n\delta - \varphi_0 - nv\alpha - 4mv\alpha) \\
&= \sum_{m=0}^{M-1} \cos(\rho) \cos(-4mv\alpha) - \sin(\rho) \sin(-4mv\alpha) \\
\rho &= n\delta - \varphi_0 - nv\alpha \\
I_n &= \cos(\rho) \sum_{m=0}^{M-1} \cos(-4mv\alpha) - \sin(\rho) \sum_{m=0}^{M-1} \sin(-4mv\alpha) \\
&= c_1 \cos(\rho) - c_2 \sin(\rho) \\
c_1 &= \sum_{m=0}^{M-1} \cos(-4mv\alpha), c_2 = \sum_{m=0}^{M-1} \sin(-4mv\alpha). \tag{4.7}
\end{aligned}$$

When using the standard four frame algorithm we then want to calculate the values I_1 - I_3 and I_0 - I_2 .

$$\begin{aligned}
I_1 - I_3 &= c_1 \cos\left(\frac{\pi}{2} - \varphi_0 - v\alpha\right) - c_2 \sin\left(\frac{\pi}{2} - \varphi_0 - v\alpha\right) - c_1 \cos\left(\frac{3\pi}{2} - \varphi_0 - 3v\alpha\right) \\
&\quad + c_2 \sin\left(\frac{3\pi}{2} - \varphi_0 - 3v\alpha\right) \\
I_1 - I_3 &= 2 \sin\left(\frac{\pi}{2} + v\alpha\right) [c_1 \sin(\varphi_0 + 2v\alpha) - c_2 \cos(\varphi_0 + 2v\alpha)], \tag{4.8}
\end{aligned}$$

similarly

$$I_0 - I_2 = 2 \sin\left(\frac{\pi}{2} + v\alpha\right) [c_1 \cos(\varphi_0 + v\alpha) + c_2 \sin(\varphi_0 + v\alpha)], \tag{4.9}$$

and therefore

$$\varphi_{measured} = \tan^{-1} \left(\frac{c_1 \sin(\varphi_0 + 2v\alpha) - c_2 \cos(\varphi_0 + 2v\alpha)}{c_1 \cos(\varphi_0 + v\alpha) + c_2 \sin(\varphi_0 + v\alpha)} \right). \tag{4.10}$$

This equation is similar to an equation shown in (Lindner, 2010) to provide a more accurate view of the motion error based on the fact each frame is actually a summation of a large number of integrations.

Using the equation for the measured phase the relationship between the real phase and the axial motion error when using the VFRI technique can be investigated. This is plotted in Figure 4.40 for integrating over 1, 2 and 3 measurements. This figure demonstrates that the sinusoidal component of the axial motion error is not affected by the VFRI technique, the offset error is simply increased. As discussed in section 2.3.4.6, this offset

is not particularly significant and it can be showed that again this offset is caused by the selection of reference frame.

This analysis has assumed both that the modulation signals are perfect sinusoids and that the amplitude is consistent over the entire measurement. If a more accurate model of the axial motion error using this technique is desired these factors must be included, however that is outside the scope of this thesis.




Figure 4.40 Theoretical Axial motion error versus phase when integrating over 1,2 and 3 measurements

4.6.6 IMPROVED THRESHOLDING METRIC

The metric used to measure the quality of the phase measurements for the initial implementation of this technique was the amplitude. It is expected that the precision will be inversely proportional to the amplitude and therefore the amplitude is not a proportional threshold metric. Furthermore, other factors that influence the precision, such as the offset, are not constant over the scene or for individual points integrated over multiple measurements. To evaluate the use of the amplitude as a quality metric, simulations were performed to measure the relationship between the standard deviation and the amplitude for different correlation amplitudes measured over multiple measurements. The results are shown in Figure 4.41. These demonstrate that the amplitude is not a good measure of the precision of the measurements. For a particular measured amplitude value, the standard deviation can have a very large spread of

values. A better measure for the precision is required if a system is desired where all the pixels will have the same output precision.

In Chapter 2 it was shown that the precision should be proportional to \sqrt{B}/A , therefore this could provide a superior measure to the amplitude alone. This metric was previously avoided as it requires more FPGA resources to implement. The simulations were rerun for this metric and are shown in Figure 4.42. This demonstrates that this ratio is a much better indicator of the standard deviation. It was therefore decided to make a second implementation of the thresholding technique using the ratio \sqrt{B}/A as the quality




Figure 4.41 Simulated standard deviation versus measured Amplitude for various correlation amplitudes




Figure 4.42 Simulated standard deviation versus \sqrt{B}/A for various correlation amplitudes

metric. This will be referred to as precision thresholding.

For simplicity, measurements using the precision thresholding technique were performed using an SR-4000 camera and post-processing the data in Matlab. The SR-4000 camera provides benefits in terms of automatically adjusting for a number of optical effects. An FPGA implementation using this metric was developed and demonstrated on our system in Chapter 6. The processing is performed by recording 1000 measurements of the contrast apparatus described in Section 4.2 using a frame time of 5 ms. The data is then stepped through, accumulating numerator and denominator registers as would happen on camera until the threshold value is reached. The phase output is then changed and the accumulators are reset. The precision of the measurements is calculated by taking the standard deviation of the phase measurements. The number of times the phase output changes over the 1000 measurements is also recorded, so the effective frame rate can be calculated. These parameters are averaged over a 40×10 region for each of the contrast areas, due to the different field of view and resolution of the SR-4000 camera.

To provide a comparison, both precision thresholding and amplitude thresholding were performed on the same data. The precision is plotted against the threshold value for both metrics in Figure 4.43. The horizontal axis for the precision thresholding has been reversed, as this threshold requires the metric to be under a certain value as opposed to amplitude thresholding where it is required to be above a value. The frame rate for these thresholds is shown in Figure 4.44. Using the precision thresholding method the precision of the dark and bright regions does not start increasing, in other words multiple frames are not integrated over, until the precision of the sloped region has become equal to the region in question. Once all the regions have the same precision, the precision of all three regions increases at the same rate with changing threshold, although there is a small offset between them. In comparison, when using amplitude thresholding, the bright and dark regions start integrating over multiple regions long before the precision of the sloped area is equal to the other two regions. The precision thresholding method is able to provide a constant precision image across the entire scene.

When using precision thresholding the frame rate of less well illuminated areas of the scene is extremely low. Due to the inverse square relationship between the number of frames integrated over and the precision, there are diminishing returns for integrating over a large number of frames. Because of this it is not necessarily always the case that precision thresholding is a better technique to use, as opposed to amplitude thresholding. Depending on the application, amplitude thresholding could be used if constant precision is not required and the extremely low frame rates that can be present in precision thresholding are undesirable.




Figure 4.43 Standard Deviation versus threshold for the three regions using the SR400 camera with precision thresholding (top) and amplitude thresholding (bottom)




Figure 4.44 Effective frame rate versus threshold for the three regions using the SR400 camera with precision thresholding (top) and amplitude thresholding (bottom)

4.7 Summary

This chapter has outlined three methods for increasing the dynamic range of indirect time of flight cameras. Optimisation of the frame time for a particular scene can improve the quality of data being recorded. However, this method is still limited by the dynamic range of the camera for a particular frame time and can unnecessarily impact the frame rate of bright objects.

Capturing data using multiple different frame times was shown to be able to produce a very high quality range image over a wide variety of reflectivities. This method requires a

very long measurement time over the entire field of view and therefore is only suitable for static scenes. High Dynamic Range Imaging is also susceptible to systematic error, due to a change in the measured phase with changing measurement time.

Using a variable frame rate can improve the range data measured in areas that have low intensity returns, while not adversely affecting the frame rate of higher intensity areas of the scene. However, the performance gains using this method can be less than the gains from simply increasing the frame time, due to non-linearity observed in the offset of indirect time of flight cameras.

Two methods of thresholding were implemented, amplitude thresholding and precision thresholding. Amplitude thresholding was shown to provide some improvement in poorly illuminated areas of the scene with moderate impact on the frame rate while precision thresholding was shown to be able to provide an approximately constant precision range image across the scene, although it has a very significant impact on the frame rate.

The methods for improving the dynamic range of indirect time of flight cameras presented in this chapter all have different advantages and disadvantages. If the measurement time is not significant Variable Frame Rate Imaging using precision thresholding is generally preferable, as it allows you to select your desired precision and is not susceptible to error introduced due to changing the frame time. If faster measurements are required, the single frame data can be outputted in parallel using this method.

Chapter 5 PHASE DETECTION ALGORITHMS

5.1 Introduction

This chapter will analyse several phase detection algorithms that have potential benefits over the standard indirect time of flight algorithm that is currently used ubiquitously in both the literature and commercial cameras. To save significant time in FPGA programming, phase calculations for this chapter were performed in Matlab using raw intensity images captured by the Victoria University Range Imaging System. Each algorithm is derived, highlighting any assumptions it makes, and is tested both in its response to motion and its linearity. A comparison is then made between the different algorithms considered. An FPGA implementation of the best performing algorithm is developed and utilised on the Victoria University Range Imaging System. This algorithm is further examined for multiple velocities and compared to the standard algorithm in terms of precision.

Surrel's method of analysing phase detection algorithms (Surrel, 1996) using their characteristic polynomial will be used to make predictions of each algorithm's behaviour, this method of analysis is discussed further in section 5.2.

To qualitatively compare the algorithms in this chapter, reasonable metrics have to be selected. Ideally we would want the RMS error to be zero however, to prevent bias from the offset, which was shown to be due to the selection of reference frame in Chapter 2, standard deviation is used instead as a measure of quality of an algorithm. This provides the RMS error value independent of the mean value. The range of the motion error is also used to measure if algorithms have outliers, which could cause significant problems in real world measurements. To make observation of the error easier, the graphs of the motion error have had their mean value subtracted from them.

Most of the algorithms analysed in this chapter come from Phase Shifting Interferometry (PSI), a field separate to indirect time of flight range imaging but one where significant research has been performed on phase detection algorithms. The conventional PSI notation has been used in this chapter. The phase steps are normally written as $n = 1...N$, rather than the steps $n = 0...N-1$ used in Chapter 2.

A well studied problem in PSI is a linear miscalibration of the phase step, meaning the intensity values have the form (Surrel, 1993)

$$I_n = \cos((n-1)(1+\varepsilon)\delta - \varphi_1) \quad (5.1)$$

where δ is the desired phase step and ε is the linear miscalibration error. By substituting (2.25) into (2.24) and rearranging, the intensity for frame n with linear motion has the form

$$I_n = \cos\left((n-1)\left(1 - \frac{v\alpha}{\delta}\right)\delta - \varphi_1\right), \quad (5.2)$$

where α is the scaling factor described in Chapter 2. This is equivalent to a linear miscalibration of $-v\alpha/\delta$ (Drayton et al., 2012 a). In phase shifting interferometry the amplitude is normally not significant and is discarded. In indirect time of flight, as the object moves the amplitude will change due to the inverse square decrease with distance. This additional error source is small compared to the linear motion error and will be addressed in section 5.14.

Motion error is measured using the linear table apparatus described in Chapter 3. For each algorithm, 100 data runs are recorded along with 5 calibration runs. Calibration runs were performed by returning the object to its position at the start of each phase measurement, measured by the encoder, and recording a corresponding static measurement. The average of the calibration runs is subtracted from each data run and the resulting error is averaged over the 100 data runs. Error bars are used to indicate ± 1 standard deviation. Linearity measurements are recorded by advancing the target in 50 mm steps, taking 100 measurements per step and averaging them. This was done over a distance of 4 m.

Because motion error is dependent on the distance moved during a measurement it is therefore dependent on both the velocity and the measurement time. To keep the

motion error consistent between algorithms, the measurement time for each algorithm is maintained at 125 ms, regardless of the number of frames the algorithm has. Due to the readout time of 3197 μs included in each frame measurement, algorithms using a higher number of frames have a slight disadvantage in repeatability as they have been integrating over a shorter period of time.

5.2 Theoretical analysis of phase algorithms

As discussed in Chapter 2, a number of papers have been published in the time-of-flight field on the impact of axial motion error (Lottner et al., 2007) (Lindner & Kolb, 2009) (Foix et al., 2011) and on analysis of harmonic components of the correlation waveform (Rapp, 2007) (Foix et al., 2011) (Godbaz et al., 2011). However, these errors have been treated as separable and the a theoretical investigation of the impact of harmonics on the motion error has not been performed. In order to evaluate the axial motion error with harmonics, a method can be adapted from phase shifting interferometry. Surrel (Surrel, 1996) proposed a method for analysing phase algorithms using their characteristic polynomial. In this chapter we will use this analysis to predict the response of the algorithms tested. This section will outline the theory behind this approach and then state the rules that Surrel demonstrated indicate the performance of a particular algorithm.

The intensity of a particular frame as a function of the phase can be written as the sum of an exponential Fourier series

$$I(\varphi) = \sum_{m=-\infty}^{\infty} \alpha_m e^{im\varphi}, \quad (5.3)$$

where α_m is the complex Fourier coefficient of the m^{th} harmonic.

For phase stepping interferometry (and indirect time of flight range imaging) we introduce a phase step δ and use a number of frames k meaning the intensity of frame k can be written as

$$I_k(\varphi + k\delta) = \sum_{m=-\infty}^{\infty} [\alpha_m e^{im\varphi}] e^{imk\delta} \quad (5.4)$$

To find the measured phase φ_m with introduced phase step δ and M phase steps, phase algorithms have traditionally been written as the arctangent of the ratio of two linear combinations of phase steps with coefficients a_k and b_k

$$\varphi_m = \tan^{-1} \frac{\sum_{k=0}^{M-1} b_k I(\varphi + k\delta)}{\sum_{k=0}^{M-1} a_k I(\varphi + k\delta)} \quad (5.5)$$

this is equivalent to the measured phase φ_m being the argument of a complex linear combination

$$\varphi_m = \arg[S(\varphi)] \quad (5.6)$$

where

$$S(\varphi) = \sum_{k=0}^{M-1} c_k I(\varphi + k\delta), \quad (5.7)$$

with $c_k = a_k + ib_k$.

Substituting (5.4) into (5.7) gives

$$\begin{aligned} S(\varphi) &= \sum_{m=-\infty}^{\infty} \left\{ \alpha_m e^{im\varphi} \sum_{k=0}^{M-1} c_k [e^{im\delta}]^k \right\} \\ &= \sum_{m=-\infty}^{\infty} \{ \alpha_m e^{im\varphi} P(e^{im\delta}) \}, \end{aligned} \quad (5.8)$$

where $P(x)$ is a polynomial of degree $M-1$

$$P(x) = \sum_{k=0}^{M-1} c_k x^k, \quad (5.9)$$

which Surrel labels the characteristic polynomial of the algorithm as it can be used to predict the algorithm's behaviour. Surrel then outlines three rules which can be used, along with an algorithm's characteristic polynomial, to predict the behaviour of an algorithm with respect to harmonics and linear motion (in his paper it is referred to as phase shift miscalibration, as shown in section 5.1 these effects are equivalent).

Surrel's three rules are as follows (Surrel, 1996):

1. Insensitivity to the m^{th} harmonic present in the intensity signal can be achieved when the complex numbers $\exp(im\delta)$ (if $m \neq 1$) and $\exp(-im\delta)$ are roots of the

characteristic polynomial. In other words, the monomials $[x-\exp(im\delta)]$ and $[x-\exp(-im\delta)]$ must appear in the factorization of the characteristic polynomial. These monomials may be identical, depending on the values of m and δ .

2. Insensitivity to the m^{th} harmonic present in the signal ($m \neq 0$) is achieved in the presence of a phase-shift miscalibration when the two complex numbers $\exp(im\delta)$ (if $m \neq 1$) and $\exp(-im\delta)$ are double roots of the characteristic polynomial. In other words, the squared monomials $[x-\exp(im\delta)]^2$ and $[x-\exp(-im\delta)]^2$ must appear in the factorization of the characteristic polynomial. Depending on the values of m and δ , these monomials may be identical.
3. The algorithmic insensitivity to the m^{th} harmonic ($m \neq 0$) is achieved in the presence of a phase-shift miscalibration when the two complex numbers, $\exp(im\delta)$ and $\exp(-im\delta)$ are roots of the order of $k + 1$ of the characteristic polynomial. The phase measured will contain no term in ε^p , $p \leq k$, as a result of the presence of this harmonic.

ε is the miscalibration of the phase shift from (5.1).

The use of complex exponentials to represent the frequency means we must necessarily handle negative frequencies. It can be shown that the motion error described in traditional theory in Chapter 2 for sinusoidal modulation signals is equivalent to the error caused by the negative fundamental frequency. For the harmonic free case, $S(\varphi)$ is

$$S(\varphi) = \frac{1}{2} e^{i\varphi} \sum_{n=0}^{N-1} [e^{i\varphi} \gamma^{(1+\varepsilon)n} + e^{-i\varphi} \gamma^{-(1+\varepsilon)n}] \gamma^{-n}, \quad (5.10)$$

where $\gamma = \exp(2\pi i/N)$ (Surrel, 1993). Knowing that $\varepsilon = -\alpha/\delta$ the equivalent phase miscalibration error can be simulated. This provides equivalent results to those shown in Chapter 2. If the same simulations are run without the negative fundamental included, the resulting phase error is only an offset value. Both simulations are shown in Figure 5.1.




Figure 5.1 Simulated miscalibration error with and without the negative fundamental

Because we are using nominally square waves for modulation of both the sensor and the light sources, we expect that the resulting correlation signal will be a triangle wave and therefore contain only odd harmonics. However, the transfer function of the illumination source, the sensor, the modulation drivers and the low pass filter placed on the sensor modulation inputs all affect the harmonic content of the signal. Therefore, to investigate the actual harmonic response of the camera, raw intensity images were recorded while the relative phase of the emitted light and the sensor modulation was stepped 64 times over the 2π phase range. This was measured over several cycles and a typical cycle is shown in Figure 5.2.




Figure 5.2 Measured raw Intensity versus Sample Number using 64 steps per cycle

A Fast Fourier Transform was performed to investigate the harmonic content of the correlation waveform. Figure 5.3 shows the relative amplitude versus harmonic number using a logarithmic scale. There are noticeable harmonic peaks at both odd and even harmonics. Table 5.1 shows the relative amplitude of each of the harmonics to the fundamental. This shows that, as expected, the third harmonic is by far the strongest. However, it has significantly lower amplitude than is expected if our simplified square wave model was used, as a triangle wave has odd harmonics with $1/m^2$ amplitude where m is the harmonic number. There is a relatively strong second harmonic that was not expected using square wave modulation however it is approximately a third of the amplitude of the third harmonic. After the third harmonic, the amplitude of subsequent harmonics quickly reduces to being insignificant.




Figure 5.3 Relative Amplitude versus Frequency for the Victoria University Range Imaging System


We are generally using low numbers of phase steps for our measurements, for higher numbers acquiring sufficient light in a short time period is not possible for our camera. Measurement of the harmonic content of the Victoria University Range Imaging System shows that it is particularly important to be insensitive to the negative fundamental and the third harmonic as these will be the strongest undesired signals. The 4th and higher harmonics are unlikely to have any measurable impact on the phase.

Several simulations were run to give an indication of the error expected depending on the harmonic content of the correlation waveform, both with and without linear motion.

Table 5.1 List of the Relative Amplitudes of harmonics and their inverses in the Victoria University Range Imaging System

Harmonic	Relative Amplitude	Inverse of the Relative Amplitude
Fundamental	1	1
2 nd	0.00948	105
3 rd	0.03224	31
4 th	0.00200	500
5 th	0.00098	1020
6 th	0.00108	926
7 th	0.00069	1449

First, simulations were run to explore the effect of harmonics in the static case. Figure 5.4 shows the linearity error of two algorithms, one that is sensitive to the negative 3rd harmonic and the positive 5th harmonic and one that is insensitive to these but is sensitive to the negative 5th and positive 7th harmonic. The first algorithm exhibits a 4 cycle error, which is expected from the literature for a third harmonic (Payne et al., 2008). The second algorithm is sensitive to the negative 5th and positive 7th harmonics and therefore a 6 cycle error is observed. Generally an algorithm sensitive to the m^{th} harmonic will have an $m-1$ cycle error and an algorithm sensitive to the negative m^{th} harmonic will have an $m+1$ cycle error.

**Figure 5.4 Simulated linearity error of algorithms with different harmonic sensitivities**

Simulations were run to find the effect of removing motion error terms without the presence of harmonics, in other words increasing the number of roots at $m = -1$, these are shown in Figure 5.5. This demonstrates that motion error due to the negative fundamental is a two cycle error within the unambiguous measurement distance.

The next most significant contribution to motion error is expected to be the third harmonic. This should have the largest amplitude of the harmonics present in our system. The negative third harmonic is implicitly included in all the following analyses.




Figure 5.5 Comparison of motion error for different multiplicity of roots at $m = -1$




Figure 5.6 Comparison of the effect of motion on an algorithm with increasing multiplicity of roots at $m = 3$ and -3

Sensitivity to the third harmonic is dependent on the number of roots at $m = 3$ and -3 . Figure 5.6 shows simulations demonstrating the effect of increasing the multiplicity of these roots. The algorithms used for these simulations have double roots at $m = -1$. A four cycle error is observed, although the relative amplitude between the peaks is heavily distorted. The effect of being sensitive to the third harmonic is much larger than the error introduced by being sensitive to the second term of the negative fundamental and therefore first order insensitivity to the third harmonic is more important for our application. Our desired phase algorithm, within the bounds of using a relatively small number of frames, should have at least a double root at $m = -1, 3$ and -3 .

As we know the actual harmonic content of our system, we can simulate what the expected linear motion response will be for our system when using the standard algorithm. As shown in Figure 5.7, there is predominantly a two cycle error. However, there is some distortion in the response caused by the third harmonic.




Figure 5.7 Simulated linear miscalibration error for the standard algorithm using the measured harmonics for our system

5.3 Standard Algorithm (N-bucket)

5.3.1 BACKGROUND

Traditionally indirect time of flight cameras have used a set of algorithms referred to as N-bucket algorithms in Phase Shifting Interferometry. These algorithms use N phase steps of $2\pi/N$ with the first frame having zero offset. The phase is calculated as

$$\varphi = \tan^{-1} \frac{\sum_{i=1}^N I_i \sin\left((i-1) \frac{2\pi}{N}\right)}{\sum_{i=1}^N I_i \cos\left((i-1) \frac{2\pi}{N}\right)}. \quad (5.11)$$

This algorithm was derived in Chapter 2. Generally the four frame algorithm is used as this simplifies the equation to

$$\varphi = \tan^{-1} \frac{I_2 - I_4}{I_1 - I_3}. \quad (5.12)$$

5.3.2 CHARACTERISTIC POLYNOMIAL ANALYSIS

The characteristic polynomial of the four frame N -bucket algorithm is

$$\begin{aligned} P(x) &= -i(x-1)(x+1)(x+i) \\ &= -i(x-1)(x-e^{\pi i})(x-e^{3\pi i/2}). \end{aligned} \quad (5.13)$$

As the phase step is $\pi/2$ we can see the single root at $e^{\pi i}$ constitutes both the positive and negative roots of $m = 2$, meaning this algorithm should be insensitive to the second harmonic in the static case. To be insensitive to the third harmonic, $e^{3\pi i/2}$ and $e^{-3\pi i/2}$ terms are required, however only the positive term is present. There is a single root at $m = -1$, but not the double root required to improve the motion response of the algorithm.

5.3.3 LINEARITY MEASUREMENTS

The data in Figure 5.8, measured using the linear table, show a sinusoidal relationship between the residuals and the actual distance. The linear table is not long enough to measure the entire unambiguous range, however, two cycles are clearly seen between 1 m and 3.5 m, which is half of the unambiguous range. Therefore this is likely the expected four cycle error. The sum of the absolute value of the residuals for this algorithm is 1.0 ± 0.3 m.

To confirm that this is a four cycle error, a second method of measuring the linearity error was performed. Instead of moving an object axially, an object is placed stationary in the camera's field of view and an additional phase offset between the sensor and illumination modulation signals is stepped from 0 to 2π in $\pi/80$ increments. The

resulting linearity error is shown in Figure 5.9 along with the linearity error from moving an object through the scene. Apart from a phase shift caused by the initial offset of the object in the moving object data, and a loss of quality in the signal for the moving object data at long range, the two methods have the same result. It is confirmed that this is a four cycle error.




Figure 5.8 Measured linearity error for the standard four frame algorithm




Figure 5.9 Comparison of measured linearity error between measurement techniques for the 4 step standard algorithm

5.3.4 MOTION ERROR ANALYSIS

As predicted by both traditional theory, shown in Chapter 2, and analysis of its characteristic polynomial, the standard phase algorithm has a sinusoidal motion error, which is shown in Figure 5.10. The observed error appears to have two cycles within the range 0 to 2π , as expected from the negative fundamental. The standard deviation of the error with distance for this algorithm is 0.044 ± 0.002 m and the range is 0.112 ± 0.007 m. These values will be used as the comparison values for the algorithms tested in this chapter. A measurement time of 125 ms and a velocity of 2 m/s is used for these experiments. Due to the preliminary nature of these measurements, only a single velocity was used. Different velocities will simply change the amplitude of the sinusoidal error observed, as shown in Chapter 2. The best performing algorithm is profiled for a number of velocities and compared to the standard algorithm in Section 5.13.




Figure 5.10 Axial motion error versus distance for the standard four frame algorithm

To be able to provide a better comparison to the simulation shown in Figure 5.7, higher resolution data over a larger range is required. Similar to how stepping the initial phase can be used to acquire data for linearity measurements, the insertion of both a stepping in initial phase and a change in the size of the phase steps can be used to simulate the effect of motion with a static scene and acquire higher resolution data on the motion error for this algorithm. Instead of setting the phase step to $\pi/2$ it is set to $21\pi/40$. This is approximately equivalent to the phase error introduced in the motion experiments,




Figure 5.11 Linear miscalibration error for the four frame standard algorithm with a linear miscalibration of $\pi/40$

however the resolution of the phase steps of the PLL used in the Victoria University Range Imaging System means the exact equivalent motion error cannot be used. This method also does not incorporate the change of intensity with distance, which is investigated in more detail in Section 5.14. The result of this linear miscalibration of the phase stepping is shown in Figure 5.11. This result validates the simulation shown in Figure 5.7.

5.4 Five frame standard algorithm

5.4.1 BACKGROUND

It has been suggested (Jongenelen et al., 2009) that the five frame N bucket algorithm has superior linearity to the four frame N-bucket algorithm. Therefore the five frame N-bucket algorithm was also characterised. Using the five frame algorithm, the equation for φ is

$$\varphi = \tan^{-1} \frac{\sum_{i=2}^5 I_i \sin\left((i-1)\frac{2\pi}{5}\right)}{I_1 + \sum_{i=2}^5 I_i \cos\left((i-1)\frac{2\pi}{5}\right)}. \quad (5.14)$$

The five frame N bucket algorithm uses five phase steps of $2\pi/5$.

5.4.2 CHARACTERISTIC POLYNOMIAL ANALYSIS

The characteristic polynomial of the five step N bucket algorithm is

$$P(x) = (x - 1) \left(x - e^{\frac{6\pi i}{5}} \right) \left(x - e^{-\frac{6\pi i}{5}} \right) \left(x - e^{\frac{2\pi i}{5}} \right). \quad (5.15)$$

As the phase step is $2\pi/5$, the two terms $\left(x - e^{\frac{6\pi i}{5}} \right)$ and $\left(x - e^{-\frac{6\pi i}{5}} \right)$ indicate that this algorithm will be insensitive to the third harmonic in the static case. There are no roots at $m = -1$ and a single root at $m = 1$, therefore this algorithm is using the negative fundamental as the detection term, and we would not expect this algorithm to be insensitive to motion error due to the negative fundamental. Insensitivity to the fifth harmonic requires a root at $e^{\frac{10\pi i}{5}} = e^{2\pi i} = e^0 = 1 = e^{-\frac{10\pi i}{5}}$ therefore the root at $x = 1$ provides insensitivity to the fifth harmonic in the static case. The roots at $m = 3$ and -3 are also roots for $m = 7$ and -7 . The root for $m = 1$ is also a root for $m = -9$ and 11 , however the roots for $m = 9$ and -11 are not present. As the negative fundamental is being used for the detection term, the normal rule for harmonics is reversed and both of these harmonics are expected to produce a 10 cycle error.

5.4.3 LINEARITY MEASUREMENTS

As predicted by analysis of this algorithm's characteristic polynomial, the four cycle error observed for the four frame algorithm is not present in this algorithm. The sum of the absolute value of the residuals is 0.5 ± 0.3 m. As expected, since the four cycle error evident in the standard four frame algorithm has been removed, the sum of the




Figure 5.12 Measured linearity error for the standard five frame algorithm

absolute value of the linear residuals is significantly smaller than for the standard four frame algorithm. Theory would suggest the next harmonic present in the resulting signal should be the negative ninth, causing a ten cycle error. However, experimental results shown in Figure 5.12 indicate this is not the case. Instead the error appears to be a two cycle error. The error has a number of unexpected features. Both the amplitude and the period of the signal appear to vary with distance. For harmonic error sources both the period and amplitude should be constant, therefore this error is not likely to have a harmonic source.

To confirm that the remaining linearity error is not coming from a harmonic source, linearity measurements were taken using a 64 phase step N-bucket algorithm. The high number of samples per period of this algorithm should make this measurement insensitive to any harmonics in the system. The linearity error is shown in Figure 5.13. This has similar non-linearity error to the five frame algorithm, supporting the argument that this is a non-harmonic error.




Figure 5.13 Measured linearity error for the 64 frame N bucket algorithm

To provide more data, the initial phase stepping method for measuring linearity error was used for the five phase step algorithm. Since the object does not move using this method, the amplitude across the range is constant. Figure 5.14 shows that the oscillatory error that was observed using the moved object method is not observed using the initial phase stepping method. The remaining error is expected to have a ten

cycle error, however no periodic waveform is observed. This can be explained by the harmonic analysis of our system in section 5.2, which shows that the amplitude of the 9th harmonic is much much smaller than the amplitude of the fundamental. It therefore is not strong enough to produce the expected ten cycle error above the noise floor of the camera.




Figure 5.14 Comparison of Linearity Error between measurement techniques for the 5 step standard algorithm

There are several potential sources for this error. As it is only observed with the moved object there is potentially multi-path interference occurring (Dorrington et al., 2011). Multi-path errors present a significant challenge for the field of indirect range imaging, however they are a separate problem from the motion errors that we are trying to address using phase detection algorithms. Another potential source for this error is the change in intensity as the object is moved, due to the inverse square drop in intensity with distance. It was decided to characterise the PMD19K sensor's response with changes in intensity. An object was placed stationary at a set distance and the frame time was increased from 8 ms to 50 ms. The results, shown in Figure 5.15, demonstrate that the intensity changes non-linearly with integration time, which is highly non-desirable. This was identified previously in Chapter 4 for the four frame algorithm. By allowing for non-linearity in the phase with changing intensity, a new model can be formed. Simulations, shown in Figure 5.16, demonstrate that non-linearity in the intensity can cause a linearity error in the phase measurement. It also displays the

features of the linearity error found in Figure 5.12, allowing for oscillations with non-uniform period and non-uniform amplitude.




Figure 5.15 Intensity versus Integration time for all four frames for the Victoria University Range Imaging System




Figure 5.16 Simulated error versus phase for non-linear pixels in an indirect time of flight camera

Measurements of the intensity with increasing integration time were performed using a SR-4000 commercial camera from Mesa Imaging in Figure 4.32. It was shown that while there was non-linearity in the offset of the correlation waveform, this camera did not experience a change in phase with changing amplitude. Therefore, as this sensor non-linearity is not a problem for the field of indirect time of flight imaging as a whole, and is instead an undesirable feature of our hardware, it does not need to be solved by the

phase detection algorithms being tested in this chapter. If the source is a multi-path error, again phase detection algorithms are not the solution for these errors. Potential solutions to multipath errors from the literature are discussed in section 2.3.4.5. For the remainder of this chapter, simulations will be used to show the theoretical linearity error for algorithms for which the primary error source is not harmonic.

5.4.4 MOTION ERROR ANALYSIS

Figure 5.17 shows the axial motion error of the five frame N bucket algorithm. As predicted, this algorithm experiences two cycle error due to the negative fundamental. The standard deviation of the error with distance for this algorithm is 0.042 ± 0.002 m and the range is 0.110 ± 0.005 m. This algorithm has a slightly lower standard deviation and range than the four frame standard algorithm, however the difference is negligible.




Figure 5.17 Axial motion error versus distance for the standard five frame algorithm

5.5 Order changed N-bucket

5.5.1 BACKGROUND

One method for improving the response of the system to motion is to change the order in which the frames are recorded (Drayton et al., 2011). If the order of the phase steps is changed from the traditional $0, \pi/2, \pi, 3\pi/2$. The actual phase at frame n is now

$$\varphi_n = \varphi_0 + m_n v \alpha \quad (5.16)$$

where m_n is an integer representing the order the frame was recorded from the set {0 1 2 3}, v is the velocity and α is a scaling coefficient including the frame time and modulation frequency to convert the velocity to the distance moved in one frame in radians. The relationship between φ_0 , the phase at the start of the measurement, and φ_m , the phase that is measured, can be generalised by following the same process as shown in Chapter 2 from

$$\varphi_0 = \tan^{-1} \frac{\tan(\varphi_m) \cos(2v\alpha) + \tan(\varphi_m) - \sin(v\alpha) - \sin(3v\alpha)}{\cos(v\alpha) + \cos(3v\alpha) + \tan(\varphi_m) \sin(2v\alpha)} \quad (5.17)$$

to

$$\varphi_0 = \tan^{-1} \frac{\tan(\varphi_m) [\cos(m_0 v \alpha) + \cos(m_2 v \alpha)] - \sin(m_1 v \alpha) - \sin(m_3 v \alpha)}{\cos(m_1 v \alpha) + \cos(m_3 v \alpha) + \tan(\varphi_m) [\sin(m_0 v \alpha) + \sin(m_2 v \alpha)]} \quad (5.18)$$

Figure 5.18 shows theoretical plots of different frame orders. Only three orderings are shown, all other permutations produce the same results to these but with different phase shifts. By shifting the position in the arctangent ratio where the additional phase error occurs, the relationship between the measured phase and the reference frame phase changes. While the offset error remains unchanged, the change in error with distance can be greatly reduced by a careful selection of sampling order. This is only reducing the effect of the negative fundamental, it assumes that no harmonics are present in the signal.




Figure 5.18 Theoretical motion error versus distance for selected orderings of frame offset (Drayton et al., 2011)

Using the simulations in Figure 5.18 we can see it would be better to use the frame order 2-0-1-3 as this should decrease the effect of the negative fundamental. However, again we need to consider the effect of harmonics. The same simulations including the third harmonic are shown in Figure 5.19. This shows that the improvements made by the algorithm for the sinusoidal modulation case are almost entirely negated by the harmonics present in the system, with only a minor improvement realised.




Figure 5.19 Comparison of motion error for different orderings of the standard algorithm with the third harmonic included

5.5.2 CHARACTERISTIC POLYNOMIAL

Because this method does not use a well ordered phase step it is not suitable to be analysed using the characteristic polynomial method. However, sufficient theory is available in section 5.5.1 to make predictions about the response of this algorithm. In the static case, $v = 0$ and therefore $\varphi_0 = \varphi_m$. Therefore, the static response should be the same as the static response for the standard algorithm. As shown in section 5.5.1, motion error due to the negative fundamental is improved with this algorithm, however, due to harmonics, only minor improvements are expected.

5.5.3 LINEARITY MEASUREMENTS

Figure 5.20 shows the linearity measurements using the order changed four frame algorithm. As predicted, the results are similar to the standard algorithm, however there

is a phase shift in the measurements and therefore the residuals. The sum of the absolute value of the residuals is 1.2 ± 0.3 m, this is slightly larger than for the standard four frame algorithm but agrees with the value within uncertainty. This algorithm has therefore not improved the linearity of the camera. As we are not measuring over the full unambiguous range of the system a phase shift can change the value of the sum of the residuals.




Figure 5.20 Linearity error for the order changed four frame algorithm

5.5.4 MOTION ERROR ANALYSIS

Figure 5.21 shows the axial motion error versus distance for the order changed four frame algorithm. The standard deviation of the error with distance for this algorithm is 0.042 ± 0.002 m and the range is 0.108 ± 0.005 m. The standard deviation and range have both decreased from the standard algorithm, however, the improvement is not significant. As predicted this algorithm has similar motion error to the standard algorithm. Since this technique has only considered the negative fundamental, there is still a significant error introduced by the unaccounted for harmonics, particularly the third harmonic.




Figure 5.21 Axial motion error versus distance for the order changed four frame algorithm

5.6 Carré Algorithm

5.6.1 BACKGROUND

Carré (Carré, 1966) proposed that the phase step could be treated as an unknown instead of a set value. The Carré algorithm uses four phase steps and assumes a constant, but unknown, phase shift of 2α between each frame. The steps are -3α , $-\alpha$, α , 3α . The phase can then be calculated as (Schreiber & Bruning, 2006)

$$I_1 = A \cos(\varphi - 3\alpha) + B = A[\cos \varphi \cos 3\alpha + \sin \varphi \sin 3\alpha] + B$$

$$I_2 = A \cos(\varphi - \alpha) + B = A[\cos \varphi \cos \alpha + \sin \varphi \sin \alpha] + B$$

$$I_3 = A \cos(\varphi + \alpha) + B = A[\cos \varphi \cos \alpha - \sin \varphi \sin \alpha] + B$$

$$I_4 = A \cos(\varphi + 3\alpha) + B = A[\cos \varphi \cos 3\alpha - \sin \varphi \sin 3\alpha] + B$$

$$(I_1 - I_4) + (I_2 - I_3) = 2A \sin \varphi (\sin \alpha + \sin 3\alpha)$$

$$(I_2 + I_3) - (I_1 + I_4) = 2A \cos \varphi (\cos \alpha - \cos 3\alpha)$$

$$\begin{aligned} \tan \varphi &= \frac{(I_1 - I_4) + (I_2 - I_3) (\cos \alpha - \cos 3\alpha)}{(I_2 + I_3) - (I_1 + I_4) (\sin \alpha + \sin 3\alpha)} \\ &= \frac{(I_1 - I_4) + (I_2 - I_3) (1 - \cos^2 \alpha)}{(I_2 + I_3) - (I_1 + I_4) (\tan \alpha - \tan \alpha \sin^2 \alpha)} \\ &= \frac{(I_1 - I_4) + (I_2 - I_3) \sin^2 \alpha}{(I_2 + I_3) - (I_1 + I_4) \tan \alpha (1 - \sin^2 \alpha)} \\ &= \tan \alpha \frac{(I_1 - I_4) + (I_2 - I_3)}{(I_2 + I_3) - (I_1 + I_4)} \end{aligned} \quad (5.19)$$

The four intensity equations can be solved to find $\tan \alpha$ using standard trigonometric identities using the following steps

$$\begin{aligned}
 I_1 - I_4 &= 2A \sin \varphi \sin 3\alpha \\
 I_2 - I_3 &= 2A \sin \varphi \sin \alpha \\
 \sin \alpha &= \frac{I_2 - I_3}{2A \sin \varphi} \\
 &= \frac{I_2 - I_3}{I_1 - I_4} \sin 3\alpha \\
 &= \frac{I_2 - I_3}{I_1 - I_4} [3\sin \alpha - 4\sin^3 \alpha] \\
 \frac{I_1 - I_4}{I_2 - I_3} &= 3 - 4 \sin^2 \alpha \\
 \sin^2 \alpha &= \frac{3(I_2 - I_3) - (I_1 - I_4)}{4(I_2 - I_3)} \\
 \tan^2 \alpha &= \frac{3(I_2 - I_3) - (I_1 - I_4)}{4(I_2 - I_3) \cos^2 \alpha} \\
 \cos^2 \alpha &= 1 - \sin^2 \alpha = \frac{(I_2 - I_3) + (I_1 - I_4)}{4(I_2 - I_3)} \\
 \tan^2 \alpha &= \frac{3(I_2 - I_3) - (I_1 - I_4)}{(I_2 - I_3) + (I_1 - I_4)} \\
 \tan \alpha &= \sqrt{\frac{3(I_2 - I_3) - (I_1 - I_4)}{(I_2 - I_3) + (I_1 - I_4)}} \quad (5.20)
 \end{aligned}$$

substituting (5.20) into (5.19) gives the final phase as

$$\begin{aligned}
 \tan \varphi &= \tan \alpha \frac{(I_1 - I_4) + (I_2 - I_3)}{(I_2 + I_3) - (I_1 + I_4)} \\
 &= \frac{(I_1 - I_4) + (I_2 - I_3)}{(I_2 + I_3) - (I_1 + I_4)} \sqrt{\frac{3(I_2 - I_3) - (I_1 - I_4)}{(I_2 - I_3) + (I_1 - I_4)}} \\
 &= \left(\frac{[3(I_2 - I_3) - (I_1 - I_4)][(I_1 - I_4) + (I_2 - I_3)]^2}{(I_2 - I_3) + (I_1 - I_4)} \right)^{1/2} \frac{1}{(I_2 + I_3) - (I_1 + I_4)} \\
 \varphi &= \tan^{-1} \frac{\sqrt{[3(I_2 - I_3) - (I_1 - I_4)][(I_1 - I_4) + (I_2 - I_3)]}}{(I_2 + I_3) - (I_1 + I_4)} \quad (5.21)
 \end{aligned}$$

Because the square root in the numerator can only take positive values, the output of the Carré algorithm is not of the range 0 to 2π . A simulation demonstrating the measured phase versus real phase is shown in Figure 5.22 for the harmonic free case. This means a transform must be done to correct the calculated phase within the

ambiguity distance. This algorithm provides an exact algebraic solution for the phase insensitive to motion error. However, it assumes that the correlation waveform is a perfect sinusoid.




Figure 5.22 Simulated relationship between measured phase and actual phase using Carré's algorithm

5.6.2 CHARACTERISTIC POLYNOMIAL ANALYSIS

Due to the square root in the numerator of Carré's algorithm, instead of using characteristic polynomial analysis, this algorithm was modelled in Matlab to predict its behaviour. Figure 5.23 shows the theoretical linearity error for Carré's algorithm. This predicts that this algorithm is sensitive to the third harmonic in the static case and will




Figure 5.23 Theoretical linearity error for Carré's algorithm

therefore have a four cycle error.

Figure 5.24 shows the theoretical response versus distance with linear motion. Two effects are visible here. Firstly, this algorithm is sensitive to the third harmonic and therefore a four cycle error is apparent. Secondly, at the phase unwrapping points caused by the square root in the denominator, the error switches polarity.




Figure 5.24 Theoretical motion error versus distance for Carrés algorithm




Figure 5.25 Linearity error for Carrés algorithm

5.6.3 LINEARITY MEASUREMENTS

Figure 5.25 shows the linearity measurements for the Carré algorithm. As predicted by the simulation in section 5.6.2 this algorithm is sensitive to the third harmonic and therefore a four cycle error is observed. The sum of the absolute value of the residuals is

1.1 ± 0.3 m. This value agrees with the value for the standard algorithm within uncertainty and therefore this algorithm has not improved the linearity error of the camera.

5.6.4 MOTION ERROR ANALYSIS

Figure 5.26 shows the axial motion error for Carré's algorithm. One of the phase unwrapping points predicted using the Matlab simulation is clearly visible. Because of the limited range over which valid data can be acquired, and the step change due to the phase wrapping point, it is difficult to determine if the four cycle error is present. The standard deviation of the error with distance for this algorithm is 0.267 ± 0.002 m and the range is 0.596 ± 0.009 m. Clearly the step change caused by the unwrapping point dominates the error for this algorithm and the metrics therefore are not particularly useful.




Figure 5.26 Axial motion error versus distance for Carré's Algorithm

It is possible to use a secondary phase calculation to deal with the phase unwrapping points. The intensity measurements from Carré's algorithm can also be used in the four frame N bucket algorithm. This can identify the region of phase that the object is in and therefore the phase can theoretically be unwrapped in real time. It is therefore interesting to investigate the error if the phase wrapping point was removed. This data was manually unwrapped in Matlab and the results are shown in Figure 5.27. The

standard deviation of the error with distance is now 0.029 ± 0.002 m and the range is 0.078 ± 0.007 m. Both the standard deviation and the range have now decreased significantly from the standard algorithm. It appears to have a distorted four cycle error.




Figure 5.27 Axial motion error versus distance for Carré's algorithm using phase unwrapping

5.7 Hariharan Algorithm

5.7.1 BACKGROUND

Hariharan (Hariharan, 1987) showed a similar approach to Carré using five phase steps of some unknown value α . Again, the amplitude is assumed to be constant across the measurement and the signals are assumed to be perfect sinusoids. The intensity frames are of the form

$$\begin{aligned}
 I_i &= A \cos(\varphi + (i - 3)\alpha) \\
 I_2 - I_4 &= A[\cos(\varphi - \alpha) - \cos(\varphi + \alpha)] \\
 &= 2A \sin \varphi \sin \alpha \\
 2I_3 - I_5 - I_1 &= 2A \cos \varphi - A \cos \varphi \cos 2\alpha - A \cos \varphi \cos 2\alpha \\
 &= 2A \cos \varphi (1 - \cos 2\alpha) \\
 &= 4A \cos \varphi \sin^2 \alpha \\
 \frac{2A \sin \varphi \sin \alpha}{4A \cos \varphi \sin^2 \alpha} &= \frac{I_2 - I_4}{2I_3 - I_5 - I_1} \\
 \frac{\tan \varphi}{2 \sin \alpha} &= \frac{I_2 - I_4}{2I_3 - I_5 - I_1} \quad (5.22)
 \end{aligned}$$

It is clear that the choice of the phase step α will have a significant effect on the measurement. As $\alpha = \alpha_{set} + \alpha_{motion}$ we therefore want to choose an α_{set} such that small changes in α caused by motion have a minimal effect on the range measurement. This is achieved when the derivative with respect to α

$$\frac{d}{d\alpha} \left(\frac{\tan \varphi}{2 \sin \alpha} \right) = \frac{-\cos \alpha \tan \varphi}{2 \sin^2 \alpha} \quad (5.23)$$

is equal to zero, yielding a phase step of $\pi/2$. The final phase algorithm is therefore

$$\varphi = \tan^{-1} \frac{2(I_2 - I_4)}{2I_3 - I_5 - I_1} \quad (5.24)$$

5.7.2 CHARACTERISTIC POLYNOMIAL ANALYSIS

The characteristic polynomial of Hariharan's algorithm is

$$\begin{aligned} P(x) &= -(x-1)(x+1)(x+i)^2 \\ &= -(x-e^0)(x-e^{\pi i})(x-e^{3\pi i/2})^2. \end{aligned} \quad (5.25)$$

This algorithm has a phase step of $\pi/2$ and therefore the double root at $e^{-i\pi/2}$ means this algorithm should not contain the first term of ε^p in φ from the negative fundamental due to linear motion. For the static case this algorithm does not provide insensitivity to the third harmonic as $e^{-3\pi i/2}$ does not occur in its polynomial.




Figure 5.28 Linearity error for Hariharan's algorithm

5.7.3 LINEARITY MEASUREMENTS

Figure 5.28 shows the linearity measurements for Hariharan's algorithm. As predicted using the characteristic polynomial, this method is sensitive to the third harmonic and therefore a four cycle error is observed. The sum of the absolute value of the residuals is 1.0 ± 0.3 m, this value agrees with the value for the standard algorithm within uncertainty and therefore this algorithm has not improved the linearity of the camera.

5.7.4 MOTION ERROR ANALYSIS

Figure 5.29 shows the axial motion error for Hariharan's algorithm. The standard deviation of the error with distance for this algorithm is 0.028 ± 0.002 m and the range is 0.083 ± 0.008 m. This algorithm shows a significant improvement in both standard deviation and range over the standard algorithm and has approximately the same error as Carré's algorithm. This is expected as, since this algorithm is insensitive to the first term of the negative fundamental, the remaining error should be primarily due to the third harmonic. Slightly more than a single cycle is observed in the data recorded, which fits with the expected four cycle error.




Figure 5.29 Axial motion error versus distance for Hariharan's algorithm

5.8 N + 1 Type B Algorithm

5.8.1 BACKGROUND

Attempts to develop generic Phase Shifting Interferometry algorithms for an arbitrary number of frames (Larkin & Oreb, 1992) (Surrel, 1993) yield a class of algorithms called N+1 algorithms. The derivation for this class of algorithms uses a set of phase steps from $n = 0 \dots N-1$ instead of $n = 1 \dots N$. This will be adjusted in the final equation to make this algorithm's notation consistent with the other algorithms described in this chapter. There are two types of N+1 algorithms, type A and type B, the difference between the two is the weighting of the first and last frames. Type B are generally considered to be useful for correcting phase step errors (Schreiber & Bruning, 2006) so they will be the focus of this section.

Σ , the first component of the Discrete Fourier Transform (DFT) of the set of intensity measurements I_0 to I_{N-1} can be expressed as

$$\begin{aligned}\Sigma &= \sum_{n=0}^{N-1} I_n e^{-\frac{2\pi n}{N}i} \\ \delta &= e^{\frac{2\pi i}{N}} \\ \Sigma &= \sum_{n=0}^{N-1} I_n \delta^{-n}.\end{aligned}\tag{5.26}$$

The phase is therefore

$$\varphi = \arg(\Sigma).\tag{5.27}$$

Each phase step is nominally $2\pi n/N$ however we allow for linear motion by introducing an error ε giving the intensity for phase step n as

$$\begin{aligned}I_n &= \cos\left(\varphi + 2\pi(1 + \varepsilon)\frac{n}{N}\right) \\ &= \frac{1}{2}\left[e^{i\varphi} \delta^{(1+\varepsilon)n} + e^{-i\varphi} \delta^{-(1+\varepsilon)n}\right].\end{aligned}\tag{5.28}$$

The non-essential offset value has been discarded and the amplitude is assumed to be 1. Substituting (5.28) into (5.26) and using φ_m as the phase calculated by the algorithm gives

$$\begin{aligned}
\Sigma &= |\Sigma| e^{i\varphi_m} \\
\Sigma &= \frac{1}{2} \sum_{n=0}^{N-1} [e^{i\varphi} \delta^{(1+\varepsilon)n} + e^{-i\varphi} \delta^{-(1+\varepsilon)n}] \delta^{-n} \\
\Sigma &= \frac{1}{2} e^{i\varphi} \sum_{n=0}^{N-1} [\delta^{\varepsilon n} + e^{-2i\varphi} \delta^{-(2+\varepsilon)n}].
\end{aligned} \tag{5.29}$$

Using $\Delta\varphi = \varphi_m - \varphi$ it therefore follows that

$$\begin{aligned}
2|\Sigma| e^{i\Delta\varphi} &= 2|\Sigma| e^{i\varphi_m - i\varphi} \\
&= 2|\Sigma| \frac{e^{i\varphi_m}}{e^{i\varphi}} \\
&= \sum_{n=0}^{N-1} [\delta^{\varepsilon n} + e^{-2i\varphi} \delta^{-(2+\varepsilon)n}].
\end{aligned} \tag{5.30}$$

Summing the geometric series and expanding the Taylor series up to the first power of ε gives

$$2|\Sigma| e^{i\Delta\varphi} \approx N + i\pi(N-1)\varepsilon + \pi\varepsilon \frac{\delta e^{-2i\varphi}}{\sin\left(\frac{2\pi}{N}\right)}, \tag{5.31}$$

therefore

$$\Delta\varphi = \frac{N-1}{N} \pi\varepsilon - \frac{\pi\varepsilon}{N \sin\left(\frac{2\pi}{N}\right)} \sin\left(2\varphi - \frac{2\pi}{N}\right). \tag{5.32}$$

We want to suppress the linear term in $\varepsilon \exp(-2i\varphi)$, hence removing the two cycle error observed due to aliasing of the negative fundamental. Fortunately a similar term can be found in the difference between the two frames $n = N$ and $n = 0$

$$\begin{aligned}
I_N - I_0 &= \cos(\varphi + 2\pi\varepsilon) - \cos \varphi \\
&= \frac{1}{2} [e^{(-\varphi - 2\pi\varepsilon)i} + e^{(\varphi + 2\pi\varepsilon)i} - e^{\varphi i} - e^{-\varphi i}] \\
&= \frac{1}{2} [e^{\varphi i} \{e^{2\pi\varepsilon} - 1\} + e^{-\varphi i} \{e^{-2\pi\varepsilon} - 1\}] \\
&\approx \frac{1}{2} [e^{\varphi i} \{2\pi\varepsilon i\} + e^{-\varphi i} \{-2\pi\varepsilon i\}] \\
&\approx \pi\varepsilon i e^{\varphi i} [1 - e^{-2\varphi i}].
\end{aligned} \tag{5.33}$$

Therefore, if we add the difference between the two, scaled by coefficient A , we can mitigate the term above. Substituting this into the equation for Σ the last term becomes

$$\pi\epsilon \frac{\delta e^{-2i\varphi}}{\sin\left(\frac{2\pi}{N}\right)} - 2Ai\pi\epsilon e^{-2i\varphi}. \quad (5.34)$$

Equating this to zero allows us to solve for A

$$\begin{aligned} 2Ai\pi\epsilon e^{-2i\varphi} &= \pi\epsilon \frac{\delta e^{-2i\varphi}}{\sin\left(\frac{2\pi}{N}\right)} \\ A &= -i \frac{1}{2} \frac{\delta}{\sin\left(\frac{2\pi}{N}\right)} \\ A &= \frac{-i \left(\cos\left(\frac{2\pi}{N}\right) + i \sin\left(\frac{2\pi}{N}\right) \right)}{2 \sin\left(\frac{2\pi}{N}\right)} \\ &= \frac{\sin\left(\frac{2\pi}{N}\right) - i \cos\left(\frac{2\pi}{N}\right)}{2 \sin\left(\frac{2\pi}{N}\right)} \\ &= \frac{1}{2} \left[1 - \cot\left(\frac{2\pi}{N}\right) \right]. \end{aligned} \quad (5.35)$$

Since there is already a factor of 1 for I_0 , the total factors for I_0 and I_N are $1 - A$ and A respectively. The equation for φ then becomes

$$\begin{aligned} \tan \varphi &= \frac{\text{Im}\{(1 - A)I_0 + AI_N\} - \sum_{n=1}^{N-1} I_n \sin\left(\frac{2\pi n}{N}\right)}{\text{Re}\{(1 - A)I_0 + AI_N\} + \sum_{n=1}^{N-1} I_n \cos\left(\frac{2\pi n}{N}\right)} \\ &= \frac{\frac{I_0 - I_N}{2} \cot\left(\frac{2\pi}{N}\right) - \sum_{n=1}^{N-1} I_n \sin\left(\frac{2\pi n}{N}\right)}{\frac{I_0 + I_N}{2} + \sum_{n=1}^{N-1} I_n \cos\left(\frac{2\pi n}{N}\right)}. \end{aligned} \quad (5.36)$$

Changing this to use the same notation as previously (starting at frame $n = 1$) this algorithm becomes

$$\varphi = \tan^{-1} \frac{\frac{I_1 - I_{N+1}}{2} \cot\left(\frac{2\pi}{N}\right) - \sum_{n=2}^N I_n \sin\left(\frac{2\pi(n-1)}{N}\right)}{\frac{I_1 + I_{N+1}}{2} + \sum_{n=2}^N I_n \cos\left(\frac{2\pi(n-1)}{N}\right)}. \quad (5.37)$$

For $N = 4$ this gives the Hariharan Algorithm, but it also provides a new solution for $N = 3$ and for higher values of N . The four step type B $N+1$ algorithm is

$$\varphi = \tan^{-1} \frac{-I_1 - 3I_2 + 3I_3 + I_4}{\sqrt{3}(I_1 - I_2 - I_3 + I_4)}. \quad (5.38)$$

5.8.2 CHARACTERISTIC POLYNOMIAL ANALYSIS

The four step type B N+1 algorithm has the characteristic polynomial

$$\begin{aligned} P(x) &= -\sqrt{3}i(x-1)\left(x - \left(-\frac{1}{2} - \frac{\sqrt{3}}{2}i\right)\right)^2 \\ &= -\sqrt{3}i(x-1)\left(x - e^{-i\frac{2\pi}{3}}\right)^2. \end{aligned} \quad (5.39)$$

As the phase step is $2\pi/3$, for the third harmonic $m\delta = 6\pi/3 = 2\pi = 0 = -m\delta$. Therefore the single root at $x = 1$ means this algorithm will be insensitive to the third harmonic, but only in the static case. There is a double root at $m = -1$, therefore this algorithm should not contain the first term of ϵ^p in φ_m due to the negative fundamental. The double root at $m = -1$ is also a double root for $m = 5$, however there is no corresponding root for $m = -5$ therefore we would expect there to be a 6 cycle linearity error for this algorithm.

It is interesting to analyse the Type A algorithm to test the literature's assertion that Type B algorithms are more suitable for dealing with linear miscalibration error. The characteristic polynomial for the four step Type A N+1 algorithm is

$$\begin{aligned} P(x) &= (x-1)\left(x - \left(\frac{1}{2} - \frac{\sqrt{3}}{2}i\right)\right)\left(x - \left(-\frac{1}{2} - \frac{\sqrt{3}}{2}i\right)\right) \\ &= (x-1)\left(x - e^{i\frac{5\pi}{3}}\right)\left(x - e^{-i\frac{2\pi}{3}}\right). \end{aligned} \quad (5.40)$$

Similar to the Type B algorithm the root at $x = 1$ means this algorithm will be insensitive to the third harmonic in the static case. What was previously a double root at $m = -1$ is now only a single root. This means this algorithm will no longer have first order insensitivity to motion error due to the negative fundamental, confirming the result from literature.

5.8.3 LINEARITY MEASUREMENTS

Figure 5.30 shows the linearity measurements for the four frame N+1 algorithm. As predicted, the four frame N+1 algorithm is insensitive to the third harmonic. There remains a two cycle error with a sum of the absolute value of the residuals of

0.3 ± 0.3 m. This has a similar linearity error as discussed in depth for the five frame N bucket algorithm in section 5.4.3, which is also insensitive to the third harmonic.




Figure 5.30 Linearity error for the four frame type B $N+1$ algorithm

Simulations were performed to find the linearity error of this algorithm if the non-harmonic error was solved. The results are shown in Figure 5.31. As predicted, a 6 cycle error is observed.




Figure 5.31 Theoretical Linearity Error versus phase for the $N+1$ algorithm

5.8.4 MOTION ERROR ANALYSIS

Figure 5.32 shows the axial motion error for the four frame N+1 algorithm. The standard deviation of the error with distance for this algorithm is 0.041 ± 0.004 m and the range is 0.11 ± 0.01 m. This exhibits what appears to be a two cycle error. A four cycle error is expected as the N+1 algorithm is insensitive to the first ϵ^p term of the negative fundamental but is only insensitive to the third harmonic in the static case. Due to the low resolution and range of the data recorded, it is possible that the distortion of the amplitude of the peaks makes them undetectable in this data. This algorithm has shown an improvement over the standard algorithm, however both Hariharan's algorithm and Carré's algorithm showed superior responses.




Figure 5.32 Axial motion error versus distance for the four frame N+1 algorithm

5.9 Novak's Algorithm

5.9.1 BACKGROUND

A series of 5 step algorithms insensitive to linear phase shift errors were proposed by Novak (Novák et al., 2008). The best performing algorithm found in their paper used five phase steps of $\pi/2$. The derivation for this algorithm is similar to that for Hariharan's algorithm discussed in section 5.7.1, the difference is in how the $\sin \alpha$ term is handled in (5.22). Previously an alpha was chosen such that $d/d\alpha$ was relatively flat. Novak instead used the five equations for the intensities I_n to find an algebraic solution for α

$$\begin{aligned}
I_5 - I_1 &= A(\cos(\varphi + 2\alpha) - \cos(\varphi - 2\alpha)) \\
&= A(\cos \varphi \cos 2\alpha - \sin \varphi \sin 2\alpha - \cos \varphi \cos 2\alpha + \sin \varphi \sin 2\alpha) \\
&= -2A(\sin \varphi \sin 2\alpha) \\
&= -4A(\sin \varphi \sin \alpha \cos \alpha) \\
I_4 - I_2 &= A(\cos(\varphi + \alpha) - \cos(\varphi - \alpha)) \\
&= -2A(\sin \varphi \sin \alpha) \\
\frac{I_5 - I_1}{I_4 - I_2} &= \frac{2 \sin \varphi \sin \alpha \cos \alpha}{\sin \varphi \sin \alpha} \\
\cos \alpha &= \frac{1}{2} \frac{I_5 - I_1}{I_4 - I_2} \tag{5.41}
\end{aligned}$$

This is then substituted into (5.22) giving

$$\begin{aligned}
\tan \varphi &= 2 \sin \alpha \frac{I_2 - I_4}{2I_3 - I_5 - I_1} \\
&= 2\sqrt{1 - \cos^2 \alpha} \frac{I_2 - I_4}{2I_3 - I_5 - I_1} \\
&= 2 \sqrt{1 - \frac{(I_5 - I_1)^2}{4(I_4 - I_2)^2}} \frac{I_2 - I_4}{2I_3 - I_5 - I_1} \\
&= \sqrt{4(I_2 - I_4)^2 - \frac{(I_2 - I_4)^2(I_5 - I_1)^2}{(I_4 - I_2)^2}} \frac{1}{2I_3 - I_5 - I_1} \\
\tan \varphi &= \frac{\sqrt{4(I_2 - I_4)^2 - (I_5 - I_1)^2}}{2I_3 - I_5 - I_1} \tag{5.42}
\end{aligned}$$

This algorithm makes the same assumptions as Carré's algorithm. For sinusoidal signals there is no motion error, however it does not take into account the effect of harmonics.

5.9.2 CHARACTERISTIC POLYNOMIAL ANALYSIS

Similarly to Carré's algorithm, Novak's algorithm includes a square root in the denominator so Matlab simulations are used to predict its behaviour. Figure 5.33 shows the theoretical linearity error of Novak's algorithm. These simulations demonstrate that this algorithm is sensitive to the third harmonic in the static case as they show a four cycle error.

Figure 5.34 shows the theoretical motion error for Novak's algorithm. This algorithm is insensitive to the negative fundamental but sensitive to the third harmonic and therefore exhibits a four cycle error. There is also a phase unwrapping step where the error changes polarity due to the square root in the denominator of this algorithm.




Figure 5.33 Theoretical linearity error for Novak's algorithm




Figure 5.34 Theoretical motion error for Novak's algorithm

5.9.3 LINEARITY MEASUREMENTS

Figure 5.35 shows the linearity error for Novak's algorithm. As predicted by the Matlab simulation, this algorithm is sensitive to the third harmonic in the static case and therefore exhibits a four cycle error. The sum of the absolute value of the residuals is 1.0 ± 0.3 m. This agrees with the value for the standard algorithm within uncertainty and therefore this algorithm has not improved the linearity of the camera.




Figure 5.35 Linearity error for Novak's algorithm

5.9.4 MOTION ERROR ANALYSIS

Figure 5.36 shows the axial motion error of Novak's algorithm. The standard deviation of the error with distance for this algorithm is 0.263 ± 0.004 m and the range is 0.57 ± 0.01 m. As predicted, a phase unwrapping point creates a step in the error similar to Carré's algorithm.




Figure 5.36 Axial motion error versus distance for Novak's algorithm

As with Carré's algorithm it is interesting to analyse the algorithm with the phase unwrapping step removed. This is shown in Figure 5.37. The standard deviation of the

error with distance is now 0.030 ± 0.002 m and the range is 0.087 ± 0.008 m. While this improves on the error from the standard algorithm it is unable to improve on a number of the other algorithms tested in this chapter.




Figure 5.37 Axial motion error versus distance for Novak's algorithm with phase unwrapped

5.10 WDFT Algorithm

5.10.1 BACKGROUND

Using the rules laid out by Surrel, an algorithm can be constructed that should provide insensitivity to harmonics as desired with the minimal number of intensity measurements. Surrel states that for an arbitrary value of N , the coefficients for this algorithm can become quite complicated. However, if you allow for one additional intensity measurement (making the root at 1 a double root), then simple results can still be found. This provides the general form shown below (Surrel, 1996) which requires $2N - 1$ phase steps of $2\pi/N$ and is insensitive to harmonics up to the order of $N - 2$ even in the presence of linear motion of the object.

$$\varphi = \tan^{-1} \frac{-\sum_{i=1}^{N-1} i(I_i - I_{2N-i}) \sin\left(i \frac{2\pi}{N}\right)}{NI_N - \sum_{i=1}^{N-1} i(I_i + I_{2N-i}) \cos\left(i \frac{2\pi}{N}\right)} \quad (5.43)$$

This is equivalent to the phase being the argument of the second DFT coefficient of a set of $2N - 1$ intensity values extending over two periods and windowed by the triangle function (Surrel, 1996).

For five steps ($N = 3$) the WDFT phase algorithm is

$$\varphi = \tan^{-1} \frac{\sqrt{3}(I_1 - 2I_2 + 2I_4 - I_5)}{I_1 + 2I_2 - 6I_3 + 2I_4 + I_5} \quad (5.44)$$

and has a phase step of $2\pi/3$.

5.10.2 CHARACTERISTIC POLYNOMIAL ANALYSIS

The five frame WDFT has the characteristic polynomial of

$$\begin{aligned} P(x) &= (x - 1)^2 \left(x - \left(-\frac{1}{2} - i\frac{\sqrt{3}}{2} \right) \right)^2 \\ &= (x - 1)^2 \left(x - e^{-i\frac{2\pi}{3}} \right)^2. \end{aligned} \quad (5.45)$$

This algorithm has five phase steps of $2\pi/3$. Therefore, similarly to the N+1 algorithm, the root required to be insensitive to the third harmonic is $x = 1$ which, in contrast to the single root in the N+1 algorithm, is a double root. Therefore this algorithm should be insensitive to the third harmonic even with linear motion. There is also a double root at $m = -1$ meaning this algorithm will have no error in the presence of linear motion due to aliasing of the negative fundamental. This algorithm is insensitive to both the major sources of error discussed in section 5.2 so the motion response should have very low error. The double root at $m = -1$ is also a double root at $m = 5$, however there is no corresponding negative root. Therefore a six cycle linearity error is expected.

5.10.3 LINEARITY MEASUREMENTS

Figure 5.38 shows the linearity error for the WDFT algorithm. As predicted this algorithm is insensitive to the third harmonic. The sum of the absolute value of the residuals is 0.3 ± 0.3 m. This agrees, within uncertainty, with the previous algorithms that are also insensitive to the third harmonic. The source of the remaining error has already been discussed in section 5.4.3.

Simulations were performed to find the linearity error of this algorithm if the non-linearity of the PMD19K sensor was solved. The results are shown in Figure 5.39. As predicted, a 6 cycle error is observed.




Figure 5.38 Linearity error for the five frame WDFT algorithm




Figure 5.39 Theoretical Linearity Error versus phase for the WDFT algorithm

5.10.4 MOTION ERROR ANALYSIS

Figure 5.40 shows the motion error for the WDFT algorithm. The standard deviation of the error with distance for this algorithm is 0.009 ± 0.004 m and the range is 0.03 ± 0.01 m. This algorithm is the best of those tested so far with the lowest standard deviation and range. This was expected as this algorithm is insensitive to the first order of motion error from both the negative fundamental and the third harmonic. The change in error of the algorithm with distance is within the uncertainty of the

measurements so it is not possible to identify if a higher harmonic, or a higher order of the negative fundamental or the third harmonic, is now the primary source of error.

To be able to identify the primary source of error in the WDFT algorithm more accurate data is required. As was used previously for the standard algorithm, phase stepping can be used to get higher resolution data. Instead of setting the phase step to $2\pi/3$ it is set to $22\pi/32$, this is the closest equivalent to the velocity used previously that can be chosen with the 320 phase steps available with the Cyclone III FPGA used in the Victoria University Range Imaging System. The result of this linear miscalibration of the phase stepping is shown in Figure 5.41. Even with this data it is still unclear what systematic error remains.




Figure 5.40 Axial motion error versus distance for the four frame WDFT algorithm

5.11 N + 3 Algorithm

5.11.1 BACKGROUND

In section 5.2 we showed the two primary sources of error were the first order errors of the negative fundamental and the third harmonic. Very encouraging results were found using the WDFT, which is insensitive to both of these error sources. It is still interesting to investigate the effect of an algorithm that is insensitive to the second order of motion error. Hibino et al. (Hibino et al., 1997) used a simultaneous equations approach to make an algorithm insensitive to a phase shift miscalibration up to the order p ($p \leq M-1$) in the




Figure 5.41 Linear miscalibration error for the five frame WDFT algorithm

presence of up to the j^{th} harmonic. In other words, allowing the phase shift for the r^{th} frame to have the form

$$\alpha_r = \alpha_{0r} \left[1 + \varepsilon_1 + \varepsilon_2 \frac{\alpha_{0r}}{\pi} + \varepsilon_3 \left(\frac{\alpha_{0r}}{\pi} \right)^2 + \cdots + \varepsilon_p \left(\frac{\alpha_{0r}}{\pi} \right)^{p-1} \right] \quad (5.46)$$

for $r = 1, 2, \dots, M$ where ε_q ($1 \leq q \leq p$) are the error coefficients. Previous algorithms have only dealt with the first order ($p = 1$), for this algorithm we are interested in the solution for $p = 2$.

To be insensitive to the j^{th} harmonic $j+4$ samples are required. Hibino showed that the first and second order terms can be removed by combinations of the first, second, second to last and last intensity values (compare this to using the first and last intensity values to remove the first order error in the $N+1$ algorithm). Solving simultaneous equations, the general form of the algorithm is shown to be

$$\varphi = \tan^{-1} \frac{\frac{1}{4}(I_1 + I_2 - I_{j+3} - I_{j+4}) \frac{\sin\left(\frac{3\pi}{j+2}\right)}{\sin^2\left(\frac{2\pi}{j+2}\right)} + \sum_{i=2}^{j+3} I_i \sin\left(\frac{2\pi}{j+2}\right) \left(i - \frac{j+5}{2}\right)}{\frac{1}{4}(I_1 - I_2 - I_{j+3} + I_{j+4}) \frac{\cos\left(\frac{3\pi}{j+2}\right)}{\sin^2\left(\frac{2\pi}{j+2}\right)} + \sum_{i=2}^{j+3} I_i \cos\left(\frac{2\pi}{j+2}\right) \left(i - \frac{j+5}{2}\right)} \quad (5.47)$$

using five steps ($j = 1$) the solution is therefore

$$\varphi = \tan^{-1} \frac{3\sqrt{3}(I_4 - I_2)}{2I_1 + I_2 - 6I_3 + I_4 + 2I_5} \quad (5.48)$$

with five phase steps of $2\pi/3$.

5.11.2 CHARACTERISTIC POLYNOMIAL ANALYSIS

The characteristic polynomial for this algorithm is

$$\begin{aligned} P(x) &= (x - 1) \left(x + \left(\frac{1}{2} + i \frac{\sqrt{3}}{2} \right) \right)^3 \\ &= (x - 1) \left(x - e^{i\frac{2\pi}{3}} \right)^3 \end{aligned} \quad (5.49)$$

this algorithm has a single root at $x = 1$. Therefore with a phase step of $2\pi/3$, it is insensitive to the third harmonic in the static case but not with linear motion. The cube root at $m = -1$ means that there will be no error terms up to ϵ^2 due to the negative fundamental. The primary source of error will be due to sensitivity to the third harmonic therefore a two cycle error is expected. The additional root at $m = -1$ will not be beneficial over having first order insensitivity, as the error will be dominated by error due to the third harmonic. The triple root at $m = -1$ is also a triple root for $m = 5$, however, since the negative root is not present, a 6 cycle linearity error is expected.

5.11.3 LINEARITY MEASUREMENTS

Figure 5.42 shows the linearity error of the five frame N+3 algorithm. This algorithm is insensitive to the third harmonic, which was predicted for the static case. The sum of the absolute value of the residuals is 0.4 ± 0.3 m. This agrees with other algorithms insensitive to the third harmonic within uncertainty. The remaining error has been explained in section 5.4.3.

Simulations were performed to find the linearity error of this algorithm if the non-harmonic error was not present. The results are shown in Figure 5.43. As predicted, a 6 cycle error is observed.




Figure 5.42 Linearity error for the five frame N+3 algorithm




Figure 5.43 Theoretical Linearity Error versus phase for the five frame N+3 algorithm

5.11.4 MOTION ERROR ANALYSIS

Figure 5.44 shows the motion error for the five frame N+3 algorithm. The standard deviation of the error with distance for this algorithm is 0.026 ± 0.002 m and the range is 0.067 ± 0.007 m. While this is an improvement over the standard algorithm, and is marginally better than Carrè's algorithm, Novak's algorithm and Hariharan's algorithm, it falls short of the results shown by the WDFT. This is due to sensitivity to the third harmonic in the dynamic case. As predicted, the first term of the third harmonic is more significant than the second term of the negative fundamental and therefore while this

algorithm shows improved motion error over the standard algorithm it has not improved on the WDFT algorithm.




Figure 5.44 Axial motion error versus distance for the five frame N+3 algorithm

5.12 Comparison of algorithms

A summary of the algorithms presented in this chapter is shown in Table 5.3.

Table 5.2 shows a performance comparison of the different phase algorithms tested in this chapter. The linearity measurements clearly separate the algorithms into two classes, algorithms that are insensitive to the third harmonic in the static case, which have a sum linearity error of 0.3 – 0.5 m, and algorithms that do not have this insensitivity, which have a sum linearity error of 1.0 – 1.2 m. Below the third harmonic there is another noise source, not due to harmonics, which becomes dominant. It has been proposed in section 5.4.3 that this error is due to non-linearity in the PMD19K sensor's phase response to changes in intensity, although there are other possibilities such as multi-path interference. Simulations show that, when this error source is not present, all of the algorithms experience a six cycle error due to the negative 5th and the positive 7th harmonics, with the exception of the 5 frame N-bucket algorithm which experiences a ten cycle error. In terms of linearity, the five frame N-bucket, four frame N+1 Type B, five frame WDFT and five frame N+3 algorithms provide significantly better performance than the other algorithms, with the five frame N-bucket algorithm being

theoretically slightly better than the others. However, due to the harmonic content of our system, the practical benefit of the five frame algorithm over the others is questionable.

A comparison of the motion error of the algorithms, shown in Figure 5.45, demonstrates that there are numerous algorithms that can be used to improve on the standard four frame algorithm. The WDFT algorithm showed the best motion error response with both the lowest standard deviation and the lowest range by a significant margin. The second best algorithm was the five frame N+3 algorithm.

The WDFT has shown the best response as it removes the first order of motion error due to both the negative fundamental and the third harmonic. It is also favourable as it is a member of the class of algorithms that are insensitive to the third harmonic in the static case and therefore it has improved linearity. Unlike Carré's algorithm and Novak's algorithm, the WDFT does not have any additional phase unwrapping points that require extra computation to handle. It is therefore selected as the best algorithm for our application.

Table 5.2 Comparison of Motion Error and Linearity of phase detection algorithms

Algorithm	Standard Deviation (m)	Range (m)	Sum of linearity residuals (m)
Standard four frame	0.044 ± 0.002	0.112 ± 0.007	1.0 ± 0.3
Standard five frame	0.042 ± 0.002	0.110 ± 0.005	0.5 ± 0.3
Order changed four frame	0.042 ± 0.002	0.108 ± 0.005	1.2 ± 0.3
Carré's algorithm (unwrapped)	0.029 ± 0.002	0.078 ± 0.007	1.1 ± 0.3
Hariharan algorithm	0.028 ± 0.002	0.083 ± 0.008	1.0 ± 0.3
N+1 type B algorithm	0.041 ± 0.004	0.11 ± 0.01	0.3 ± 0.3
Novak's algorithm (unwrapped)	0.030 ± 0.002	0.087 ± 0.008	1.0 ± 0.3
WDFT algorithm	0.009 ± 0.004	0.03 ± 0.01	0.3 ± 0.3
N+3 algorithm	0.026 ± 0.002	0.067 ± 0.007	0.4 ± 0.3

Table 5.3 Summary of phase detection algorithms

Algorithm	Phase Steps	Phase Equation $\tan \varphi =$	Characteristic Polynomial $P(x) =$
N Bucket (4 step)	$\pi n/2$ $n = 0, 1, 2, 3$	$\frac{I_2 - I_4}{I_3 - I_1}$	$-i(x-1)(x+1)(x+i)$
N Bucket (5 step)	$2\pi n/5$ $n = 0, 1, 2, 3, 4$	$\frac{\sum_{i=1}^5 I_i \sin\left((i-1)\frac{2\pi}{5}\right)}{\sum_{i=1}^5 I_i \cos\left((i-1)\frac{2\pi}{5}\right)}$	$(x-1)\left(x - e^{\frac{6\pi i}{5}}\right)\left(x - e^{-\frac{6\pi i}{5}}\right)\left(x - e^{\frac{2\pi i}{5}}\right)$
Carré's algorithm	$(2n-3)\alpha$, $n = 0, 1, 2, 3$ $\alpha = \text{arbitrary}$	$\frac{\sqrt{3(I_2 - I_3) - (I_1 - I_4)[(I_1 - I_4) + (I_2 - I_3)]}}{(I_2 + I_3) - (I_1 + I_4)}$	-
N+1	$2\pi n/3$ $n = 0, 1, 2, 3$	$\frac{I_1 - 3I_2 + 3I_3 - I_4}{\sqrt{3}(I_1 - I_2 - I_3 + I_4)}$	$-\sqrt{3}i(x-1)\left(x + \left(\frac{1}{2} + \frac{\sqrt{3}}{2}i\right)\right)^2$
Harikaran's	$(n-2)\pi/2$ $n = 0, 1, 2, 3, 4$	$\frac{2(I_2 - I_4)}{2I_3 - I_5 - I_1}$	$-(x-1)(x+1)(x+i)^2$
Novak's	$(n-2)\pi/2$ $n = 0, 1, 2, 3, 4$	$\frac{\sqrt{4(I_2 - I_4)^2 - (I_5 - I_1)^2}}{2I_3 - I_5 - I_1}$	-
WDFT	$2\pi(n+1)/3$ $n = 0, 1, 2, 3, 4$	$\frac{\sqrt{3}(I_1 - 2I_2 + 2I_4 - I_5)}{I_1 + 2I_2 - 6I_3 + 2I_4 + I_5}$	$(x-1)^2\left(x + \left(\frac{1}{2} + i\frac{\sqrt{3}}{2}\right)\right)^2$
N+3	$2\pi(n-2)/3$ $n = 0, 1, 2, 3, 4$	$\frac{3\sqrt{3}(I_4 - I_2)}{2I_1 + I_2 - 6I_3 + I_4 + 2I_5}$	$(x-1)\left(x + \left(\frac{1}{2} + i\frac{\sqrt{3}}{2}\right)\right)^3$




Figure 5.45 Comparison of axial motion error between all algorithms

5.13 Multiple Velocity Measurements

The measurements in this chapter were repeated to ensure the algorithm is robust over a range of positive and negative velocities. Velocities of 0.5 m/s, 1 m/s, 1.5 m/s and 2 m/s were used. These measurements were performed for the four frame standard algorithm and the five frame WDFT algorithm. Due to the fixed length of the table, the distance over which data could be acquired for slower speeds is larger, as the distance required for acceleration and deceleration is reduced.

Figure 5.46 shows the motion error for multiple velocities for the four frame standard algorithm and Figure 5.47 shows the same for the five frame WDFT. For the standard algorithm, as expected, as the velocity increases the amplitude of the approximately sinusoidal error increases. In comparison, for the five frame WDFT, there is no discernable difference in error between the different velocities with all of them being within the precision of our camera. The WDFT has therefore been demonstrated to provide phase measurements that are not measurably influenced by a wide range of axial velocities.




Figure 5.46 Comparison of axial motion error versus distance for the four frame standard algorithm for multiple velocities




Figure 5.47 Comparison of axial motion error versus distance for the five frame WDFT algorithm for multiple velocities

5.14 Change in amplitude with distance

Previously we have assumed that the amplitude remains constant over the measurement. However, in reality there will be an inverse square decrease in amplitude with distance. While experimental results have been very promising, due to the acceleration time of the linear table and saturation of the camera, close range data was not measured and short distances is where the change in amplitude is expected to be most significant.

Data were simulated in Matlab and a graph comparing the standard algorithm to the WDFT both with and without the inverse square decrease is shown in Figure 5.48. This




Figure 5.48 Effect of inverse square decrease on the motion error for both the four frame standard and five frame WDFT algorithms

demonstrates that while the change in amplitude has an effect on the motion error, the impact is small for distances 1.5 m and further when using the standard algorithm. For the WDFT algorithm there is additional error for distances closer than 1 m. With the inclusion of the inverse square decrease in amplitude with distance the WDFT algorithm still provides a large improvement over the standard algorithm, even at close distances.

5.15 FPGA Implementation

As the algorithm that showed the best overall performance, the five frame WDFT algorithm was implemented on the Victoria University Range Imaging System's FPGA. Because the intensity values and their coefficients are stored as 16-bit signed integers on the FPGA the equation actually implemented was

$$\varphi = \tan^{-1} \frac{5461 \times \sqrt{3}(I_1 - 2I_2 + 2I_4 - I_5)}{5461 \times (I_1 + 2I_2 - 6I_3 + 2I_4 + I_5)} \quad (5.50)$$

to use the maximum range of the 16 bit coefficient.

The data acquisition was rerun to acquire data for the FPGA implementation. A comparison between the motion error of the FPGA and Matlab implementations is shown in Figure 5.49. The scale has been maintained from the previous measurements to allow easy comparison.




Figure 5.49 Comparison of axial motion error for the WDFT between Matlab and FPGA implementations

The FPGA implementation has successfully reproduced the results of the WDFT algorithm processed using Matlab. Confidence tests can be used to establish if the small differences observed are significant. To be able to perform confidence tests on the metrics used in this chapter we must first establish the form of their distribution. Figure 5.50 shows the distribution of standard deviations for the WDFT with processing done in Matlab over the 100 data runs. Similarly Figure 5.51 shows the distribution for the WDFT processed on the FPGA. To test if the standard deviations are normally distributed a Kolmogorov-Smirnov test was performed. For the Matlab processing this gave a P value of 0.4895 and for FPGA processing this gave a P value of 0.5738. The null hypothesis that the data comes from the normal distribution is not rejected. Parametric tests can therefore be used to compare these data.

Using an unpaired t test to compare the standard deviations gives a P value of 0.3725. The null hypothesis that the mean values of these metrics are the same cannot be rejected. We conclude that there is no statistically significant difference between these two implementations of the WDFT algorithm.

5.16 Comparison of precision

To ensure the quality of the phase measurements has not been compromised, a comparison between the precision of the standard algorithm and the WDFT was performed. So as to not bias the results due to the readout time of the sensor, the five frame standard algorithm was used.

Figure 5.52 shows a comparison of the standard deviation for the two algorithms. At long distances the precision of the WDFT appears slightly worse overall than the five frame standard algorithm. Since precision is related to the intensity of the received light, and therefore the distance, a paired test is used to determine if the difference between these two algorithms is significant. The Spearman correlation coefficient was calculated to confirm the pairing which gave a p value of 0.9324, this gives a probability of < 0.01% that a correlation coefficient this large or larger would be observed if the data were not correlated. To test if a t test is valid, a Kolmogorov-Smirnov test was performed, this gave a P value of 0.0124, meaning it is very likely that the data are not normally distributed. A Wilcoxon matched pairs test was therefore used to compare the data




Figure 5.50 Distribution of standard deviations for the WDFT with processing done in Matlab




Figure 5.51 Distribution of standard deviations for the WDFT with processing done on the FPGA

instead of a paired t test. This gives a P value of 0.0001 and therefore the two have statistically significantly different medians. The median for the WDFT is 0.013 ± 0.009 m and the median for the standard algorithm is 0.012 ± 0.008 m.

While there is a statistically significant difference between the two, compared to the numerous other factors that impact the precision of indirect time of flight cameras, specifically the large influence of frame time and the reflectivity of the imaged object, this small change in precision does not impact the usefulness of the WDFT algorithm. Taking the most easily controllable of these factors, the frame time, the median precision for the standard algorithm with an frame time of 30 ms, instead of the 25 ms




Figure 5.52 Comparison of the standard deviation of range measurements versus distance for both the four frame standard algorithm and the five frame WDFT algorithm

used previously, is 0.008 ± 0.006 m. This small change, where legitimate frame times for the Victoria University Range Imaging System range from 10 ms to 500 ms or larger, has a much greater effect than changing between the two algorithms does.

5.17 Application specific algorithms

While an algorithm has been found that effectively eliminates axial motion error for our camera, the flexibility of phase detection algorithms means that there is potential for other algorithms to provide benefits in situation specific applications. This section will outline two potential areas where algorithms other than the five frame WDFT may be beneficial.

One application where specialised algorithms could be applied is where it is expected that non-linear motion will occur, in particular where an object is accelerating at a constant rate. From Surrel (Surrel, 1996), we know that to be insensitive to acceleration error, higher multiplicity of roots is required. Starting from the WDFT, it is reasonably easy to construct algorithms that have improved insensitivity to the error from accelerating objects. The multiplicity of the double roots at $m = -1, 3$ and -3 is increased, with each increase requiring two additional phase steps. The results, measured by manipulating the phase step size to imitate acceleration, are shown in Figure 5.53. The WDFT (double root) is plotted for comparison. This demonstrates that increasing the multiplicity of roots can improve the error due to acceleration significantly if required.




Figure 5.53 Comparison of acceleration error between algorithms with different multiplicity of roots at $m = -1, 3$ and -3 for an acceleration of $\pi/80$ per frame

Another application where custom algorithms could be used is where linear motion is not a concern. In this case the linearity of the camera can be improved while still only using four frames by selecting the location of roots of the phase algorithm to mitigate linearity error instead.

From our results this can be achieved by implementing the four frame type B $N+1$ algorithm discussed in Section 5.8. While the linearity of this algorithm is not as good as the five frame standard algorithm, it is significantly better than the four frame standard algorithm as the first harmonic it is sensitive to is the negative 5th. For our system, since




Figure 5.54 Comparison of linearity error using the phase stepping method

the higher harmonics are so weak it is likely that the measurable difference between the two will be extremely small. Figure 5.54 shows a comparison of the measured linearity error for the two algorithms using the phase stepping method. This demonstrates that the error is similar, although the four frame N+1 algorithm is slightly worse.

By observing the characteristic polynomial of the four frame N+1 algorithm we can see that there is a double root at $m = -1$ that is not necessary for linearity, this root can be removed and the same linearity maintained, in this case the algorithm reduces to the three frame standard algorithm.

5.18 Summary

This chapter has investigated the implementation of different phase detection algorithms in indirect time of flight cameras. Algorithms were tested both for static linearity and motion error. Several algorithms were found that can improve on the performance of the standard algorithm. The best performing algorithm for motion error was the Windowed Discrete Fourier Transform, which improved the standard deviation of the motion error from 0.044 ± 0.002 m to 0.009 ± 0.004 m and the range of the motion error from 0.112 ± 0.007 m to 0.03 ± 0.01 m. This algorithm was also beneficial for the linearity, improving the sum of the residuals of the linearity from 1.0 ± 0.3 m to 0.3 ± 0.3 m over the standard four frame algorithm. The linearity would be further improved if the phase non-linearity of the PMD19K sensor was solved. This algorithm was implemented and run in real time on the FPGA in the Victoria University Range Imaging System. This showed equivalent performance demonstrating this algorithm can be implemented in real time with a trivial increase in computational effort over the standard algorithm. It was also shown that the WDFT was effective at eliminating motion error across a range of velocities. The precision of the five frame WDFT was compared to the standard algorithm and it was found that the new algorithm had decreased the precision of the range calculations slightly. However, the change is minimal compared to other factors that influence the precision. The idea of custom algorithms for application specific problems was also investigated, and it was shown that phase detection algorithms provide flexibility to be adapted to a number of applications.

Chapter 6 REAL SYSTEM IMPLEMENTATION

6.1 Introduction

In this chapter the compact version of the Victoria University Range Imaging System, developed as part of this research and described in Chapter 3, is combined with the algorithm and operational improvements outlined in Chapters 4 and 5 into a real world system. A number of real world environments are imaged with the system and applications are demonstrated, specifically 3D modelling and implementation of the system on a mobile robot.

6.2 Combination of improvements

Four operating modes are implemented in the VHDL code for the Cyclone III FPGA on the Victoria University Range Imaging System, which can be selected in real time. All calculations are performed on the FPGA with no additional computational power or memory elements required. The four modes implemented on the Victoria University Range Imaging System are:

- the standard operating mode with no improvements
- the standard phase detection algorithm with Variable Frame Rate imaging (precision thresholding was used for the results in this chapter) (VFRI)
- the WDFT phase detection algorithm with standard dynamic range (WDFT)
- the WDFT phase detection algorithm with Variable Frame Rate imaging (WDFT + VFRI).

To demonstrate the improvements possible using these different modes, linearity measurements were performed for all four operating modes and are shown in Figure 6.1. Table 6.1 shows calculated parameters for the different operating modes.




Figure 6.1 Linearity comparison of different operating modes implemented on the Victoria University Range Imaging System

As expected, the modes with Variable Frame Rate imaging enabled have much higher precision at long range than modes that do not. The precision is determined by the threshold value used. In this case the threshold has simply been chosen to effectively demonstrate the difference in operating modes. The system implementing VFRI has improved the median precision by a factor of three, and the minimum precision by an order of magnitude over the standard system.

The WDFT + VFRI system has slightly better median precision than the VFRI system while the minimum precision is significantly worse. The minimum precision of the WDFT + VFRI system is heavily influenced by a single outlier that, if removed, reduces the minimum precision to 0.036 ± 0.004 m, which is similar although slightly worse than the system using VFRI only. With the exclusion of the outlier, the WDFT + VFRI system improves the precision by an order of magnitude over the system using only the WDFT.

Table 6.1 Precision and linearity of different operating modes of the Victoria University Range Imaging System

Operating mode	Median Precision (m)	Minimum Precision (m)	Sum linearity error (m)
Standard	0.05 ± 0.01	0.21 ± 0.02	1.2 ± 0.3
WDFT	0.09 ± 0.02	0.36 ± 0.06	0.3 ± 0.3
VFRI	0.015 ± 0.003	0.0227 ± 0.0005	1.24 ± 0.06
WDFT + VFRI	0.0114 ± 0.0006	0.05 ± 0.02	0.34 ± 0.06

Operating modes using the WDFT algorithm have better linearity than the modes using the standard algorithm, where sinusoidal errors are visible due to the third harmonic. From the data in Table 6.1, the linearity error of the WDFT system is a quarter of the linearity error observed for the standard system. As discussed in Chapter 5, this error could be further reduced by removing amplitude related phase error. Within uncertainty, the WDFT + VFRI system has the same linearity error that was observed for the system only implementing the WDFT technique.

The linearity and precision improvements provided by the VFRI and WDFT techniques agree with the experiments presented previously in Chapter 4 and Chapter 5. Here the two techniques are combined and the measurements demonstrate that the WDFT + VFRI system has successfully reproduced the positive aspects of both techniques. Using this system high precision, high accuracy data can be acquired without requiring any calibration of the camera.

As discussed in Section 4.6, the VFRI technique is known to decrease the frame rate of the measurements. Figure 6.2 shows the impact of the technique on the frame rate of the operating modes that implement this technique. For distances up to 1.6 m for the VFRI system and 1.5 m for the WDFT + VFRI system, the frame rate is the maximum rate of 8 FPS as integrating for a single frame is sufficient to acquire the desired quality. For longer distances the frame rate decreases in order to provide the constant precision. For the VFRI system the worst frame rate (at a distance of 4.9 m) is 0.144 FPS, for the WDFT + VFRI system it is 0.032 FPS. The lower frame rates for the WDFT + VFRI are expected as the WDFT algorithm was shown to have slightly worse precision than the standard algorithm.




Figure 6.2 Frame rate versus distance for both the VFRI and the WDFT + VFRI operating modes

There is some quantisation appearing in the frame rates, as the possible precision values are quantised by the requirement to integrate over an integer number of frames. This results in some flat areas in the recorded frame rate over a range of distances, for example for the VFRI system between 1.8 m and 2.1 m. For this range of distances the precision of the pixel is too low using a single measurement but exceeds the threshold using two measurements, therefore the same frame rate of 4 FPS is observed for all of these distances.

6.3 Comparison of range images of real scenes

Two static real world scenes were imaged using the Victoria University Range Imaging System to demonstrate the improvements due to the new operating modes. To simplify the comparison, two modes have been used, the standard operating mode and the WDFT + VFRI mode. A photograph of the scene is also included to provide context to the image. The frame time for both systems was selected as the longest frame time for which no pixels in the scene saturate.

The first real world scene (referred to as Scene 1) was a view from the floor of the Victoria University of Wellington Range Imaging lab. This is a view that a mobile robot exploring the lab might encounter. A photograph of this scene is shown in Figure 6.3. A red box has been included to indicate the approximate field of view of the Victoria University Range Imaging System when recording these measurements. This field of

view is set by the lens used in the system and could be expanded (or contracted) if required. A number of features are noticeable in this scene. There is the edge of a chair overhanging the range imaging system's view, although the base and column of the chair is not visible to the camera. A discarded object is visible in the bottom right area of the scene. There are also some boxes in the background. This scene was chosen as it contains both objects that are at a number of different distances and objects made of materials with significantly different optical properties. The curved edges of the object in the bottom right of the scene will test the spatial resolution of the system. This scene is typical of what may be observed in the real world and provides a test of the dynamic range of the system. The measurement time for this scene was 600 ms.



Figure 6.3 Photo of Scene 1

Figure 6.4 shows a typical range image of Scene 1 using both the standard mode and the WDFT + VFRI mode. The two images are similar however there are some differences. Due to the limited dynamic range of the system the precision of the measurements was low in some areas of the scene when using the standard mode. The degradation is particularly noticeable on the black part of the chair (the blue coloured area in the top left of Figure 6.4), where low intensity returns result in the standard mode image appearing speckled due to the low precision of the measurements in this area. In the WDFT + VFRI image this speckled appearance is not present.




Figure 6.4 Typical capture of Scene 1 using the standard system (left) and the WDFT + VFRI system (right)

To demonstrate the difference in precision between the two systems for Scene 1, the standard deviation for each pixel was calculated over 200 measurements and is shown for the standard system in Figure 6.5 and the WDFT + VFRI system in Figure 6.6. For areas with high intensity both images are the same. For the chair seat, which was identified as an area that had low intensity, the standard system has an average precision of 0.05 ± 0.01 m and the WDFT + VFRI system has an average precision of 0.010 ± 0.002 m. Even higher precision could be achieved using the WDFT + VFRI system by changing the threshold, although this would come with the trade-off of decreased frame rates.

The second example scene is a view from a desktop. Notable features in this scene are a power supply unit, which combines a highly reflective front bezel with a low reflectivity front display screen. Behind the power supply is a window with closed blinds. The sensor used in the Victoria University Range Imaging System has poor background light suppression technology compared to more modern sensors, and therefore light bleeding around the gaps in the blinds is likely to cause problems for our sensor.

A photograph of this scene is shown in Figure 6.7. Again a red box has been included to indicate the approximate field of view of the Victoria University Range Imaging System when taking these measurements. This scene will test the ability of the system to reject background illumination and also highlight the effect of objects transparent to the wavelength of light used in the system. The measurement time for this scene was 600 ms.




Figure 6.5 Standard deviation of range measurements for Scene 1 using the standard system




Figure 6.6 Standard deviation of range measurements for Scene 1 using the WDFT + VFRI system

Figure 6.8 shows the measured range for both systems. Again both images look similar, however for the standard system, the blinds in the background and the power supply display a degradation in the quality of the measurements. As expected, where sunlight is leaking through the blinds, the distance measurement cannot be performed. As mentioned previously, more modern sensors have improved background light suppression to mitigate this issue. Some manufacturers provide specific background light suppression models of their cameras, where additional circuitry is included on the sensor IC to improve performance in high lighting conditions.




Figure 6.7 Photo of Scene 2

The precision for Scene 2, calculated over 1000 measurements, is shown for the standard system in Figure 6.9 and the WDFT + VFRI system in Figure 6.10. Where the sunlight has saturated the sensor the standard deviation is very high, as it is essentially a random measurement. On the dark front panel of the power supply there is significant improvement in the precision from 0.10 ± 0.04 m to 0.019 ± 0.003 m due to the VFRI technique as the low returning light level causes poor precision for the standard system.




Figure 6.8 Typical capture of Scene 2 using the standard system (left) and the WDFT + VFRI system (right)




Figure 6.9 Standard deviation of range measurements for Scene 2 using the standard system




Figure 6.10 Standard deviation of range measurements for Scene 2 using system with WDFT + VFRI

To show the effect of the sunlight, the measurements were repeated at night and are shown in Figure 6.11. Measurements were only taken with the WDFT + VFRI system. The measurements are the same as those in Figure 6.8 with the exception of the area that was previously spoiled by sunlight, which is now properly imaged.

A final characteristic to note in Scene 2 is the distortion in the top of the power supply above the display screen. Close observation shows there is a slight dip here and not a straight edge. A high angle photograph of the scene, shown in Figure 6.12, shows that the glass of the display, which is transparent to red light, extends across the top of the

power supply unit. Therefore in this region the blinds behind the supply are being imaged (with some multipath interference occurring due to some amount of reflected light from the screen).




Figure 6.11 Typical measurement of Scene 2 using the WDFT + VFRI system at night




Figure 6.12 High angle view of Scene 2

Scene 2 has highlighted an issue for these cameras, and distance measurement systems in general, in real world environments, namely the detection of transparent objects. Glass doors, walls and windows are significant hazards for mobile robots and are difficult to detect using many of the standard approaches for robotic sensor systems, they are occasionally difficult even for human beings. More research is required to find solutions for this problem, but perhaps augmenting the high resolution electromagnetic wave

based range imaging system with an ultrasonic sensor would be the simplest solution for mobile robotics.

6.4 3D Modelling

One promising application for indirect time of flight cameras is the modelling of objects. Three dimensional modelling is used for a wide variety of applications, such as profiling objects of cultural significance, modelling for films and reproducing objects with 3D printers. The object profiled in this section is a foam model hand, which is shown in Figure 6.13. This object has a number of differently oriented curved surfaces including surfaces both parallel and perpendicular to the viewing axis of the camera meaning a distance profile of this object should be a good test of the depth and spatial resolution of the camera.



Figure 6.13 Photograph of foam hand

The profile of the hand was measured by recording 500 frames of the hand using a measurement time of 500 ms and averaging them. The hand was then removed and a further 500 frames were measured and averaged to provide a calibration measurement. The subtraction of the two measurements provides the profile of the hand. The profile of the hand, using the WDFT + VFRI system, is shown in Figure 6.14.

When doing 3D modelling of static objects the scene will generally be averaged over a number of measurements to provide a high precision measurement. As VFRI is essentially a targeted averaging technique the advantage of VFRI is not immediately

obvious in this application. However, commonly when modelling objects there will be a minimal level of precision that is desired for the model. While averaging over the entire scene for a very large number of measurements will provide this precision, the number of measurements required is difficult to ascertain, as it is highly scene dependent. By using VFRI, the level of precision can be selected using the threshold, instead of needing an arbitrarily large number of frames to ensure the desired precision is achieved. The VFRI mode can be configured to output a signal when all the pixels in the scene have achieved the desired precision and calculated a phase value, therefore measuring to the required precision in minimal time.




Figure 6.14 Profile of model hand measured using the WDFT +VFRI system

A higher depth resolution can be achieved by increasing the modulation frequency. Frequencies above 40 MHz are not possible using our hardware, however future models of this system could implement higher frequencies by replacing the modulation drivers and sensor. Higher modulation frequencies also have a shorter unambiguous measurement distance and therefore the linearity improvements shown by the WDFT will be more important.

6.5 Implementation on a mobile platform

The previous real world scenes have been static, however this thesis is primarily interested in mobile robotics applications. In this section the system was mounted on a

Pioneer 2 mobile robot and the suitability of the system was evaluated. A photograph of this configuration is shown in Figure 6.15.




Figure 6.15 Photograph of Pioneer 2 robot setup front view (left) and side view (right)

6.5.1 DRIVING TOWARDS A WALL

The simplest test to perform using the robot is to drive in a straight line. A photograph of the initial scene is shown in Figure 6.16. Some inconsequential parts of the scene have been obscured to protect confidential information.




Figure 6.16 Photograph of the scene used for the straight line driving test

After the initial setup, some non-idealities of the robot were found while running this test. When instructed to move in a straight line at a constant speed the robot was unable to maintain that constant speed at the beginning of the motion profile. The motors would cycle between moving too fast and too slow, resulting in a jerky motion. This occurs five times, with a period of 2 s and a distance of 0.4 m, until the robot settles to a constant velocity, taking half of the available 4 m driving distance. This caused the recorded measurements to not be satisfactorily repeatable and therefore direct statistical comparisons between the two modes is problematic. The pioneer robot used in these experiments is an old model, and perhaps a newer model would provide more repeatable measurements. While they make repeatable measurements difficult, non-idealities such as these are present in mobile robotics and are issues that these cameras will have to deal with.

Due to the long distances and dark objects in the scene, a long integration time of 250 ms and low velocity of 0.2 m/s were used for these measurements to provide reasonable precision measurements of the scene. In terms of motion error this is equivalent to moving at a higher speed with a lower integration time.

Because the entire scene will be changing rapidly, the VFRI technique was not used for these measurements. When implementing this technique the non-VFRI image can be outputted concurrently, provided an increase in memory requirements is acceptable. It would therefore be possible to compare the standard and VFRI pixels in real time to identify moving pixels and choose the standard pixel in these cases, however this is outside the scope of this thesis.

There was also considerable lateral movement of the robot as it followed the nominally straight path. Figure 6.17 shows four consecutive frames recorded during the data run. Objects such as the chair back in the bottom left hand side of the scene should move outwards as the robot moves forwards with a constant angular field of view, however observation of frame 3 in particular shows that at times these objects have moved back towards the centre of the frame as the robot has weaved slightly to the left during the acquisition. This movement has caused some significant phase detection issues at the edges of objects as no method for dealing with lateral motion has been included in this

system. Methods of dealing with this error were discussed in section 2.3.4.6 The distance images are heavily influenced by this lateral motion and the random nature of it further impedes the recording of repeatable results for direct numerical comparison between the two systems.

The robot was set up and driven along nominally the same trajectory for both systems. A

Consecutive range images demonstrating weaving motion




Figure 6.17 Four consecutive range images indicating lateral movement from first (top) to last (bottom)

comparison frame is shown in Figure 6.18 for the starting point. As the frame time had to be set such that the camera would not saturate at the end of the motion profile (when it is close to the far wall) the precision for dark objects such as the chair backs and monitors is poor in this image. Also, the distance to some objects is greater than 5 m and therefore the measured distance has wrapped and the far wall, for example, is measured as being at a very short distance. This could be corrected by using multiple modulation frequencies to extend the unambiguous distance as discussed in section 2.2.3.2.

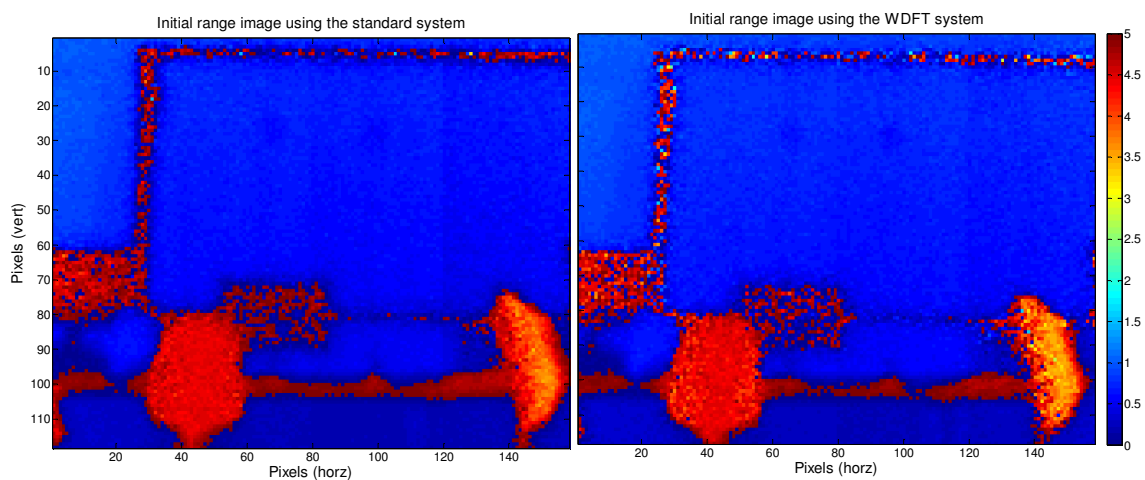


Figure 6.18 Initial range measurement for straight line driving test measurements using the standard system (left) and WDFT system (right)

Figure 6.19 demonstrates a comparison of four consecutive images taken in two separate data runs, one using the standard system and one using the system implementing the WDFT algorithm. There are a couple of notable features in these images. Both the systems suffer from some issues at the edges of certain objects both due to the camera moving forward, and therefore these edges shifting, and due to lateral motion of the robot as described earlier in this section.

Observation of the measured distance of the chair on the left hand side of the scene, from the perspective of the camera, appears to show the WDFT algorithm is providing a superior distance measurement. To investigate this, the average distance measured to the chair (the left hand side where some edge issues occur in some frames has been excluded from these calculations) is shown in Figure 6.20. These measurements show that on average the WDFT provided a smoother trajectory for the chair than the

standard algorithm. Using a linear fit of the data, the relative velocity of the chair was calculated as -0.18 ± 0.04 m/s using the standard system and -0.167 ± 0.009 m/s using the WDFT + VFRI system. As the change in measured distance is much smoother for the WDFT algorithm, the linear fit is much better and the slope measurement has higher precision.




Figure 6.19 Comparison of range images for the straight line test using the standard algorithm (left) and the WDFT (right)




Figure 6.20 Comparison of the average measured distance of the seatback over four measurements between the standard operating mode and the WDFT mode

Because the chair is not directly in front of the camera, a correction has to be made to get a value to compare with the set velocity of the robot. A model of the system is shown in Figure 6.21. The robot starts at frame 1 at the position indicated and is moving along D_{perp} , the initial angle to the chair is θ . Using the initial D_{meas} and estimating the initial angle using the field of view of the camera, the distance D_{para} can be calculated, and is a constant over the measurements. Using the calculated D_{para} and the measured distance, D_{perp} can be calculated for the four frames and the change in these measurements provides the estimated velocity of the robot.

Performing this correction provides velocities of -0.21 ± 0.05 m/s and -0.20 ± 0.01 m/s. Both of these values agree with the set velocity of 0.2 m/s within uncertainty, however the WDFT system was able to measure the velocity with less uncertainty due to a better linear fit of the data. While this result is promising, due to the irreproducibility of the robot's path, this result cannot be confirmed until it is demonstrated repeatedly on a more consistent robotic platform or can be compared with a ground truth measurement.

Due to the repeatability issues identified for the robot when following the simple path of moving in a straight line, more complex paths were not tested.




Figure 6.21 Model of the Robot and chair system

6.5.2 DETECTION OF THIN OBSTACLES

In this section a 5.6 mm diameter coaxial cable was strung between two chairs to test the detection of thin obstacles. These thin obstacles are one of the reasons why a full field of view imager is highly desirable for mobile robotics. They are easily missed and reasonably common in real world environments. The repercussions for missing such an obstacle can be very severe for robots both in terms of damage to the robot and safety for people in the area. A photo of the set up is shown in Figure 6.22.

A thin object such as this coaxial cable is unlikely to be detected by any system that only scans along a 2D line. The distance at which the cable will be detected is determined by



Figure 6.22 Photograph of the thin object detection scene

the spatial resolution of the 3D sensor.

Using the system with the WDFT algorithm enabled, the robot was driven towards the wire. Similar to Section 6.5.1, VFRI was not used as guaranteed high frame rates are required. The first frame at which some of the wire is detected is shown in Figure 6.23. The wire has not been completely detected, however some elements of it have. The distance to the wire for this frame is 0.95 m.




Figure 6.23 Initial partial detection of wire using the WDFT system

In the frame following the initial detection, shown in Figure 6.23, the wire was fully detected, this is shown in Figure 6.24. The wire now appears as a solid object in the distance image and is at a distance of 0.75 m.

Using our system the 5.6 mm diameter wire can be detected at a distance of 0.75 m. This distance is determined by the spatial resolution of the sensor and the optical configuration. As sensor resolutions increase this distance will also increase, however the introduction of a larger field of view will decrease this distance.

6.5.3 EVALUATION OF SUITABILITY FOR MOBILE ROBOTICS

The experiments performed in this chapter have demonstrated that while considerable progress has been made in developing a compact, configurable ranging camera using




Figure 6.24 Comprehensive detection of wire using the WDFT system

indirect time of flight technology, there is still some significant future work required to make these cameras suitable for implementation in mobile robotics.

Detection of thin objects was demonstrated, however lateral motion of the robot caused considerable problems with the range images. Future work is required to combine the improvements demonstrated in this thesis with a method for mitigating lateral motion error to make the system viable on mobile robots. A mobile robotic platform with smoother acceleration and deceleration, as well as solving the issues with weaving while trying to move in a straight line would also improve the performance. Issues were identified with phase wrapping that also need to be addressed, as it is likely that a mobile robot will encounter distances longer than the unambiguous measurement distance of the camera. Solutions to this problem have been developed but must be efficiently included with the other solutions described in this chapter. Having sufficient background light suppression to be able to operate in sunlight will also be important for a large number of mobile robotics applications. This can be achieved through a combination of improvements in both the pixel architecture and filters used in these cameras.

Chapter 7 CONCLUSIONS

7.1 Thesis summary

In this thesis, a compact indirect time of flight range imaging system was developed that addresses some of the limitations of existing commercial cameras for mobile applications. As well as the hardware development of the system, two important contributions were made in the operation of indirect time of flight cameras. It was shown that the standard phase detection algorithm could be replaced with a phase detection algorithm that provides benefits both in terms of linearity and the amelioration of axial motion error. A method of operation was developed and experimentally verified where the dynamic range of the camera is improved by varying the effective integration time across the scene depending on the estimated quality of each individual pixel. It was also demonstrated how these improvements could be combined together into a real world system.

7.1.1 DEVELOPMENT OF THE COMPACT VICTORIA UNIVERSITY RANGE IMAGING SYSTEM

One of the significant outputs of this thesis was the development of the compact Victoria University Range Imaging System. The prototype system available at the start of this thesis was bulky, had complex power requirements and included an expensive FPGA development board. Because of these characteristics the system was not suitable for mounting on a mobile platform.

The core design requirement of the new system was to provide a compact, configurable, relatively inexpensive indirect time of flight system that was capable of being powered from a single unregulated power supply and could handle camera resolutions of up to 1 MegaPixel.

The final design was based on four modular boards assembled in a stacked configuration and met all of the requirements set out at the beginning of the project. Modularising the

system into several boards for the different subsystems of the design makes the system more configurable, as an individual board can be replaced in order to upgrade the system. It also decreases the physical footprint of the system.

The system is powered using a single unregulated DC input of 10 – 20 V. Unregulated power is distributed throughout the boards and individually regulated on each board, remaining consistent with the design principle of modularity. This voltage range makes the system suitable to be powered by a single lead acid battery, a common method of providing power for mobile robots.

The processing on the system was performed using a Cyclone III FPGA. This relatively inexpensive FPGA is capable of performing all the real time processing required for the system and, with the addition of on board DDR2 memory, is capable of handling cameras up to 1 MegaPixel.

A comparison between the prototype system and the newly developed system is shown in Figure 7.1. The prototype system has dimensions of 400 mm × 200 mm × 250 mm and the compact system has dimensions of 100 mm × 160 mm × 150 mm (100 mm × 160 mm × 400 mm with cables, which could be easily reduced if required).




Figure 7.1 Comparison between the prototype system (left) and newly developed system (right) (McClymont, 2010)

7.1.2 IMPLEMENTATION OF THE WINDOWED DISCRETE FOURIER TRANSFORM

The standard indirect time of flight phase detection algorithm makes a number of unreasonable assumptions about the properties of both the correlation waveform and

the objects being measured. The invalidation of these assumptions leads to significant systematic errors in the distance measurements recorded using these cameras.

The first assumption made by the standard algorithm is that the correlation waveform is a perfect sinusoid. In real world systems there will always be some harmonic components present in the correlation waveform. These harmonics express themselves as a sinusoidal linearity error with distance. Instead of taking a calibration approach, the implementation of a phase detection algorithm that does not assume a sinusoidal correlation waveform can ameliorate this linearity error without the need for a calibration step.

The second assumption made by the standard phase detection algorithm is that the actual distance to the object is not changing during the measurement time. In real applications, particularly mobile applications, objects in the scene or the camera itself may move, invalidating this assumption. The standard algorithm can be replaced with an algorithm that is insensitive to linear motion of objects in the scene. The error due to objects moving linearly towards or away from the camera, as well as error due to motion of the mobile platform itself, can therefore be mitigated without the requirement for optical flow algorithms that are processor intensive.

Following a literature review and theoretical evaluation, a number of algorithms were experimentally tested and the Windowed Discrete Fourier Transform (WDFT) was selected as the best performing algorithm for general mobile robotics applications. A number of other algorithms for special cases were also explored. The WDFT algorithm was implemented on the FPGA in real time and was demonstrated to be robust over a range of velocities. It was shown that the WDFT offers comparable precision to the standard algorithm.

The implementation of the WDFT algorithm requires trivial additional processing power, does not require significant hardware changes to current indirect time of flight cameras and provides significant benefits over the standard algorithm used ubiquitously throughout the literature and in commercial cameras.

7.1.3 IMPLEMENTATION OF VFRI

Due to the varying reflectivity of, and distance to, objects in real scenes, the intensity of the returning illumination can vary largely between different objects. The dynamic range of indirect time of flight cameras is therefore crucial to acquiring a high precision measurement of an entire scene.

A method for improving the dynamic range of indirect time of flight cameras was proposed in this thesis, called Variable Frame Rate Imaging. This method involved integrating each pixel over a number of integration periods determined by the estimated quality of the pixel. Variable Frame Rate Imaging sacrifices the consistency of the frame rate over the scene in return for improved dynamic range.

Two implementations of Variable Frame Rate Imaging, using different quality metrics, were tested as part of this thesis. The first metric used was the amplitude of the correlation waveform. This proved successful in improving the precision of poorly illuminated areas of a scene, however there was still a large disparity in the precision of different areas of the scene.

The second metric used was the square root of the offset of the correlation waveform divided by its amplitude. It was experimentally verified that this metric was proportional to the precision and using this thresholding an arbitrary minimum precision could be set for the phase measurements across an entire scene. As well as improving the dynamic range of the system, using a precision based threshold is significant for a number of applications such as 3D modelling, where a certain level of precision may be required for the entire model.

As well as the experimental application of this method, a theoretical framework was developed to describe its operation. This theory verifies the experimental results measured using this technique and is useful for demonstrating the cause of the decreased frame rates experienced when compared to increasing the integration time of the camera.

7.1.4 COMBINATION INTO A REAL SYSTEM

A system for real world measurements, capable of operating in real time without any additional processing, was implemented. This combined the improvements from both the Windowed Discrete Fourier Transform and Variable Frame Rate Imaging. A compact real world system was demonstrated that could provide high precision and high accuracy over a large range of distances without performing a calibration or requiring significant additional computational power.

The potential for high precision measurement of high contrast scenes and the ability to measure high quality 3D models of objects was demonstrated. A number of tests were performed using a mobile robot and areas for future work became apparent. The system was successful in demonstrating that a compact configurable indirect time of flight camera suitable for mounting on mobile robots had been designed. It also demonstrated that these cameras are suitable for detecting thin hazards that point or line range measurement systems may not detect. However, before such devices can be used effectively, the improvements in this thesis must be combined with a technique for dealing with lateral motion error.

7.2 Publications Arising from this Thesis

The work in this thesis has resulted in the primary authorship of a journal article submitted to the IEEE Sensors Journal and papers accepted for publication in the proceedings of four conferences. It has also resulted in a book chapter being accepted for publication in the Springer-Verlag series 'Studies in Computational Intelligence'. This work has also contributed to a further conference and book chapter publication where I am listed as a co-author.

Drayton, B.M.M., Carnegie, D.A., & Dorrington, A.A., 2013. Phase algorithms for reducing axial motion and linearity error in indirect time of flight cameras. In *IEEE Sensors Journal*. Submitted 24/10/2012. Resubmitted with minor changes 30/01/2013. Accepted 30/03/2013

Drayton, B., Carnegie, D.A. & Dorrington, A.A., 2011. The development of a time of flight range imager for mobile robotics. In *International Conference on Automation, Robotics and Applications (ICARA)*. Wellington, 2011.

Drayton, B.M.M., Carnegie, D.A. & Dorrington, A.A., 2012. Characterisation of the Victoria University Range Imaging System. In *Australasian Conference on Robotics and Automation*. Wellington, 2012.

Drayton, B.M.M., Carnegie, D.A. & Dorrington, A.A., 2012. Improved Phase Detection Algorithms for Indirect Time of Flight Range Imaging Cameras. In *IEEE International Symposium on Robotic and Sensors Environments*. Magdeburg, 2012.

Drayton, B.M.M., Carnegie, D.A. & Dorrington, A.A., 2013. Variable Frame Time Imaging for Indirect Time of Flight Range Imaging Cameras. In *IEEE International Instrumentation and Measurement Technology Conference*. Minneapolis, 2013.

McClymont, J., Carnegie, D., Jongenelen, A. & Drayton, B., 2010. The Development of a Full-Field Image Ranger System for Mobile Robotic Platforms. In *DELTA Conference*., 2010.

Drayton, B., Carnegie, D.A., & Dorrington, A.A., 2013. Design and Characterisation of a Full-Field Range Imaging Camera for use in Mobile Applications. In *Recent Advances in Robotics and Automation, Series Studies in Computational Intelligence*. Publication pending.

Carnegie, D.A., McClymont, J.R.K., Jongenelen, A.P.P., Drayton, B., Dorrington, A.A., & Payne, A.D., 2011. Design and Construction of a Configurable Full-Field Range Imaging System for Mobile Robotic Applications. In Mukhopadhyay, S.C, Lay-Ekuakille, A. And Fuchs, A. (Eds): *New Developments in Applications in Sensing Technology* pp133-155, Series Lecture Notes in Electrical Engineering, Vol 83, Springer-Verlag, 2011.

7.3 Future Work

This thesis has made significant advances in allowing indirect time of flight cameras to be used for mobile robotics applications. Through real world testing a number of further

requirements have been identified that must be addressed before these cameras can be widely used.

The most significant issue identified in the real world tests is the requirement to combine the WDFT algorithm, which addresses axial motion error, with a solution for lateral motion error. There are a number of methods being developed for this primarily based on optical flow techniques, however algorithmic solutions should also be explored. As five measurements are recorded for the WDFT, and algorithms were demonstrated that could operate using four measurements with the same phase step, there is potential to shrink the number of pixels affected by lateral motion error significantly by simultaneously calculating two four sample results using the first and last four samples of the five sample WDFT. Using this system, lateral motion errors due to the first or last measurement shifting between objects could be rectified, potentially reducing the number of pixels affected by lateral motion error by 40%. However, the viability of this procedure is yet to be explored.

Phase wrapping due to the unambiguous measurement distance of the camera was observed in the real world measurements. A solution to this problem, such as using multiple modulation frequencies to extend the range, should be added to the system.

There are a number of minor improvements that could be made to the hardware of the system, however the most significant limiting factor is the PMD19K-2 sensor. This sensor is now old technology and a replacement sensor is needed to provide higher resolutions and frame rates. Acquiring a new sensor is complicated by the increasing commercialisation of indirect time of flight systems meaning vendors are no longer willing to provide sample sensors for small volumes or research.

Another limitation of the current sensor is the configuration of the modulation drivers. The 40 MHz input frequency limitation means it would be preferable to replace the modulation drivers. Different modulation drivers could also combine the split modulation regions of the sensor, mitigating one source of error.

Further investigation is required into the decreased performance from integrating over multiple frames. There is potential for improving the phase measurement further through statistical modelling of the measurement process.

7.4 Conclusions

Original contributions made by this thesis include:

- The design and construction of a configurable real time indirect time of flight system capable of being mounted on a mobile robot.
- The characterisation of a number of phase detection algorithms and the introduction of theory for predicting the efficacy of algorithms to the indirect time of flight field of research.
- The selection and extensive testing of the WDFT algorithm for the purpose of decreasing linearity error and axial motion error.
- Exploration of the potential for implementing customised phase detection algorithms for specific applications.
- The development of Variable Frame Rate Imaging, a technique where the dynamic range of indirect time of flight cameras can be improved by varying the effective integration time across the scene.
- The development of a theoretical framework for Variable Frame Rate Imaging.
- The extension of the Variable Frame Rate Imaging technique to allow for minimum precision images to be recorded using indirect time of flight cameras.

The objective of this research was to design a configurable indirect time of flight camera that could be mounted on a mobile robot, solve a number of already known issues for time of flight cameras in this domain and to validate this camera using real world tests. The constructed system was successfully mounted on a mobile robotic platform and shown to have promising features for mobile robotics, albeit with future work required.

The final output of this thesis is an indirect time of flight camera with significantly improved precision, linearity and axial motion error insensitivity without the requirement of calibration procedures.

REFERENCES

Acroname Robotics, 2012. *MESA 5m USB SwissRanger product page*. [Online] Available at: <http://www.acroname.com/robotics/parts/R317-SR4000-CW.html> [Accessed 13 December 2012].

Alenya, G., Dellen, B. & Torras, C., 3D modelling of leaves from color and ToF data for robotized plant measuring. In *IEEE International Conference on Robotics and Automation*. Shanghai, 2011, pp. 3408-3414.

Altera, 2011. [Online] Available at: <http://www.buyaltera.com> [Accessed 20 July 2011].

Analog Devices, 2010. [Online] Available at: http://www.analog.com/static/imported-files/data_sheets/AD5301_5311_5321.pdf [Accessed 12 November 2010].

Besl, P.J., Active, optical range imaging sensors. *Machine Vision and Applications*, 1988, 1(2), pp. 127-152.

Binder-Steckver, 2011. [Online] Available at: <http://www.binder-connector.de/en/> [Accessed 20 July 2011].

Blanc, N. et al., Miniaturized smart cameras for 3D-imaging in real-time. In *Proceedings of IEEE Sensors*. Vienna, 2004, pp. 471-474.

Broadbent, E., Stafford, R. & MacDonald, B.A., Acceptance of Healthcare Robots for the Older Population: Review and Future Directions. *International Journal of Social Robotics*, 2009, 1(4), pp. 319-330.

Buttgen, B. et al., 2005. *CCD/CMOS Lock-In Pixel for Range Imaging: Challenges, Limitations and State-of-the-Art*. [Online] Mesa Imaging Available at: <http://www.mesa-imaging.ch/publications.php> [Accessed 10 October 2012].

Carnegie, D., Cree, M., Dorrington, A. & Payne, A., A high resolution full-field range imaging system for robotic devices. In *IASTED International Conference on Robotics and Applications*. Cambridge, Massachusetts, USA, 2005, pp. 306-311. ACTA Press.

Carré, P., Installation et utilisation du comparateur photoélectrique et interférentiel du Bureau International des Poids et Mesures. *Metrologia*, 1966, 2(1), pp. 13-23.

Chiabrando, F., Chiabrando, R., Piatti, D. & Rinaudo, F., Sensors for 3D Imaging: Metric Evaluation and Calibration of a CCD/CMOS Time-of-Flight Camera. *Sensors*, 2009, 9(12), pp. 10080-10096.

Chiabrando, F., Piatti, D. & Rinaudo, F., New Sensors for Cultural Heritage Metric Survey: the ToF Cameras. In *CIPA Symposium*. Prague, 2011, pp. 300-307.

Christie, S. et al., Design and development of a multi-detecting two-dimensional ranging sensor. *Measurement Science and Technology*, 1995, 6(9), pp. 1301-1308.

Conroy, R.M., Dorrington, A.A., Künnemeyer, R. & Cree, M.J., Range imager performance comparison in homodyne and heterodyne operating modes. In *Proc. SPIE*, 7239., 2009, p.723905.

Digi-key, 2012. *D-Imager web page*. [Online] Available at: <http://www.digikey.co.nz/product-detail/en/EKL3105/255-3502-ND/2816058> [Accessed 12 13 2012].

Dorrington, A.A. et al., Separating true range measurements from multi-path and scattering interference in commercial range cameras. In *Proc. SPIE*, 7864., 2011, p.786404.

Drayton, B., Carnegie, D.A. & Dorrington, A.A., The development of a time of flight range imager for mobile robotics. In *International Conference on Automation, Robotics and Applications (ICARA)*. Wellington, 2011, pp. 470-475.

Drayton, B.M.M., Carnegie, D.A. & Dorrington, A.A., Improved Phase Detection Algorithms for Indirect Time of Flight Range Imaging Cameras. In *IEEE International Symposium on Robotic and Sensors Environments*. Magdeburg, 2012 a, pp. 156-161.

Drayton, B.M.M., Carnegie, D.A. & Dorrington, A.A., Characterisation of the Victoria University Range Imaging System. In *Australasian Conference on Robotics and Automation*. Wellington, 2012 b, pp. 1-8.

- Drayton, B.M.M., Carnegie, D.A. & Dorrington, A.A., Variable Frame Time Imaging for Indirect Time of Flight Range Imaging Cameras. In *IEEE International Instrumentation and Measurement Technology Conference*. Minneapolis, 2013, pp. 609-613.
- Falie, D. & Buzuloiu, V., Noise Characteristics of 3D Time-of-Flight Cameras. In *International Symposium on Signals, Circuits and Systems*. Iasi, 2007, pp. 1-4.
- Foix, S., Alenya, G. & Torras, C., Lock-in Time-of-Flight (ToF) Cameras: A Survey. *IEEE Sensors*, 2011, 11(9), pp. 1917 - 1926.
- Frank, M. et al., Theoretical and experimental error analysis of continuous-wave time-of-flight range cameras. *Optical Engineering*, 2009, 48(1), p.013602.
- Fuchs, S. & Hirzinger, G., Extrinsic and depth calibration of ToF-cameras. In *IEEE Conference on Computer Vision and Pattern Recognition*. Anchorage, 2008, pp. 1-6.
- Godbaz, J.P., Cree, M.J. & Dorrington, A.A., Understanding and Ameliorating Non-Linear Phase and Amplitude Responses in AMCW Lidar. *Remote Sensing*, 2011, 4(1), pp. 21-42.
- Gokturk, B. & Rafii, A., Method and system to increase dynamic range of time-of-flight (TOF) and/or imaging sensors. *U.S. Patent 7 379 100*, 2008.
- Gokturk, S.B., Yalcin, H. & Bamji, C., A Time-Of-Flight Depth Sensor - System Description, Issues and Solutions. In *Computer Vision and Pattern Recognition Workshop*. Washington, 2004, p.35.
- Guomundsson, S.A., Aanaes, H. & Larsen, R., Environmental Effects on Measurement Uncertainties of Time-of-Flight Cameras. In *International Symposium on Signals, Circuits and Systems*. Iasi, 2007, pp. 1-4.
- Hahne, U. & Alexa, M., Exposure Fusion for Time-Of-Flight Imaging. In *Pacific Conference on Computer Graphics and Applications*. Kaohsiung, 2011, pp. 1887-1894. Eurographics Association.
- Hariharan, P., Digital phase-stepping interferometry: effects of multiply reflected beams. *Applied Optics*, 1987, 26(13), pp. 2506-2507.

Hibino, K., Oreb, B.F., Farrant, D.I. & Larkin, K.G., Phase shifting algorithms for nonlinear and spatially nonuniform phase shifts. *Journal of the Optical Society of America*, 1997, 14(4), pp. 917-930.

Hussmann, S., Hermanski, A. & Edeler, T., Real-Time Motion Artifact Suppression in TOF Camera Systems. *IEEE Transactions on Instrumentation and Measurement*, 2011, 60(5), pp. 1682-1690.

iC-Haus GmbH, 2011. *iC-HKB 155 MHz Laser Switch datasheet*. [Online] Available at: http://www.ichaus.de/upload/pdf/Hk_files.pdf [Accessed 18 Oct 2011].

Jakevicius, L. & Demcenko, A., Ultrasound attenuation dependence on air temperature in closed chambers. *Ultragarsas (Ultrasound)*, 2008, 63(1), pp. 18-22.

Jongenelen, A.P.P., *Development of a Compact, Configurable, Real-time Range Imaging System*. PhD Thesis. Victoria University, Wellington. 2010.

Jongenelen, A.P.P., Carnegie, D.A., Payne, A.D. & Dorrington, A.A., Development and characterisation of an easily configurable range imaging system. In *International Conference Image and Vision Computing New Zealand*. Washington, USA, 2009, pp. 379-384.

Kahlmann, T., Remondino, F. & Ingensand, H., Calibration for Increased Accuracy of the Range Imaging Camera Swissranger. In *Image Engineering and Vision Metrology*. Dresden, 2006, pp. 136-141.

Katz, D.S. & Some, R.R., NASA Advances Robotic Space Exploration. *Computer*, 2003, 36(1), pp. 52-61.

Khoshelham, K. & Elberink, S.O., Accuracy and Resolution of Kinect Depth Data for Indoor Mapping Applications. *Sensors*, 2012, 12(2), pp. 1437-1454.

Kim, Y.M., Chan, D., Theobalt, C. & Thrun, S., Design and calibration of a multi-view TOF sensor fusion system. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. Anchorage, 2008, pp. 1-7.

- Lange, R. & Seitz, P., Solid-State Time-of-Flight Range Camera. *IEEE Journal of Quantum Electronics*, 2001, 37(3), pp. 390-397.
- Larkin, K.G. & Oreb, B.F., Design and assesment of symmetrical phase-shifting algorithms. *Journal of the Optical Society of America A: Optocs, Image Science, and Vision*, 1992, 9(10), pp. 1740-1748.
- Lehmann, M. et al., Smart pixels for future 3D-TOF sensors. In *IEEE workshop on charge-coupled devices and advanced images sensors*. Karuizawa, 2005, pp. 193-196.
- Lindner, M., *Calibration and Real-Time Processing of Time-of-Flight Range Data*. PhD Thesis. University of Siegen, Seigen. 2010.
- Lindner, M. & Kolb, A., Lateral and Depth Calibration of PMD-Distance Sensors. In *Second international conference on Advances in Visual Computing*. Heidelberg, 2006, pp. 524-533. Springer-Verlag Berlin.
- Lindner, M. & Kolb, A., Calibration of the intensity-related distance error of the PMD TOF-camera. In *SPIE Intelligent Robots and Computer Vision XXV*. Boston, 2007, p.67640.
- Lindner, M. & Kolb, A., Compensation of Motion Artifacts for Time-of-Flight Cameras. In *DAGM Workshop on Dynamic 3D Imaging*. Siegen, Germany, 2009, pp. 16-27. Springer-Verlag Berlin.
- Lindner, M., Kolb, A. & Ringbeck, T., New Insights into the Calibration of ToF-Sensors. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. Anchorage, 2008, pp. 1-5.
- Lottner, O., A.Sluite, Hartmann, K. & W.Weih, Movement Artefacts in Range Images of Time-of-Flight Cameras. In *International Symposium on Signals, Circuits and Systems*. Iasi, 2007, pp. 1-4.
- Macron Dynamics Inc., 2011. [Online] Available at: <http://catalog.macrondynamics.com/> [Accessed 20 Jul 2011].
- May, S. et al., Three-dimensional mapping with time-of-flight cameras. *Journal of Field Robotics*, 2009, 26(11-12), pp. 934-965.

May, S., Werner, B., Surmann, H. & Pervolz, K., 3D time-of-flight cameras for mobile robotics. In *Proc. IEEE International Conference on Intelligent Robots and Systems*. Beijing, 2006, pp. 790-795.

McClymont, J., *Development of Extrospective Systems for Mobile Robotic Vehicles*. Masters Thesis. Victoria University, Wellington. 2010.

McClymont, J., Carnegie, D., Jongenelen, A. & Drayton, B., The Development of a Full-Field Image Ranger System for Mobile Robotic Platforms. In *IEEE International Symposium on Electronic Design, Test and Application*. Queenstown, 2010, pp. 128-133.

Mesa Imaging, 2011. *Swissranger API Documentation*. [Online] Mesa Imaging (1.6.1) Available at: <http://www.mesa-imaging.ch/drivers.php> [Accessed 16 August 2012].

Murphy, R.R., Trial by fire [rescue robots]. *IEEE Robotics and Automation Magazine*, 2004, 11(3), pp. 50-61.

Niclass, C. et al., Single-photon synchronous detection. *IEEE Journal of Solid-state circuits*, 2008, 44(7), pp. 114-117.

Novák, J., Novák, P. & Antonín, M., Multi-step phase-shifting algorithms insensitive to linear phase shift errors. *Optics Communications*, 2008, 281(21), pp. 5302-5309.

Payne, A.D., Dorrington, A.A. & Cree, M.J., Illumination waveform optimization for time-of-flight range imaging cameras. In Remondino, F. & Shortis, M.R., eds. *Proc. SPIE, 8085*. Munich, 2011, p. 80850.

Payne, A.D., Dorrington, A.A., Cree, M.J. & Carnegie, D.A., Improved linearity using harmonic error rejection in a full-field range imaging system. In *Proc. SPIE, 6805*. San Jose, 2008, p.68050.

Payne, A.D., Dorrington, A.A., Cree, M.J. & Carnegie, D.A., Characterization of modulated time-of-flight range image sensors. *Three-Dimensional Imaging Metrology*, 2009, 7239, p.723904.

Payne, A.D., Dorrington, A.A., Cree, M.J. & Carnegie, D.A., Improved measurement linearity and precision for AMCW time-of-flight range imaging cameras. *Applied Optics*, 2010, 49(23), pp. 4392-4403.

Payne, A.D. et al., Multiple Frequency Range Imaging to Remove measurement ambiguity. In *9th Conference on Optical 3-D Measurement Techniques*. Vienna, 2009, pp. 139-148.

PMD Technologies, *PhotonICs PMD 19K-2 Datasheet*. Datasheet. PMD Technologies. 2008.

PMD, T., 2010. *PMD[vision] CamCube 3.0 Datasheet*. [Online] Available at: http://www.pmdtec.com/fileadmin/pmdtec/downloads/documentation/datenblatt_camcube3.pdf [Accessed 28 February 2011].

Pomares, J., Gil, P. & Torres, F., Visual Control of Robots Using Range Images. *Sensors*, 2010, 10(8), pp. 7303-7322.

Radmer, J., Fuste, P., Schmidt, H. & Kruger, J., Incident Light Related Distance Error Study and Calibration of the PMD-Range Imaging Camera. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. Anchorage, 2008, pp. 1-6.

Rapp, H., *Experimental and Theoretical Investigation of Correlating TOF-Camera Systems*. Masters Thesis. University of Heidelberg, Germany. 2007.

Ringbeck, T. & Hagebeuker, B., A 3D time of flight camera for object detection. In *Optical 3D measurement techniques*. Zurich, 2007, pp. 135-146.

Routh, H.F. & Skyba, D.M., Functional imaging with ultrasound. *Medica Mundi*, 2001, 46, pp. 59-64.

Schreiber, H. & Bruning, J.H., 2006. *Phase Shifting Interferometry Chapter, in Optical Shop Testing*. Third Edition ed. Hoboken, NJ, USA: John Wiley & Sons, Inc.

SensComp, 2011. *Complete Ultrasonic Ranging Solutions webpage*. [Online] Available at: http://www.senscomp.com/products/complete-ultrasonic-ranging-solutions.php#x1_ranging_module_pro_developers_kit [Accessed 12 February 2013].

Softkinetic, 2012. *Softkinetic DepthSense camera page*. [Online] Available at: <http://www.softkinetic.com/en-us/solutions/depthsensecameras.aspx> [Accessed 13 December 2012].

Surrel, Y., Phase stepping: a new self-calibrating algorithm. *Applied Optics*, 1993, 32(19), pp. 3598-3600.

Surrel, Y., Design of algorithms for phase measurements by the use of phase stepping. *Applied Optics*, 1996, 35(1), pp. 51-60.

Voth, D., A new generation of military robots. *IEEE Intelligent Systems*, 2004, 19(4), pp. 2-3.

Wiedemann, M., Sauer, M., Driewer, F. & Schilling, K., Analysis and Characterization of the PMD Camera for Application in Mobile Robotics. In Chung, M.J. & Misra, P., eds. *Proceedings of the 17th World Congress of the International Federation of Automatic Control*. Seoul, 2008, pp. 13689-13694.

Appendix A PCB LAYOUTS

A.1 Compact system




Figure A.1.1 FPGA Board Top Layer




Figure A.1.1: FPGA Board Mid Layer 1




Figure A.1.2: FPGA Board Mid Layer 2




Figure A.1.3 : FPGA Board Bottom Layer

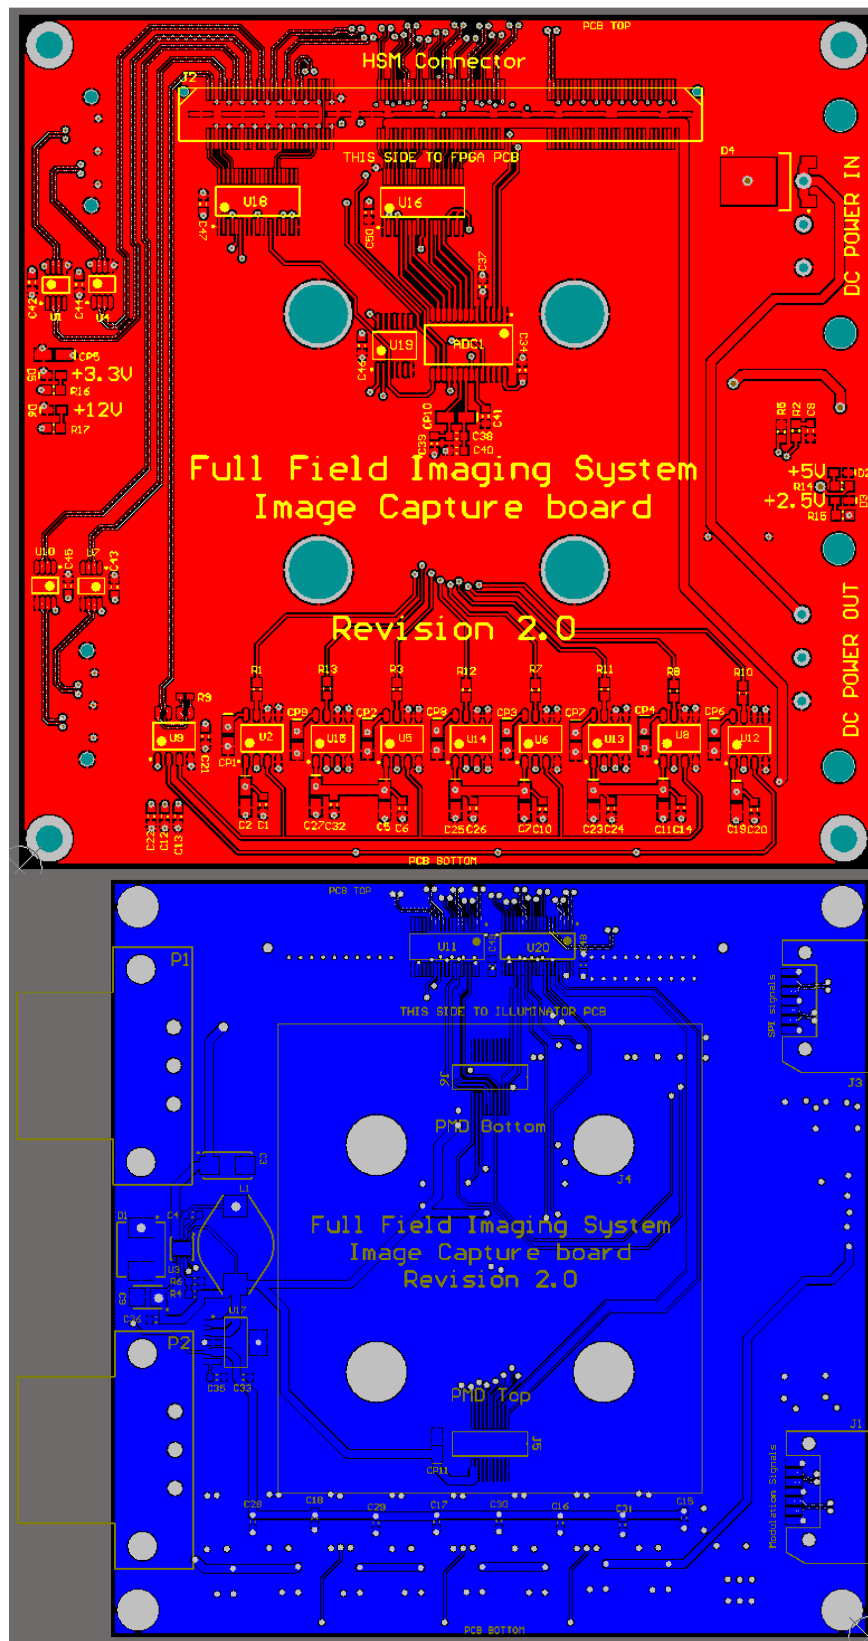


Figure A.1.4 : Image Capture Board top (top) and bottom (bottom) layers

A.2 Compact system revision 2.1




Figure A.2.1 : FPGA Board Rev 2.1 Top Layer

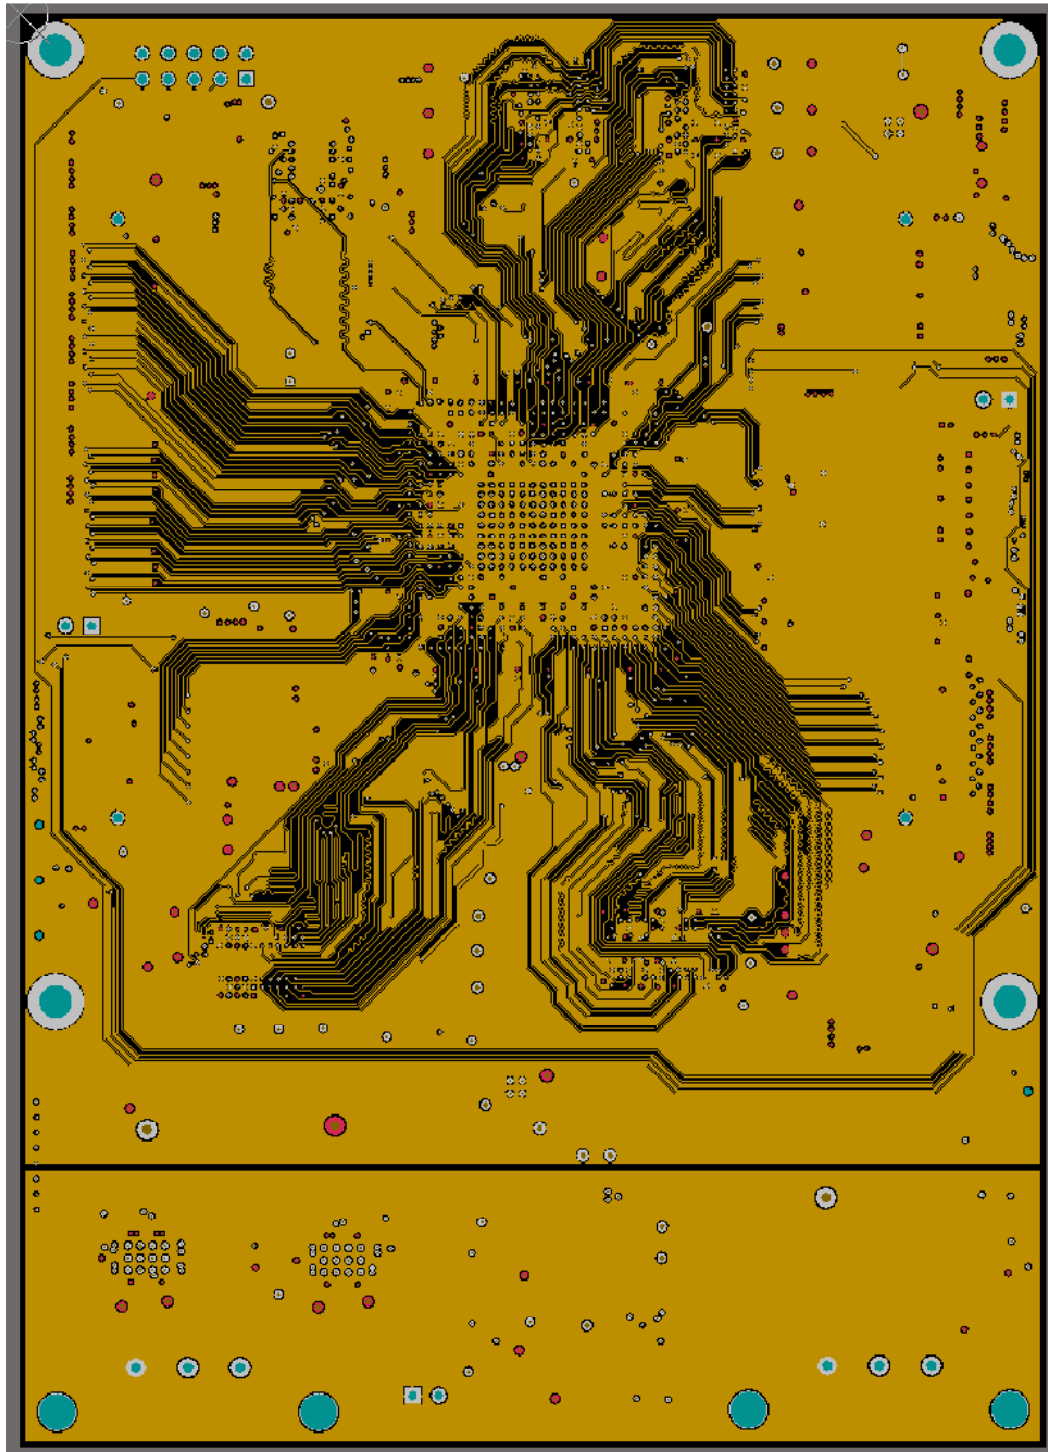


Figure A.2.2: FPGA Board Rev 2.1 Mid Layer 1




Figure A.2.3 : FPGA Board Rev 2.1 Mid Layer 2




Figure A.2.4 : FPGA Board Rev 2.1 bottom layer

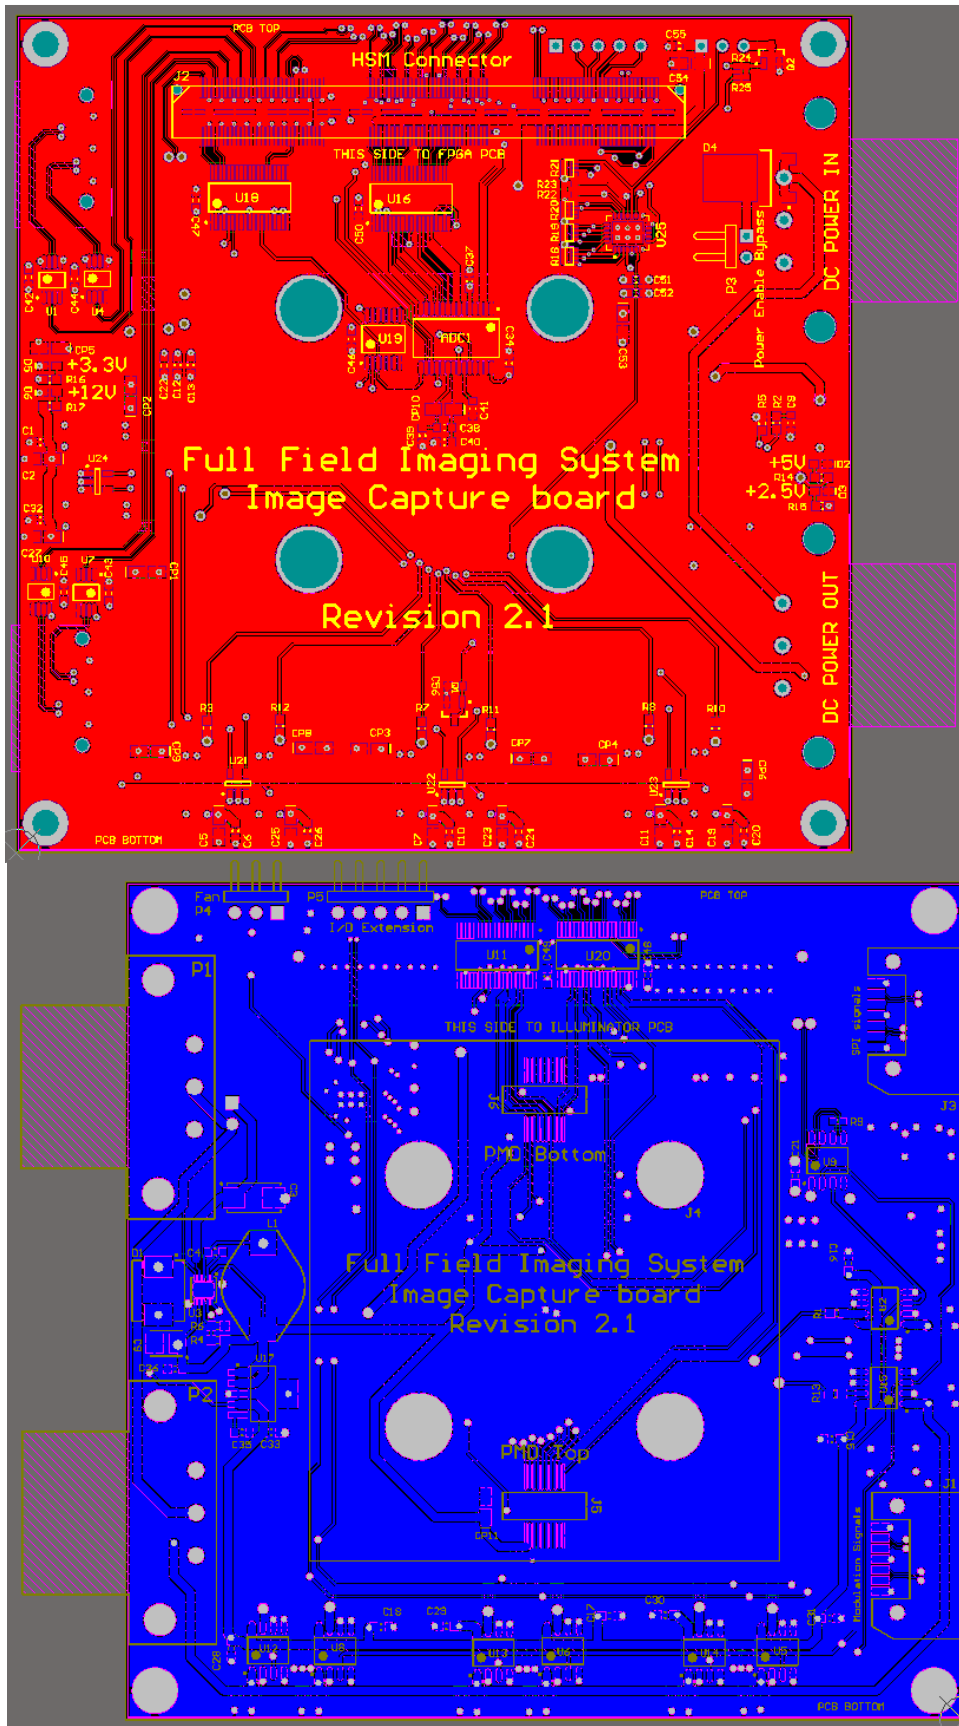


Figure A.2.5 : Image Capture Board revision 2.1 data layers