Development of a Full-Field Time-of-Flight Range Imaging System

A thesis submitted in fulfilment

of the requirements for the degree

of

Doctor of Philosophy in

Physics and Electronic Engineering

at

The University of Waikato

by

Andrew Dean Payne

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

2008

Abstract

A full-field, time-of-flight, image ranging system or "3D camera" has been developed from a proof-of-principle to a working prototype stage, capable of determining the intensity and range for every pixel in a scene. The system can be adapted to the requirements of various applications, producing high precision range measurements with sub-millimetre resolution, or high speed measurements at video frame rates. Parallel data acquisition at each pixel provides high spatial resolution independent of the operating speed.

The range imaging system uses a heterodyne technique to indirectly measure time of flight. Laser diodes with highly diverging beams are intensity modulated at radio frequencies and used to illuminate the scene. Reflected light is focused on to an image intensifier used as a high speed optical shutter, which is modulated at a slightly different frequency to that of the laser source. The output from the shutter is a low frequency beat signal, which is sampled by a digital video camera. Optical propagation delay is encoded into the phase of the beat signal, hence from a captured time variant intensity sequence, the beat signal phase can be measured to determine range for every pixel in the scene.

A direct digital synthesiser (DDS) is designed and constructed, capable of generating up to three outputs at frequencies beyond 100 MHz with the relative frequency stability in excess of nine orders of magnitude required to control the laser and shutter modulation. Driver circuits were also designed to modulate the image intensifier photocathode at 50 V_{pp}, and four laser diodes with a combined power output of 320 mW, both over a frequency range of 10–100 MHz. The DDS, laser, and image intensifier response are characterised. A unique method of measuring the image intensifier optical modulation response is developed, requiring the construction of a pico-second pulsed laser source. This characterisation revealed deficiencies in the measured responses, which were mitigated through hardware modifications where possible. The effects of remaining imperfections, such as modulation waveform harmonics and image intensifier irising, can be calibrated and removed from the range measurements during software processing using the characterisation data.

Finally, a digital method of generating the high frequency modulation signals using a FPGA to replace the analogue DDS is developed, providing a highly integrated solution, reducing the complexity, and enhancing flexibility. In addition, a novel modulation coding technique is developed to remove the undesirable influence of waveform harmonics from the range measurement without extending the acquisition time. When combined with a proposed modification to the laser illumination source, the digital system can enhance range measurement precision and linearity.

From this work, a flexible full-field image ranging system is successfully realised. The system is demonstrated operating in a high precision mode with sub-millimetre depth resolution, and also in a high speed mode operating at video update rates (25 fps), in both cases providing high (512×512) spatial resolution over distances of several metres.

Acknowledgements

I owe thanks to many people for providing academic and emotional support throughout the duration of this project. First and foremost I want to express my gratitude for the continual inspiration, motivation, help and support I have received from my supervisors. To Dale Carnegie, who encouraged me to continue my academic career and the challenge of undertaking a PhD, and to Michael Cree and Adrian Dorrington who took over the reins through the later stages of study. I have really enjoyed being part of the team in such an enthusiastic environment, where I am captivated by the continual generation of new ideas and advances, never quite knowing what to expect next!

This work would not have been possible without the financial support provided by the Tertiary Education Commission (Top Achiever Doctoral Scholarship), and the University of Waikato (Waikato Doctoral Scholarship), or without the kind generosity of Geoff Austin and the guys in the Department of Physics at The University of Auckland who provided a stimulating place to work for the duration of my study.

I would like to give special thanks to my parents, Dean and Wendy, who have always supported me in what I've done. Your valiant attempts to understand my work are highly commendable, and I appreciate all of the encouragement you have given me over the years.

And last, but definitely not least, I want to acknowledge the support of my wife Jenny – look Jenny I finally made it! I dedicate this thesis to you, and your infinite patience and understanding as deadlines were extended, and my work was continuously "almost finished". Without your support I would have never gotten this far. Thank you!

To all others who have contributed but whose names I fail to mention here, thank you very much.

Table of Contents

List of Tables

List of Figures

1. Introduction

Seeing the world in three dimensions is something that we often take for granted, a virtue that has proven challenging for computer machine-vision applications where imaging sensors often provide a flat, two dimensional view of the world. Numerous means have been employed to extend image measurements into the three-dimensional realm, leading to familiar solutions such as stereo vision and laser scanning. The continual emergence of new applications requiring depth data often surpasses the capabilities of these existing products and technologies, commanding attributes such as higher accuracy and precision, reduced acquisition time, and smaller size and costs.

Some typical applications where range data forms an integral part of the sensing system include obstacle detection for automotive safety (Hsu et al., 2006), navigation and path planning of autonomous vehicles through unknown terrain (Juberts and Barbera, 2004; Kolski et al., 2006), or search and rescue applications (Bostelman et al., 2005). In other applications, hand or body motion can be tracked, for example as an input for controlling interactive games (Breuer et al., 2007), fruit is readily segmented from cluttered surroundings for automated harvesting (Jimenez et al., 2000), and real-time generation of TV special effects have been demonstrated (Kawakita et al., 2004c). These types of applications involve motion, either within the scene, or of the sensor itself, therefore requiring rapid acquisition to avoid measurement error due to motion blur.

Alternatively, emphasis on spatial and depth resolution is important for applications such as quality control of manufactured products (Wolf et al., 2000), spacecraft inspection

(Nellums et al., 2006), surveying (Nagihara et al., 2004), facial recognition (Chang et al., 2003) and medical imaging, for example where facial scanning can be used to identify syndromes for genetic diseases (Hammond et al., 2004). The static or limited motion contained in these scenes alleviates the need for high speed capture, but precision is vital.

Each of these existing, and future, applications have a variety of distinct requirements which cannot always be met using traditional range measurement techniques. Geometric based systems, such as stereo imaging, can experience occlusion problems when the sensor view is obstructed by objects in the scene, and laser scanning systems are inherently slow due to the mechanical scanning employed. A new technology has recently emerged to address some of these issues; a time of flight system capable of providing simultaneous depth capture of an entire scene. Rather than serially scanning a collimated laser beam, the acquisition time can be reduced by expanding the illumination to encompass the entire scene, while the flight time of each point in the field of view is measured in parallel with an array of receivers.

Time of flight distance measurement is conceptually straightforward: a signal is emitted from a source, and the time required for the signal to propagate to a target, and be reflected back to a receiver, is measured. The target distance is then calculated from the known speed of the signal. This is a technique which can be found in nature; dolphins, whales, bats and other animals use a time of flight method known as echolocation for navigation and foraging by emitting sounds and listening for the delayed reflections. The difficulty is that in an optical measurement system, light travels at 299,792,458 ms^{-1}; hence measuring the propagation delay over short ranges is a formidable challenge. For example, a distance of 1 cm requires measurement of the time delay with resolution better than 70 ps (taking into account the fact that the light travels the twice the distance, to and from the target). To overcome the difficulty of performing measurements on such a short time scale, *indirect* time of flight techniques have been developed, where the object range is encoded into a variable other than time, such as intensity, frequency, or phase, greatly simplifying the electronic requirements of the measurement system. One such indirect time of flight system utilises a heterodyne measurement technique, where the range is encoded into the phase of a low frequency signal (Carnegie et al., 2005), demonstrating centimetre range resolution. The heterodyne system uses an amplitude modulated light source operating at a

high frequency to flood illuminate the entire scene. The flight time of the light travelling to, and being reflected back from, objects within the scene results in retardation of the illumination modulation envelope phase proportional to the distance travelled. A high speed optical shutter acts as a mixer to produce a low frequency beat signal output using the heterodyne technique. The low frequency signal is then digitised using a standard digital video camera, from which the phase, and therefore range, can be calculated individually for every pixel from the captured time variant intensity sequence

The range imaging system demonstrated by Carnegie et al. uses a pair of synchronised bench top frequency synthesisers to generate the two high frequency modulation signals to control the illumination and optical shutter modulation. An array of amplitude modulated LEDs, operating at frequencies up to 20 MHz, are utilised to flood illuminate the scene, and the reflected light is focused onto the input window of an image intensifier operating as a high speed optical shutter. The optical gain of the image intensifier is controlled by a 4 V_{pp} sinusoidal voltage applied to the photocathode, providing optical gain modulation at frequencies up to 30 MHz. The output from the image intensifier is focused onto a CCD camera, allowing acquisition software on a PC to capture the image frames and calculate the phase (range) independently for each pixel.

1.1 Objectives

The objective of this work is to enhance the existing range imaging system, creating a high performance, flexible range measurement device. To adapt to the needs of existing and upcoming applications, the system may be required, for example, to perform range imaging with millimetre resolution in one instance, while in another instance produce measurements at video update rates. The measurement spatial resolution is, at least in principle, independent of the acquisition period due to the parallel nature of the system; hence a resolution of at least 512×512 pixels should be attained for either high precision or high speed applications. The previously demonstrated heterodyne system produces relative range measurements due to an arbitrary phase offset in each measurement, requiring the inclusion of a known reference object within the scene to determine absolute range. This limitation should be overcome to allow absolute range measurements to be made in arbitrary scenes. To achieve these goals requires implementation of a number of

key hardware and processing enhancements to the existing heterodyne range imaging system.

To achieve millimetre depth resolution, the illumination source and shutter are required to operate at higher modulation frequencies, approaching 100 MHz. The existing hardware is unable to operate at these speeds; hence development of a new illumination source and image intensifier driver circuit is required. By increasing the image intensifier modulation voltage by an order of magnitude, range measurement precision and spatial resolution can be enhanced; hence an image intensifier driver circuit capable of operating at high speed (100 MHz) and high voltage (50 V) is desired. The maximum operating frequency of the LEDs in the existing system is constrained by their response time. An alternate source, such as laser illumination, must therefore be investigated and implemented. To generate absolute range measurements an electronic timing reference signal is required, which must be synchronised to the high frequency illumination and shutter modulation signals. To achieve this, and to alleviate the dependency on bulky bench-top signal generators, a circuit is required to generate three synchronised frequency-locked signals derived from a single clock source, producing output frequencies up to 100 MHz with precise (1 Hz) frequency tuning.

Once these hardware enhancements have been completed, characterisation is required to verify correct operation of the individual components in the system, as well as the performance of the system as a whole. By characterising components individually, undesirable traits can be identified which may otherwise be obscured in a measurement of the entire system. By understanding the causes of adverse characteristics in the measured responses, their effects can, in some cases, be mitigated through modification of the hardware. However, this is not always practical, and in these instances the characterisation data can be used to simulate and calibrate the system to reduce the resultant range measurement error.

To increase the prospect for commercialisation, miniaturisation of the hardware is desirable. The heterodyne method is inherently analogue in nature, requiring sinusoidal waveforms at different frequencies, using analogue amplifiers, filters, etc., resulting in the use of complicated hardware. In contrast, digital systems can be highly integrated into an IC design, simplifying the hardware requirement, while reducing manufacturing cost, size,

and complexity. A digital phase stepping technique is investigated to emulate the analogue heterodyne signals, reducing the need for analogue components in the hardware design. A unique digital control method is developed, and the design is implemented in a FPGA. The new digital technique overcomes limitations within the measured illumination and image intensifier responses, increasing range measurement linearity compared to the analogue control system, without the need to extend the acquisition period.

1.2 Outline

This work is divided into eight chapters:

Chapter 2 investigates optical range measurement techniques, with emphasis on full-field methods capable of simultaneous capture of the entire scene. The basic working principles of interferometric, geometric (stereo vision and triangulation), time of flight and other range measurement techniques are examined, and the limitations of each method are discussed. Direct and indirect time of flight measurement methods are discussed and evaluated, where indirect methods encode object range into a variable other than time in order to simplify the otherwise stringent electronic requirements. The influence of undesirable contaminants, such as ambient light, as well as sensor limitations such as quantisation error and limited dynamic range are identified as appropriate for each technique. The reported performance of these methods is reviewed for a variety of commercially available systems and systems under development.

Parameters affecting measurement precision for the heterodyne system are identified in **Chapter 3**, namely the modulation frequency, the modulation index, and the signal to noise ratio. Optimisation of these terms is therefore crucial to achieving range measurements with the desired millimetre resolution. Range ambiguity is addressed, where, due to phase being cyclic, the same phase value can result from objects at different distances. Finally, deviations from ideal sinusoidal illumination and shutter waveforms are examined as these can compromise measurement linearity due to the inclusion of higher order harmonics.

Chapter 4 outlines the electronic enhancements required for the range imaging system to achieve the desired millimetre measurement performance. Requiring two frequency-locked signals with a small tuning difference to control the illumination source and optical

shutter, a circuit is designed using direct digital synthesiser (DDS) ICs, with outputs operating at frequencies up to 100 MHz. Addition of a third reference timing signal is also incorporated to perform absolute range measurements. An image intensifier is used as a voltage-dependent optical shutter element, and the operating mechanism is explained. Chapter 3 informs us that measurement precision is dependent, among other parameters, on the modulation frequency; therefore a driver circuit is designed to enhance the image intensifier operating frequency range from 20 MHz to 100 MHz. The image intensifier modulation voltage is also increased by an order of magnitude to enhance the spatial resolution of the system.

A corresponding increase in modulation frequency is required from the illumination source. **Chapter 5** finds that the response time of standard LEDs inhibit operation at the modulation frequencies required, and that laser diodes offer high output power and intensity modulation bandwidth at relatively low cost. A circuit is designed to drive four laser diodes, with the output illumination coupled to the system through optical fibres. Laser eye safety concerns are also addressed. A number of auxiliary signals are required to control the various electronic components within the system, suited to implementation using a FPGA.

In Chapter 3 we find that the shape of the generated illumination and shutter waveforms influence the measurement precision and linearity; hence **Chapter 6** sets out to characterise and optimise the response of the laser and image intensifier components, as well as quantifying the performance of the DDS and the combined range measurement system. The DDS and laser are simply measured using standard test equipment, whereas optical characterisation of the image intensifier response is more complicated. A unique measurement method is developed using a picosecond pulsed laser source. From the characterisation results, modifications are made to the hardware to mitigate range measurement errors and maximise measurement precision.

Chapter 7 investigates the effect of timing errors, known as jitter, in the high frequency modulation signals, and proposes generating the signals from a digital-only source. A digital design can be highly integrated into an IC, reducing the size, cost, and complexity of the resultant system. A number of different digital control schemes are investigated to determine the effect on range measurement precision and accuracy, resulting in the

generation of a novel control scheme capable of not only replacing the analogue system, but also outperforming it. The new technique is implemented in an FPGA, and the experimental results are compared to those predicted through simulation.

The project is concluded in **Chapter 8** with some example acquisitions of real life scenes, a brief summary, and an assessment of further possible development and improvements.

1.3 Publications Arising from this Thesis

The work presented in this thesis has directly resulted in the submission of a patent application, the primary authorship of a journal paper and four papers published in conference proceedings. A second journal paper titled "Enhancing Linearity and Precision in a Full-Field Range Imaging Camera" is in preparation and will be submitted to Applied Optics. This work has also resulted in contributions towards a further journal paper and four conference papers.

List of publications:

- Payne, A. D. (2007). "Harmonic rejection sampling methodology," New Zealand Provisional Patent Application No. 562739, filed 19 October 2007.
- Payne, A. D., Dorrington, A. A., Cree, M. J. and Carnegie, D. A. (2008). "Characterizing an image intensifier in a full-field image ranging system," *IEEE Sensors Journal*, accepted for publication.
- Payne, A. D., Dorrington, A. A., Cree, M. J. and Carnegie, D. A. (2008). "Improved Linearity Using Harmonic Error Rejection in a Full-Field Range Imaging System," *Proc. SPIE - 3D Image Capture and Applications VII*, vol. 6805, pp. 68050D.
- Payne, A. D., Dorrington, A. A., Cree, M. J. and Carnegie, D. A. (2006). "Image Intensifier Characterisation," *Proc. Image and Vision Computing New Zealand 2006 (IVCNZ'06)*, pp. 487-492.
- Payne, A. D., Carnegie, D. A., Dorrington, A. A. and Cree, M. J. (2006). "Full Field Image Ranger Hardware." *Proc. Third IEEE International Workshop on Electronic Design, Test and Applications (DELTA'06)*, pp. 263-268.

- Payne, A. D., Carnegie, D. A., Dorrington, A. A. and Cree, M. J. (2005). "A Synchronised Direct Digital Synthesiser," *Proc. First International Conference on Sensing Technology (ICST)*, pp. 174-179.
- Dorrington, A. A., Cree, M. J., Carnegie, D. A., Payne, A. D., Conroy, R. M., Godbaz, J. P. and Jongenelen, A. P. P. (2008). "Video-rate or high-precision: a flexible range imaging camera," *Proc. SPIE - Image Processing: Machine Vision Applications*, vol. 6813, pp. 681307.
- Dorrington, A. A., Cree, M. J., Payne, A. D., Conroy, R. M. and Carnegie, D. A. (2007). "Achieving sub-millimetre precision with a solid-state full-field heterodyning range imaging camera," *Measurement Science & Technology*, vol. 18, pp. 2809-2816.
- Dorrington, A. A., Cree, M. J., Carnegie, D. A., Payne, A. D. and Conroy, R. M. (2007). "Heterodyne range imaging as an alternative to photogrammetry," *Proc. SPIE 6491 – Videometrics IX*, San Jose, CA, USA, pp. 64910D.
- Cree, M. J., Dorrington, A. A., Conroy, R. M., Payne, A. D. and Carnegie, D. A. (2006). "The Waikato Range Imager," *Proc. Image and Vision Computing New Zealand (IVCNZ'06)*, pp. 233-238.
- Dorrington, A. A., Cree, M. J., Carnegie, D. A. and Payne, A. D. (2005). "Selecting Signal Frequencies for Best Performance of Fourier-based Phase Detection," *Proc. Twelfth New Zealand Electronics Conference*, pp. 189–193.

2. Optical Range Measurement

This chapter focuses on optical range measurement techniques. The short wavelength of light allows optical ranging systems to produce measurements with excellent angular resolution, outperforming competing technologies such as ultrasound or microwave due to diffraction limitations. The chapter outlines some of the techniques developed to generate a range map, or 3D image of a scene.

Emphasis is given to simultaneous full-field acquisition methods, where the system is capable of making measurements in parallel, reducing the acquisition time while producing high resolution results. Some techniques are inherently capable of achieving this goal, while others are an adaptation of single point measurement techniques into a full-field configuration.

2.1 Full-Field Range Measurement Techniques

A number of optical ranging techniques are presented, each with their associated advantages and disadvantages. An overview is given in Figure 2.1, where the techniques are divided into various categories. Optical range measurement techniques are often categorised as “active” or “passive”, however, the classification scheme chosen here does not employ that distinction as measurement deficiencies, such as shadowing, can be common to both active and passive systems.

The requirements of a particular application inherently determine the suitability of each technique, including the range to be measured (μm–km), accuracy, acquisition time, cost,

whether absolute or relative measurements are required, etc. Applications of interest for this project are typically in the range of 1–100 m, and require an absolute measurement to be made with high accuracy. Low cost is an important factor to allow the sensor to be used in a wide range of applications.

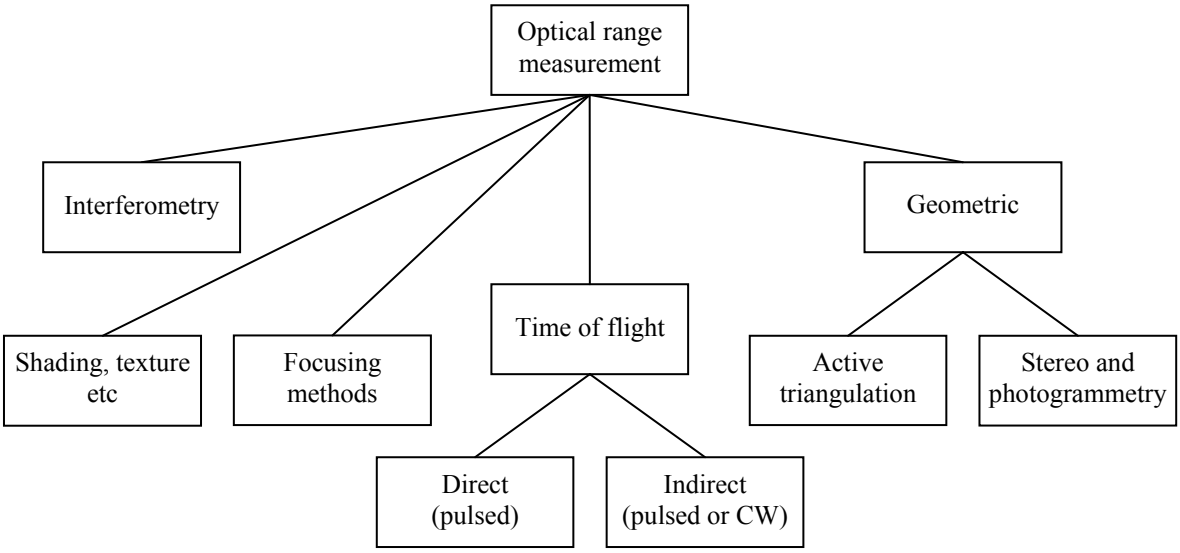

Figure 2.1: Optical range finding techniques.

2.1.1 Interferometric, Shading and Focusing Methods

Interferometric systems use a coherent waveform (light source) which is split into two paths, a reference path and a measurement path (Hariharan, 1985). At the detector, the beams are superposed, creating another waveform whose amplitude is dependent on the path-length difference between the two originating beams. A detector measures the amplitude of the resultant wave, and the length of the reference path is adjusted by a known amount to change the phase of the reference beam. From three or more measurements, the phase of the resultant waveform can be determined and the measurement path length can be calculated. The phase measurement can only be determined over a 2π radian range, resulting in an ambiguous distance measurement. The ambiguous range corresponds to half the source wavelength, which is of the order of hundreds of nanometres for visible light. Various techniques are available to extend the ambiguous range, however interferometry is predominantly suited to applications requiring highly accurate (nm) measurements over short distances (μm–cm); hence is not ideal for our applications of interest.

The shape of an object can be estimated from a number of cues within a single image, including shading, texture, intensity, etc. These techniques do not determine an absolute distance, but rather object profiles within the scene are estimated. Parameters such as the position of light sources within the scene, and the object's surface properties affect the shape determination, leading to suitability for object recognition applications, where a database of known objects (shape, colour, texture, etc.) have been predefined and can be compared to that in the recorded image (Worthington and Hancock, 2000). These techniques are not particularly suitable for true range measurement of random objects unless parameters such as the light source position can be precisely controlled (in which case it falls under the category of geometric range measurement).

Focusing methods make use of image blurring which occurs when an object is at a different distance to that of the focal plane (Gökstorp, 1994). The simplest case of a thin lens system is shown in Figure 2.2, where an object at the focal plane distance will produce a sharp image at the sensor plane, while an object at a different depth will incur blurring of the image.

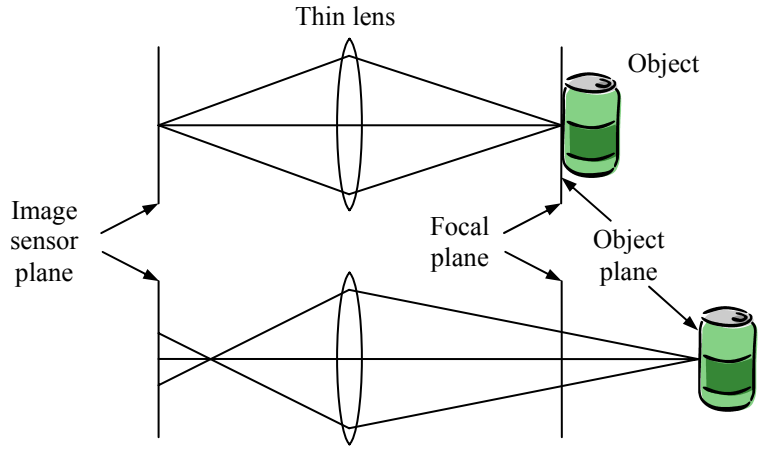

Figure 2.2: Principle of using focusing for range measurements for a simple thin lens. Top: The object lies on the focal plane, so a sharp image is produced at the image sensor plane; Bottom: Moving the object away from the focal plane produces a blurred image at the image sensor.

Accurately deconvolving the defocusing operator from within the image is the primary limitation to achieving an accurate depth measurement. Extraction of object features or object texture is required (Ens and Lawrence, 1993), constraining the types of scene that the method can be applied to. The depth measurements can be ambiguous, as image blurring occurs both within and outside the focal distance, and the level of defocusing is inversely proportional to the object distance from the focal plane, hence long range

measurements have poor accuracy. Focusing methods are therefore most suitable for operating within a limited depth range, where high precision is not required.

2.1.2 Geometric Range Measurements

Viewing a scene from multiple viewpoints gives a human depth perception, and as such is very well known and studied (Howard and Rogers, 1995). A triangle is formed between two known reference points and the unknown object position, as shown in Figure 2.3. From the known separation value x, and the measured angles α and β, the range r can be calculated.

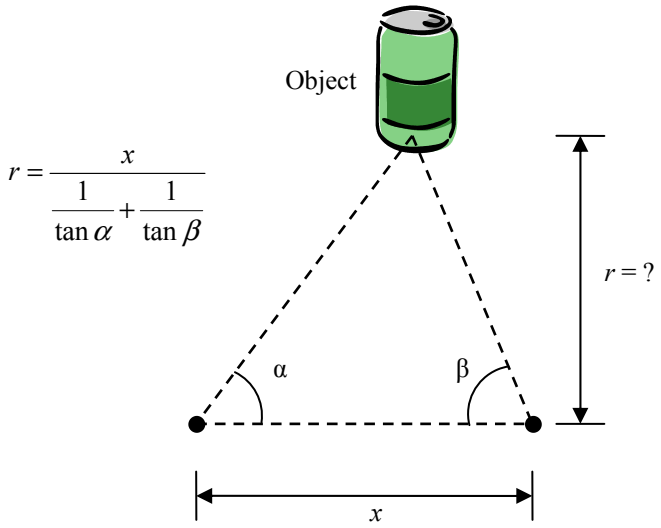

Figure 2.3: Principle of geometric range finding.

Stereo disparity based methods (photogrammetry) use this principle, capturing multiple images of a scene from different viewpoints (Atkinson, 2001). Complex processing algorithms then identify features in the images, and from the knowledge of the camera position and angle, calculate the range to each feature. The acquisition time is potentially very short, but the system speed can be limited by the intensive computational processing required to perform the 2D-correlation to identify and register feature sets between different images. This processing time is compounded if a significant number of images (required for high quality results) are used, simultaneously captured from multiple viewpoints. Scenes with rich contrast are preferred as this improves the likelihood of successful feature extraction; whereas shadowing in a scene can reduce the contrast and limit performance. An occlusion problem occurs when an object obstructs a feature from a particular viewpoint, preventing a measurement from being made. Although multiple

viewpoints can overcome the limitations of occlusion (and shadowing), the increase in computational processing is significant, implying that relatively flat scenes are preferred so that as little as two captures can be used. As shown in Figure 2.3, the range accuracy is determined from the baseline x, and the ability to accurately resolve the angles α and β. The resolution of the captured images inherently limits the resolution of the determined angles, and therefore for a particular configuration (fixed baseline) as the measured distance increases the range accuracy decreases; i.e. the measurement precision is range dependent. To improve the resolution at increased range, the baseline must be increased with two implications; 1) the system size becomes proportional to the distances being measured, and 2) the likelihood of occlusion of some features in the scene increases. The passive nature of this technique is advantageous in many applications, although in industrial applications a light source can be used to reduce shadowing.

Active triangulation uses the same geometric method as photogrammetry, but one receiver is replaced with a light source which projects a point or image onto the scene. Measurement of the angle of the reflected light, combined with knowledge of the projected light angle, allows the object range to be calculated (Webster, 1999). The active system has the ability to spatially and temporally localise the measurement by projecting the 'feature' onto the scene, reducing the so-called correspondence problem between multiple viewpoints, hence simplifying the required computation (Webster, 1999). The simplest system projects the light onto a single point, which is systematically 2D (raster) scanned over the scene by mechanical means, with the reflection detected by a 1D line sensor. The acquisition speed is severely constrained by the serial nature of this approach, largely due to the use of a galvanometer used to move scanning mirrors.

By expanding the projected point into a line, only 1D scanning over the scene is required and the acquisition time is therefore much shorter, at the added expense of requiring a 2D sensor such as a video camera. In these cases, the measurement spatial (angular) resolution is defined by the resolution of the scanning system, although precise angular scanning is susceptible to errors due to vibration.

A further increase in speed can be achieved by extending the technique to a full-field configuration through the use of structured lighting (Salvi et al., 2004) where a 2D pattern is projected onto the scene; hence no mechanical scanning is required. The complexity of

the pattern defines the number of 'features' projected on to the scene, which in turn defines the resolution of the resultant range map. The pattern is deformed by the object shape, but must remain unambiguous when reflected from the scene. This constrains the complexity of the pattern; hence also constrains the spatial resolution achievable. Coloured light can be used to improve the pattern density with less chance of incurring ambiguity, although this only offers a relatively small gain in resolution. The acquisition time is comparable to the stereo disparity method, and the computational time is comparatively small since the projected features are relatively easy to extract, making this an extremely fast method. Despite the achievable speed, the inherent limitations of the geometric method (occlusion, inverse range accuracy, and physical baseline problems) limit the usefulness of this technique, with it being most suited to short range and shape profiling applications.

2.1.3 Time of Flight Range Measurements

Time of flight measurement involves transmitting a light pulse and measuring the time delay for the pulse to travel to a target and be reflected back to a receiver (Webster, 1999), as illustrated in Figure 2.4. As the speed of light is known, the target range can be calculated using Equation (2.1), where δt is the time delay, and c is the speed of light. The division by two is due to the round trip distance travelled.

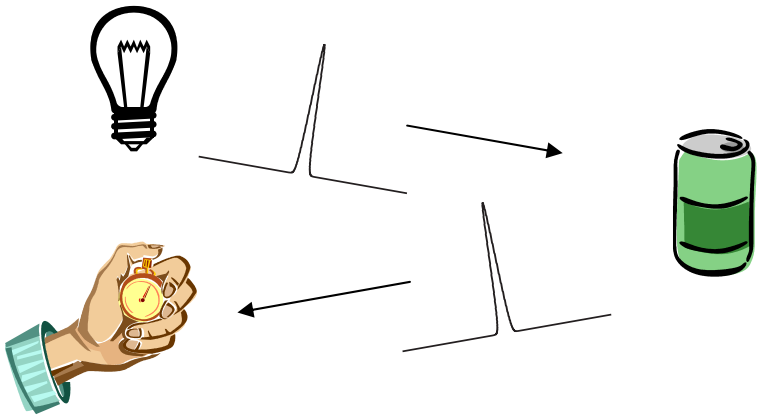

Figure 2.4: Principle of time of flight ranging. A pulse of light is emitted, it reflects from an object back to a receiver, and the time delay is measured.

$$r = \frac{\delta t\, c}{2} \tag{2.1}$$

As the emitter and detector can be placed very close together, the light path is approximately coaxial, and problems due to shadowing and occlusion are minimised.

Unlike geometric methods, measurement accuracy has no direct dependency on the target distance* (Strand, 1985; Webster, 1999), suiting the time of flight method to applications requiring long range measurements. From Equation (2.1), achieving millimetre range resolution necessitates resolving the time delay to better than 7 ps. This requires the use of a light source capable of delivering a high power pulse over a very short time, and an optical receiver with high bandwidth, sensitivity and stability. A major advantage of using the pulsed light source is that average power can remain relatively low, allowing eye safety requirements to be met, while providing a high SNR during the short pulse time.

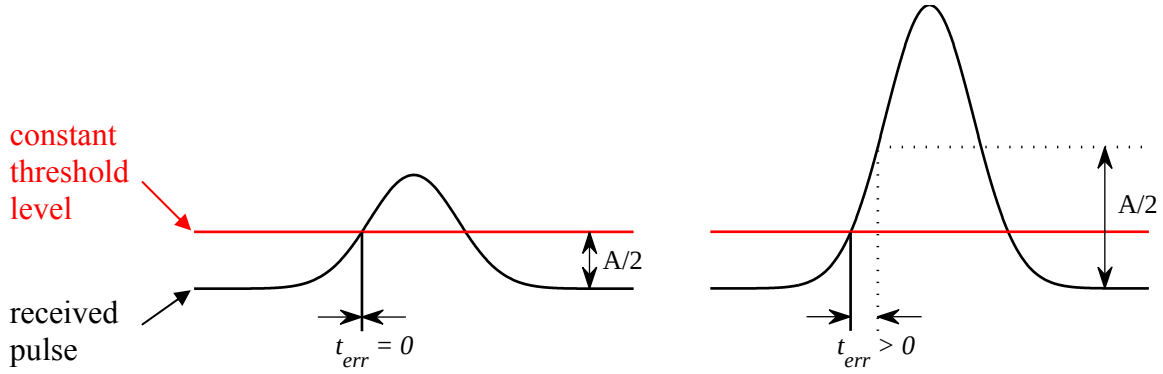

Figure 2.5: A fixed threshold level can cause a timing error due to the transition speed of the received pulse. Left: The trigger occurs correctly near half maximum amplitude (A/2). Right: The trigger occurs early causing an error in the measurement.

The optical receiver must also be connected to high speed timing electronics; if a simple counter is used it must be driven by a clock at 150 GHz to measure the time delay with the ~7 ps resolution required for millimetre precision ranging from a single light pulse. The electronic requirements are also complicated by the fact that the received pulse rise and fall times are not instantaneous; the bandwidth of the light source and receiver limit the transition speed, and temporal dispersion occurs due to atmospheric disturbances (Yang et al., 2007). If a constant threshold level is used to trigger the stop time of the counter (e.g. using a comparator), an amplitude dependent error can occur as illustrated in Figure 2.5. To overcome this error the position of the peak can be used to trigger the counter, or the

* Atmospheric absorption and scattering will reduce the signal amplitude (hence SNR) for long range measurements, reducing measurement precision, but this does not reduce the measurement accuracy. Measurement error due to refraction is ignored for this simple comparison.

amplitude can be measured to dynamically vary the threshold voltage. Such methods further complicate the required high speed electronics.

As with the geometric methods, it is desirable to work with higher dimensional light sources and sensors to generate a full field system to reduce the acquisition time. Projecting the illumination over the entire scene and using a 2D sensor array allows measurement of the entire scene to be performed with a single light pulse. The structured lighting method projects a pattern over the scene and encodes range spatially using the projected 'features' – limiting the achievable spatial resolution. In comparison, the time of flight method homogeneously illuminates the scene and encodes range temporally. The resulting spatial resolution is therefore only limited by the resolution of the 2D receiver.

The complexity of the receiver makes constructing a small area, high resolution array difficult. Ailisto et al. (Ailisto et al., 2002) used 18 optical fibres, which were bunched together in a bundle behind the receiver optics (to form an array), and split into separate fibres to connect to individual photodiodes, thus allowing the receiver electronics to remain relatively large. This method does not scale well; achieving high resolution range images would significantly increase the complexity and cost of the receiver fibre bundle and electronic receivers. The authors overcame the poor resolution by mechanically moving the receiver head, taking two measurements to achieve a resolution of 6×6 pixels. The timing electronics provide each measurement with a resolution of 9.4 cm, and adjust the threshold depending on the received signal amplitude (Ailisto et al., 2002).

Researchers at M.I.T.'s Lincoln Laboratory have built an array using avalanche photodiodes which are connected to external timing circuits (Albota et al., 2002), removing the requirement of a fibre bundle while still allowing relatively large electronic circuits to be used. A later revision provided a sensor resolution of 32×32 pixels, and bonded the photodiode array to CMOS timing electronics rather than using external counters (Aull et al., 2002; Marino and Davis, 2005). The timing resolution is limited by an effective clock speed of 2 GHz (Marino and Davis, 2005) giving a maximum theoretical single pulse range precision of ~8 cm, although the design does not compensate for variations in signal amplitude.

Niclass et al. (Niclass et al., 2005) have built an avalanche photodiode array of 32×32 pixels, however only one external time-to-digital converter is used, hence the pixel measurements must be performed serially, negating the key advantage of the array. Advanced Scientific Concepts Inc (ASC) (Stettner et al., 2004) have produced avalanche photodiode arrays with 128×128 pixels. ASC record both a digital time and analogue temporal information about the received pulse, providing precision of approximately 6 inches (15 cm) from just the digital data, or sub-inch (<2.54 cm) precision if the analogue data are also processed (Stettner et al., 2004). Raytheon have also developed avalanche photodiode arrays (Halmos et al., 2003), with latest generation devices reportedly being developed with 256×256 pixel resolution (Andressen, 2005).

The time of flight systems described in this section have an advantage that, assuming adequate SNR is achieved, a single pulse is sufficient for measurement of the entire scene and the range values can be read from the sensor with little or no post processing required. Another advantage is the ease of processing multiple signal returns, for example from a partially transparent object. The electronic speed/bandwidth limits the range resolution to tens of centimetres; to enhance the measurement resolution, the time delay of a number of pulses can be averaged together, also increasing the measurement SNR. Even for state of the art devices, the sensor array spatial resolution is typically quite low due to the complexity of combining a high speed counter within each pixel. To manufacture a large array, an avalanche photodiode (APD) or PIN photodiode array is bump bonded to a readout integrated circuit which contains the timing and peak detection circuitry (Halmos et al., 2003; Stettner et al., 2004). This complication increases the manufacturing costs; however the use of a hybrid array allows exotic materials to be used to construct the photodiode to operate with light at longer wavelengths, alleviating eye safety constraints.

2.1.4 Shuttered Light Pulse

To simplify the system requirements (primarily bandwidth constraints) a method of indirect time of flight known as range gating segmentation, or shuttered light pulse, has been utilised by a number of groups (Christie et al., 1995; Schroeder et al., 1999; Iddan and Yahav, 2001; Medina et al., 2006). In its most basic form it is used for underwater viewing (Swartz, 1994), where water absorbs and scatters light making a standard video camera ineffective at viewing the ocean floor. A pulsed light source emits a high energy

burst, sending out a "light wall" towards the ocean floor as illustrated in Figure 2.6. The light is reflected back from the seabed, carrying an imprint of the sea floor. As the light wall reaches the camera, a high speed shutter is rapidly opened then closed, allowing an image to be integrated over a short time period. Any light backscattered from particles in the water are temporally rejected, while light which is reflected from the sea floor contributes to the resultant image providing a high signal to noise ratio. The depth to the sea floor must be known in advance so that the camera shutter can be opened at the correct time to expose the sea bed, rather than the intermediate volume of water, where the required time delay is calculated by rearranging Equation (2.1).

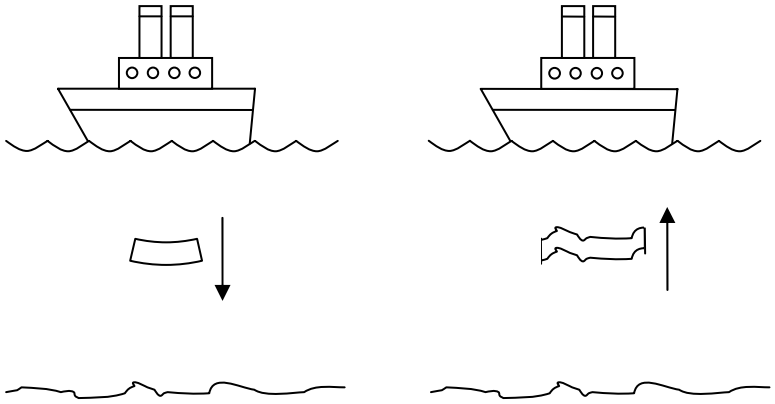

Figure 2.6: A pulsed "light wall" is reflected off the ocean floor, and a camera with a high speed shutter takes an image of the seabed.

To create a range measurement the underwater viewing technique is reversed; rather than using predetermined depth knowledge to maximise the image intensity, an image is taken and the received intensity is used to determine the depth (Busck, 2005). Figure 2.7 illustrates the basic theory of the method. In this implementation a light source emits a short pulse of light which is spread out of over the full field of view, and the high speed shutter is opened. As before, the light experiences a time delay due to the distance travel to each object in the scene, however this time the shutter is closed before all the reflected light is received. The image contains more light for close range objects making them appear bright, while distant objects appear dim; the range is therefore negatively proportional to the recoded intensity. A large number of shuttered pulses are usually integrated together to form a single frame, enhancing the SNR.

It is apparent that the intensity of an imaged object also depends on a number of other parameters (colour, texture, etc.), and that a distant highly reflective object may appear

brighter than a closer absorbent object, hence a method of correcting for these factors is required. To overcome this problem, a second frame can be acquired where the shutter time is extended, allowing all of the transmitted light pulse to be collected (Christie et al., 1995; Iddan and Yahav, 2001). The second frame is then use to normalise the measurement. Equation (2.2) is used to calculate the range, where T is the width of the illuminating pulse and I_1/I_2 is the ratio of the amplitude of the two captured frames.

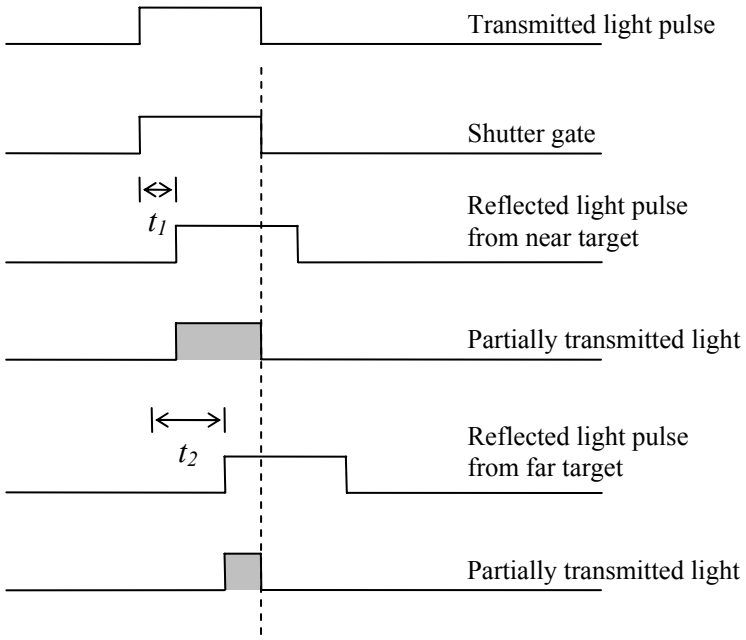

Figure 2.7: Principle of range gating. The shaded area is the amount of light integrated at the sensor; a near object appears bright, while a far object appears dark.

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

Indirectly measuring time of flight with this technique has the advantage that the electronic bandwidth requirement can be orders of magnitude below that of a direct measurement system, while maintaining similar range precision.

Two significant problems of interest are considered by Christie et al. (Christie et al., 1995), dynamic range limitations and ambient lighting. The light intensity reaching a target decreases as $1/r^2$ where r is the distance between the light source and the target. As the method described above also reduces the light proportionally with distance, the intensity rapidly decreases for distant targets. The depth measurement is therefore compromised by the limited dynamic range of a typical image sensor. Christie et al. suggest delaying the

shutter gate pulse as shown in Figure 2.8 so that the measured intensity increases with distance. This reduces the dynamic range of a given scene, allowing measurements to be performed over a larger depth variation. A variation of this scheme, where the shutter is gated both early and late, has been implemented, generating two frames capturing the near and far objects respectively (Izhal et al., 2005; Ushinaga et al., 2006). The sum of the two images then provides an 'extended gating' image used to normalise the object reflectivity when calculating the ratio I_1/I_2. The advantage of this variation is that both images can be captured simultaneously rather than sequentially by using a custom CMOS/CCD sensor.

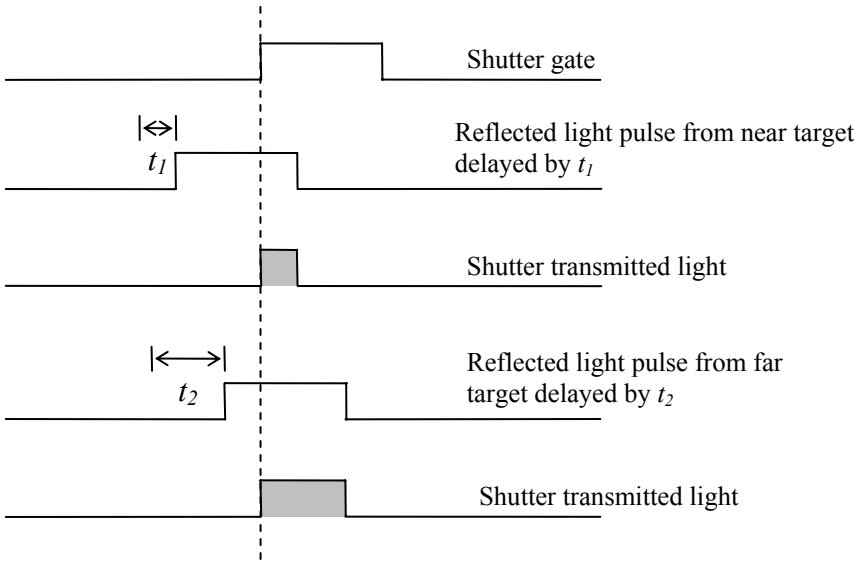

Figure 2.8: Range gating where the shutter signal is delayed. Intensity is proportional to range, enhancing distant targets and reducing sensor dynamic range requirements.

Ambient light causes a non-linear error in both of the described shuttered light pulse methods. The ratio I_1/I_2 in Equation (2.2) becomes $(I_1+A_1)/(I_2+A_2)$, where A_1 and A_2 are the contribution from ambient background lighting in each frame. Although the ratio of A_1 to A_2 can be estimated from the exposure time (assuming the ambient lighting level remains constant), the ambient lighting value cannot be determined from the two captured image frames. The use of optical interference filters and/or higher laser power pulses coupled with a reduced camera aperture are suggested, however this can only reduce the effect of ambient light and not eliminate it.

A number of groups (Schroeder et al., 1999; Gonzalez-Banos and Davis, 2004; Schrey et al., 2004; Medina et al., 2006) take an extra measurement with the light source off (a background measurement) to calculate the three unknown parameters, range, reflectance

and ambient lighting. It is suspected that the sensors in (Medina et al., 2006) and (Schrey et al., 2004) will still suffer from some effects of ambient light as the shutter pulse width varies between captured frames, hence the background level cannot easily be cancelled from the measurement. Schroeder et al. (Schroeder et al., 1999) overcome this problem by extending the shutter gate width to be larger than the light pulse width, and keep the three shutter pulses the same width. The first partially intersects the incoming reflected light, the second captures the entire reflected light pulse, and the third is captured outside the light pulse, allowing a simple subtraction to remove the effect of ambient light as shown in Figure 2.9 and Equation (2.3).

A similar method is used in (Gonzalez-Banos and Davis, 2004), where three narrow shutter periods are used to partially capture different segments of a wider light pulse. It should be noted that despite the compensation applied, if ambient light is allowed to reach the sensor the dynamic range will decrease and the measurement precision will be reduced. It is therefore still preferable to reduce the level of background light reaching the sensor, using practices such as placing an optical filter over the receiver to reject light which does not match the emitter wavelength.

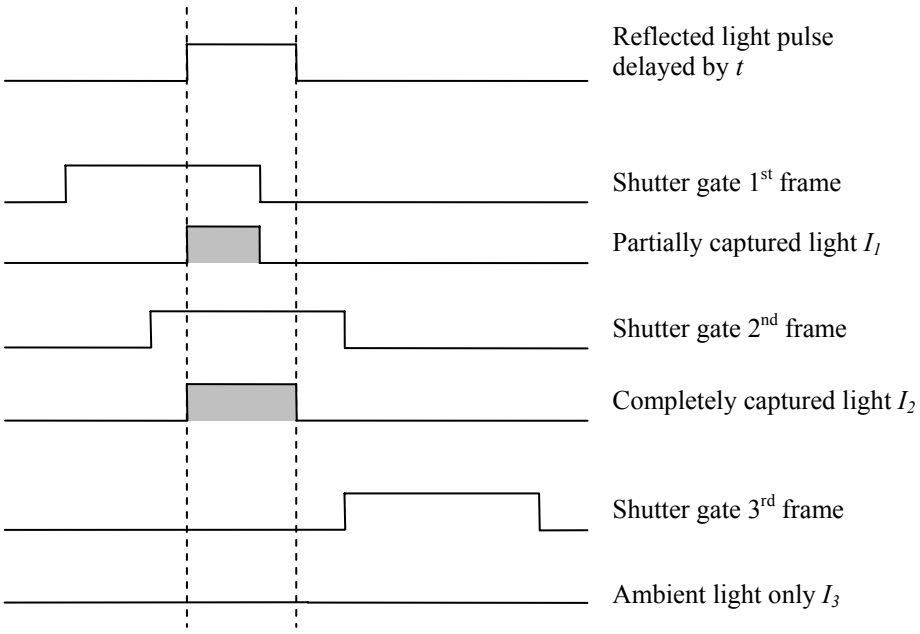

Figure 2.9: Keeping the shutter pulse width constant and capturing 3 images allows range gating while simplifying cancellation of ambient light.

$$r = \frac{cT\{1-(I_1-I_3)/(I_2-I_3)\}}{2} \tag{2.3}$$

Despite utilising an additional image frame, this method of removing the effect of ambient light is more susceptible to random noise than the previous technique. Error propagation of Equation (2.3) shows that the noise on each individual captured image contributes to the error of the calculated range; however the benefit of ambient light cancellation outweighs this disadvantage.

2.1.5 Time-Slicing (Range Gating Segmentation)

The previous shuttered light pulse methods emphasise speed, namely achieving a range map while acquiring as few frames as possible. In contrast, it is possible to use a larger number of measurements to achieve higher range precision. Busck et al. use a pulsed light source and very high speed shutter, while taking a number of images where the shutter is delayed by a fixed amount between each frame in what they call a time-slicing technique (Busck and Heiselberg, 2004). The shutter speed utilised is much faster than the previous methods where a relatively long illumination pulse was used, allowing finer 'range slices' to be captured. The light is reflected from an object and the shutter gate is initially misaligned, causing a black image as shown in Figure 2.10. A fixed delay is incrementally added to each frame, causing the light pulse and shutter to temporally align with a corresponding increase in image intensity, followed by decay as the shutter pulse again is misaligned. The resultant intensity profile follows a triangular shape, illustrated in the bottom half of Figure 2.10, where the image number of the peak is proportional to the object range. Viarani et al. use a similar configuration (Viarani et al., 2004).

The authors perform two calculations on every pixel; Equation (2.4) gives a summed intensity and Equation (2.5) gives the average two-way travel time, where I_i is the intensity for frame i and t_i is the shutter time delay for frame i (Busck and Heiselberg, 2004; Busck, 2005). Rearranging Equation (2.5) to solve for r gives the range to each object, and I can be used to give a standard 2D intensity image.

This technique reduces some of the limiting factors of the previous methods, albeit at the expense of a long acquisition time. The system performance is no longer directly constrained by the dynamic range of the sensor; rather than distance being proportional to

intensity it is found by locating the centre of the triangular peak. A weighted average is used to find the peak location, which relies on a number of image frames. This has two distinct advantages: firstly the SNR of the range measurement is not limited by the SNR of a single frame, and secondly determination of the peak location (and hence range) is no longer constrained to discrete steps by quantisation of an intensity measurement, and is much closer to a continuous measurement. The SNR will improve as a larger number of images are taken, where the N pulses contributing to the triangular peak region improve SNR proportional to $\sqrt{N}$.

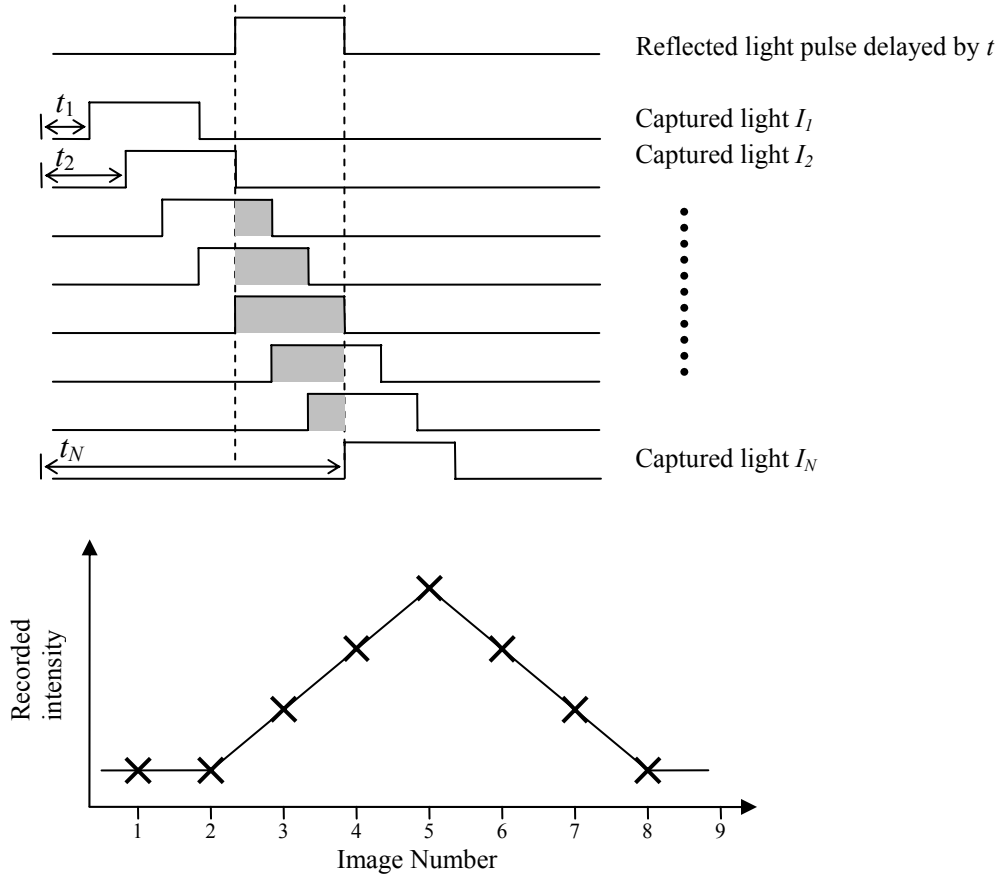

Figure 2.10: Range gating with N images. Top: The shaded area illustrates the intersection between the light pulse and the shutter; Bottom: The object range is determined from the location of the peak of the recorded intensity profile.

$$I = \sum_i I_i \tag{2.4}$$

$$\langle t \rangle = \frac{2r}{c} = I^{-1} \sum_i I_i t_i \tag{2.5}$$

The centroiding method of calculating the range, Equation (2.5), is likely used by Busck et al. as only two values need to be stored for each pixel ($\sum I_i$ and $\sum I_i t_i$), and are updated as each frame arrives. Discarding each data frame after it has been processed greatly simplifies the data acquisition hardware requirements as only minimal storage capacity is required. The weighted average method is susceptible to errors generated by a DC shift in the captured waveform, which is caused by sensor dark current and ambient lighting. Busck et al. subtract a background value from each pixel before it is processed, although it is unclear how the background value is acquired in (Busck and Heiselberg, 2004).

In (Viarani et al., 2004), a single frame at the beginning of each acquisition is exposed with the light source turned off to measure the ambient light which is then subtracted from the following image frames; however even a small deviation from the true background level will cause an error to accumulate over all the successive image frames, leading to a significant error in the final range measurement. As the waveform shape can be estimated prior to data acquisition (from the light pulse and shutter gating profiles) it would be beneficial to use nonlinear regression to fit the known curve to the experimental data, locating the triangular peak with less influence on the background intensity. This approach has been used by Andersson (Andersson, 2006) who uses least-squares parameter fitting, however the computational requirement for such a method is significantly increased over that of Equation (2.5). Andersson also recommends as an alternative to use simple thresholding, detecting the midpoint of the waveform rising edge; however errors due to laser pulse shape distortions or sensor nonlinearities are not compensated for using a simple thresholding technique.

The intensity of the reflected light will decrease by $1/r^2$ as range increases, and this intensity problem is compounded by the fact that the method does not make maximum use of the reflected light as most of the time the sensor is gated closed, leading to very low intensities for long range targets (Busck and Heiselberg, 2004). To reduce the effect of this problem, because each gated image taken resolves a different known depth window in the scene, the gain of the sensor can be increased as the range increases (Busck and Heiselberg, 2004). This helps to enhance the SNR of distant objects (hence measurement precision) by improving the system dynamic range. A corresponding algorithm would

need to be implemented to incorporate the dynamic receiver sensitivity variable in the range calculations.

The requirement of collecting a large number of images (~40 frames) leads to a long acquisition time (~1 second); hence this technique is most suitable for applications incorporating a static scene. Despite this limitation, excellent precision is attainable using this method.

2.1.6 AMCW Illumination

The previous direct and indirect time of flight systems have utilised a square pulse illumination source. Kawakita et al. use an amplitude modulated continuous wave (AMCW) with a triangular profile, while the shutter is pulsed once every period (Kawakita et al., 2000; Kawakita et al., 2004c; Kawakita et al., 2006). Two images are acquired with the light source phase shifted by 180°, refer Figure 2.11. The first image is captured during the light source ascent, where the intensity of each object decreases with range; followed by the second image captured during the light source descent where the intensity increases with range.

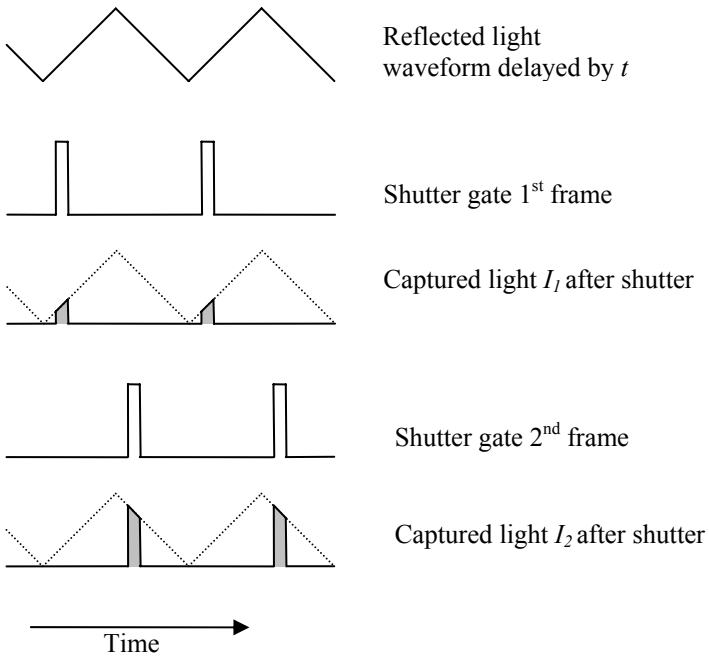

Figure 2.11: Modulation used by Kawakita et al. (Kawakita et al., 2004c; Kawakita et al., 2006). The illumination source emits a triangular modulated waveform, which is sampled twice. The intensity of the first sample decreases with range, while the second sample increases with range.

$$r = \frac{c}{2}\left[t_s - \frac{T}{2}\left(\frac{I_1 / I_2}{1 + I_1 / I_2} \right) \right] \tag{2.6}$$

The ratio of these two images is used to calculate the range, as given by Equation (2.6), where I_1/I_2 is the intensity ratio of the two captured frames. The additional term t_s is a shutter pulse delay, and T now signifies the period of the modulated light source waveform (rather than its pulse width as used earlier). As with the two frame method presented earlier, ambient light introduces a non-linear error into the range measurements. The authors use a dichroic mirror (Kawakita et al., 2000; Kawakita et al., 2006) or prism (Kawakita et al., 2004c) to separate visible and infrared light from the modulated source, and an optical filter is also used to further reduce ambient infrared light which does not match the selected source wavelength. The application presented is a depth-mapping TV camera where the spectrum of ambient light sources in a television studio can be controlled. By operating at an infrared wavelength, a fluorescent light bulb can be used for general lighting without interfering with the range measurement (Kawakita et al., 2004c). Experimental results in (Kawakita et al., 2000) verify precision degradation occurs in environments with a source containing infrared light (sunlight); "The external light leaking through the dichroic mirror shifted the operating point of the probe camera, and errors started appearing near the maximum measurable range". Due to the typically short distances (<10 m) and high powered light source (1 W average) in (Kawakita et al., 2004c), the error due to ambient light may be minimal for some applications outside the studio where high precision measurements are not required.

Test results presented in (Kawakita et al., 2004c; Kawakita et al., 2006) show a nonlinear curve between the calculated and true distance measurements, and the authors suggest this is due to the shape of the illumination modulation waveform. The limited bandwidth of LEDs (tens of MHz) restricts control over the waveform shape; the optical output is likely to approach a sinusoidal shape when modulated at high frequencies despite using triangular injection current. Although laser diodes with a much higher bandwidth (~GHz) could be used to replace the LED light source, the author retains the benefits of using LEDs (discussed further in Section 5.1.1) by performing a correction using a configurable nonlinear amplifier (Kawakita et al., 2006). The correction function is generated from an experimental calibration dataset. A FPGA is used as the signal processor to calculate the

range (Kawakita et al., 2004a) using Equation (2.6), and it would be expected that this correction could be performed within the FPGA (by means of a lookup table or similar), removing the requirement of a nonlinear amplifier.

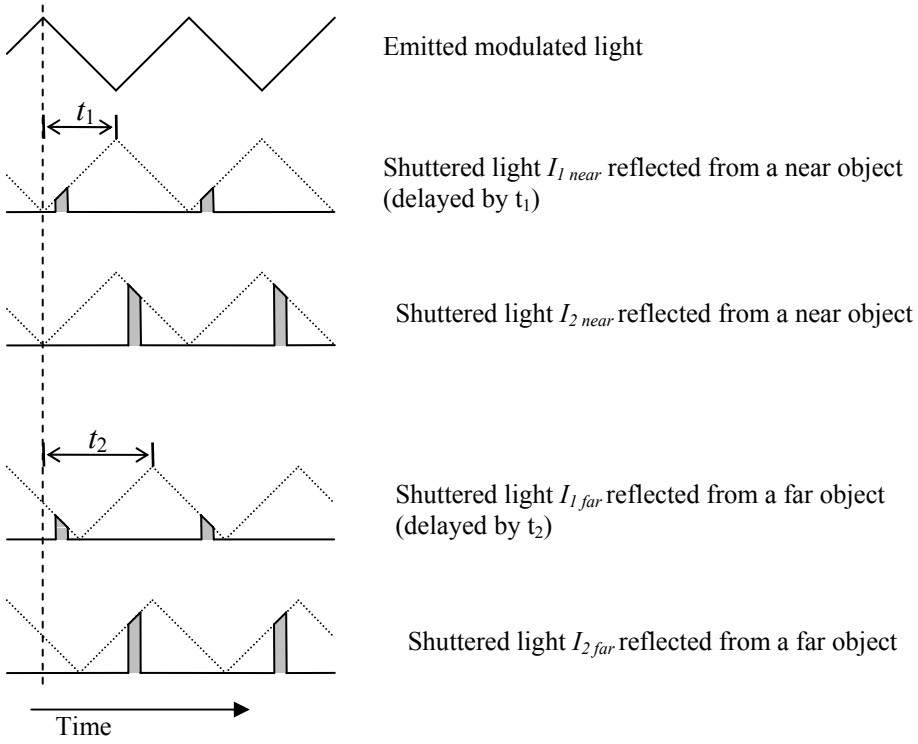

Figure 2.12: Objects at different distances are indistinguishable due to aliasing if the range exceeds $cT/4$.

An assumption is made *a priori* that the first captured frame occurs on the light source rising edge, while the second frame is captured on the falling edge. If the time of flight exceeds $T/2$ (where T is the light source modulation period) then this assumption is no longer valid and ambiguity occurs over the true distance to the object. This is illustrated in Figure 2.12 where a near object delays the light by t_1, and a far object delays the light by t_2. The shuttered light (grey area) integrated by the camera for each image is identical for the two objects despite their different ranges. This problem is avoided in the TV studio by selecting a modulation frequency low enough to encase all possible object distances (the farthermost distance typically being a rear wall), with the maximum distance possible calculated as $cT/4$, where c is the speed of light. For example, 10 MHz gives a working distance of 7.5 m. However, reducing the modulation frequency (enlarging T in Equation (2.6)) increases uncertainty in the resulting range measurement (assuming the

SNR remains constant). A trade off therefore exists between the maximum range depth and range precision.

For applications where high precision is essential then the maximum possible modulation frequency is desired, however both the LED light source and the shutter are already operated near their maximum speed. To achieve a 20 times increase in the modulation frequency, Kawakita et al. modified the method shown in Figure 2.11, to modulate the gain of the shutter while pulsing the light source (opposite to the previous method). The LED light source is replaced with a pulsed laser diode which operates more than two orders of magnitude faster than the LED array. The method of operation is the same as the previous system – taking one measurement during the ascending modulation, a second during the descending modulation, and using the ratio to calculate the distance. The experimental results show a significant increase in the range precision (Kawakita et al., 2004b), reducing the measurement uncertainty from centimetres to millimetres.

The maximum measurable range depth for this method remains equal to $cT/4$, which for the reported modulation period of 1 ns (Kawakita et al., 2004b) equates to 75 mm. As the object is unlikely to be within 75 mm of the camera, it is possible to measure the relative range of objects at a greater distance (rather than making an absolute measurement). This is useful to determine the shape of an object within a limited depth of field, for example in facial recognition applications where facial features only span a small depth, but the person may be metres from the camera. The relative depth measurements are still constrained by the condition that the first pulse must occur on the modulation rising edge, limiting the possible depth ranges from $2ncT/4$ to $(2n+1)cT/4$, where n is an integer 0,1,2… For example, the 1 ns modulation period given by the authors allows correct relative measurements in the 300–375 mm range, but not within the range 350–425 mm. This issue is not addressed as the published results were selected to fall within the correct operating range (Kawakita et al., 2004b).

The illumination modulation deviates from the triangular shape at high frequencies in the earlier system; in the new design the triangular shutter modulation also deviates from its ideal shape causing a similar non-linear range error. The experimental results published actually span an approximate distance of 450–550 mm, or 100 mm variation despite the theoretical maximum variation of 75 mm for the 1 ns modulation period. This is due to the

difficulty of generating such a short shutter modulation period, and in practice the period was slightly longer (Iizuka, 2007). The problems of range non-linearity and ambient light have been inherited from the previous design. Ambient light again affects the measurements by introducing an error, although this design is likely to intensify the problem as the shutter is constantly (partially) open compared to the low duty cycle pulses used previously.

Despite the shortcomings, the AMCW range cameras designed provide excellent speed, producing high spatial resolution (1280×720) depth images at video frame rates.

2.1.7 Phase Stepped AMCW Homodyne

An alternative to the triangular illumination waveform is to use a sinusoidal profile. A minimum of three images are taken at different phase offsets, although typically four samples are taken as shown in Figure 2.13. The round trip time for the illumination signal to be reflected from an object in the scene back to the receiver adds a phase delay to the signal modulation envelope. By determining the change in phase, the distance to the object can be calculated.

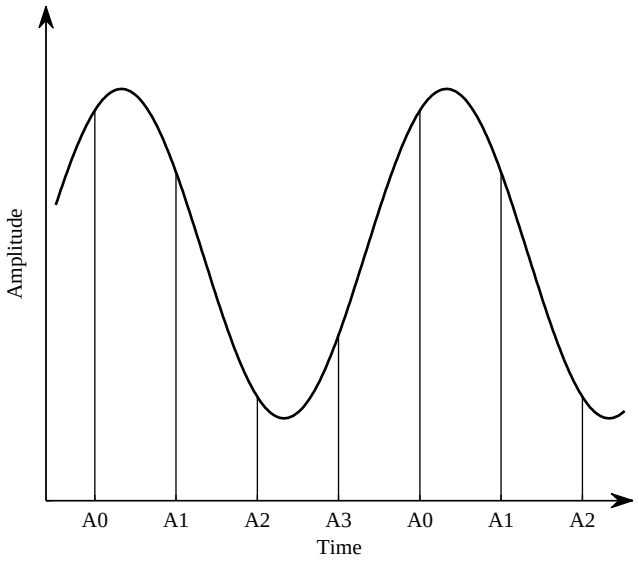

Figure 2.13: Reflected sinusoidal illumination is sampled four times to determine the modulation envelope phase.

Rather than pulsing the illumination source, or the shutter, the homodyne method modulates both the light source and the shutter with a continuous signal. For this discussion, both of the waveforms are initially assumed to be sinusoidal, hence the shutter

is capable of providing some form of gain or attenuation rather than discrete "open" or "closed" states.

The transmitted light source signal can be written as

$$P_T(t) = P_0[1 + m_T \sin(\omega_T t)] \tag{2.7}$$

where $P_T(t)$ is the transmitted power, P_0 is the average power, m_T is the transmission modulation index or modulation depth (refer Section 3.1.2), and ω_T is modulation frequency ($\omega_T = 2\pi f_T$). The light is reflected off objects in the scene, and the round trip delay adds a phase shift to the modulation envelope, φ, given by

$$\varphi = \frac{2\omega_T r}{c} \tag{2.8}$$

where r is the range to the object, and c is the speed of light. The received optical signal is then

$$P(t) = P_{av}[1 + m_T \sin(\omega_T t + \varphi)] \tag{2.9}$$

where P_{av} is the average received power and takes into account the average transmitted power, reflectivity of the object, receiver efficiency, the $1/r^2$ intensity dependence on range, and other system parameters. The reflected signal is focused onto the receiver, which is gain modulated by

$$G(t) = G_{av}[1 + m_R \sin(\omega_R t)] \tag{2.10}$$

where G_{av} is the average gain, m_R is the receiver modulation index and ω_R is the receiver modulation frequency. This process multiplies the two signals together before they reach the image sensor, resulting in

$$\begin{aligned} I(t) &= P_{av}[1 + m_T \sin(\omega_T t + \varphi)] \cdot G_{av}[1 + m_R \sin(\omega_R t)] \\ &= P_{av} G_{av} \{\, 1 + \tfrac{1}{2} m_T m_R [\cos((\omega_T - \omega_R)t + \varphi) - \cos((\omega_T + \omega_R)t + \varphi)] \\ &\qquad + m_T \sin(\omega_T t + \varphi) + m_R \sin(\omega_R t) \,\}. \end{aligned} \tag{2.11}$$

When the light source and receiver gain modulation frequencies are equal ($\omega_T = \omega_R$), also known as homodyne detection, the first cosine term becomes independent of t. The remaining cosine terms, which have a temporal dependency, are low pass filtered from the

signal due to the image sensor's long integration period (the integration time may typically be 5–6 orders of magnitude larger than the period of the high frequency terms and can therefore be ignored). The signal of interest reaching the image sensor is then

$$I = P_{av} G_{av} \left(1 + \tfrac{1}{2} m_T m_R \cos(\varphi)\right). \tag{2.12}$$

Equation (2.12) can be rewritten in the simplified form

$$I = A\cos(\varphi) + B \tag{2.13}$$

where A is related to the modulation index of the illumination source and receiver gain, B is the average signal intensity (also contributed to by background light), and φ is associated to the object range as described in Equation (2.8). The resultant intensity I is independent of time; hence a static image is produced that can be captured with a standard camera. As there are three unknown variables, at least three measurements of the scene must be made. This is achieved by adding a relative phase shift between the illumination and receiver waveforms, modifying Equation (2.13) to become

$$I_i = A\cos(\varphi + \theta_i) + B \tag{2.14}$$

where θ_i is a phase shift for frame i. Monson et al. (Monson et al., 1999) show that θ_i should be spaced over 2π radians for optimal performance. The simplest case utilises N evenly spaced samples over 2π rad as given by Equation (2.15).

$$\theta_i = \frac{(i-1)\cdot 2\pi}{N}, \quad i = 1,2,\ldots,N \tag{2.15}$$

If the samples are evenly spaced, a discrete Fourier transform can be performed and the variables B, A, and φ in Equation (2.13) can be calculated from the DC component, and the amplitude and phase of the first frequency bin. The entire Fourier spectrum does not need to be calculated to achieve this, leading to the simplification

$$A = \frac{\sqrt{[\Sigma(I_i \cos\theta_i)]^2 + [\Sigma(I_i \sin\theta_i)]^2}}{2} \tag{2.16}$$

$$B = \frac{\Sigma I_i}{N} \tag{2.17}$$

$$\varphi = \arctan\left(\frac{\Sigma(I_i \cos\theta_i)}{\Sigma(I_i \sin\theta_i)}\right) \tag{2.18}$$

where i=1,2,…,N. Despite requiring a minimum of three recorded images, four are typically used instead (Spirig et al., 1997; Lange et al., 2000; Oggier et al., 2004; Schwarte, 2004; Hsu et al., 2006). One reason for this is that the numerical processing required is greatly simplified as the cosine and sine terms can be chosen to have values from the set [−1, 0, 1], removing up to $2N$ multiplications from the required calculations. Equations (2.16) to (2.18) then simplify to (2.19) to (2.21). From the measured phase φ, Equation (2.8) can be used to determine the object range, and the value A can be used to generate a standard intensity image.

$$A = \frac{\sqrt{(I_1 - I_3)^2 + (I_2 - I_4)^2}}{2} \tag{2.19}$$

$$B = \frac{I_1 + I_2 + I_3 + I_4}{4} \tag{2.20}$$

$$\varphi = \arctan\left(\frac{I_1 - I_3}{I_2 - I_4}\right). \tag{2.21}$$

This technique offers some significant advantages over the previous triangular modulation methods used by Kawakita et al. (Kawakita et al., 2004c; Kawakita et al., 2006). Equation (2.21) allows the signal phase, hence range, to be measured independent of ambient lighting, thereby removing the associated non-linear error. Another significant advantage is that the system bandwidth requirements are further relaxed as neither the illumination source or receiver is pulsed, reducing the system complexity and cost.

The measured phase φ wraps around 2π if the range depth exceeds $cT/2$, where T is the modulation period, creating an ambiguous range measurement. The measurement precision is proportional to the modulation frequency used (discussed further in Section 3.1); hence high frequency operation is preferable. This leads to a similar problem to that with the triangular modulation method; high precision measurements can only be made over a short range depth. However, relative range measurements can be made at distances outside this range depth, with a significant advantage over the triangular modulation method; an arbitrary range window can be used. Previously, the first sample had to occur on the rising edge of the triangular waveform; hence the unambiguous range was limited to one quarter of the modulation signal wavelength. The sinusoidal modulation scheme can resolve the phase φ over the full 0–2π range; hence the

unambiguous range is doubled, and relative range measurements can be made at any absolute distance.

Most of the previous techniques encoded the range into the intensity of the received signal; in this instance range measurement precision is constrained by the dynamic range of the sensor. A CCD typically has 8–12 bit resolution, which after removing 3–4 bits for natural scene intensity variation, the remaining number of bits in which to encode the range is limited. In contrast, the homodyne method encodes the range into a continuous variable (phase), alleviating this restriction. Assuming an extremely poor quality signal where only 2 bits of information is retrieved by the CCD camera, techniques such as the shuttered light pulse, where range is encoded into intensity, can determine four unique range distances from the four possible intensity values. In comparison, from a phase encoded signal where four samples are taken, 24 unique range values can be determined. An example is shown in Figure 2.14, where a sinusoidal waveform is phase shifted and 12 unique (2 bit) values are captured by the CCD over a π rad range. A similar increase is experienced for higher precision measurements, for example 8 bit resolution will produce 256 range possibilities

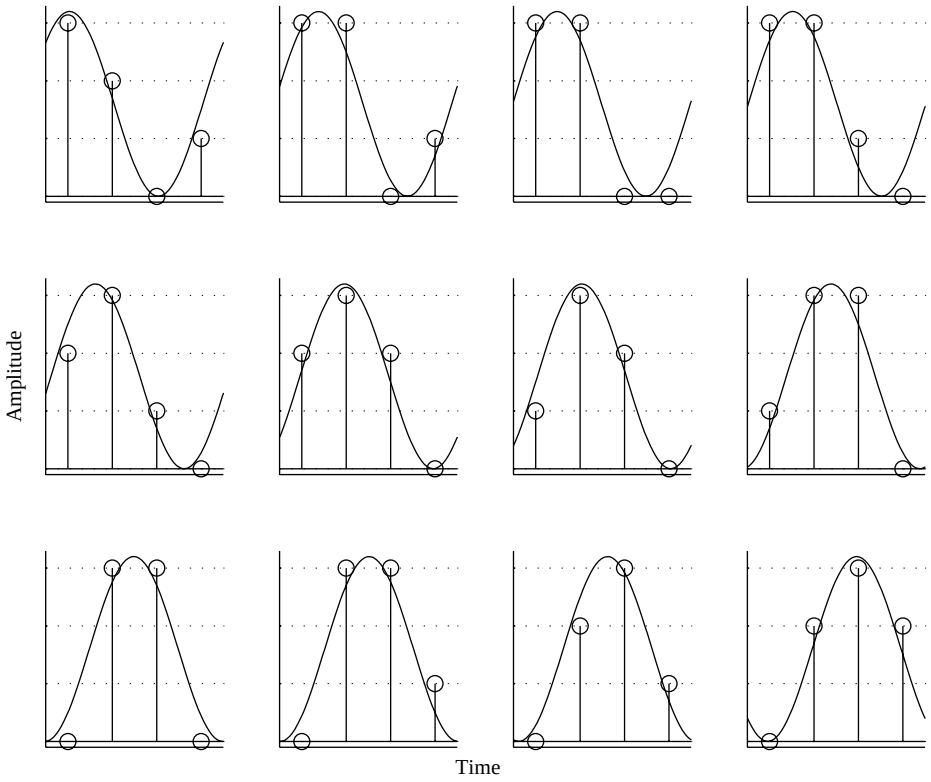

Figure 2.14: Capturing data with 2 bit precision (4 intensity levels) allows 24 unique phase angles to be determined (the first 12 are shown spanning π rad).

for intensity encoded range measurements, compared to approximately 2000 possibilities for a phase encoded technique.

The trade-off over the shuttered light pulse and triangular modulation techniques is that more samples are required; implying an increase in acquisition time. Some custom CCD/CMOS sensors have been developed which allow capture of two or four phase shifted measurements simultaneously per pixel, alleviating the speed issue at the expense of spatial resolution (Lange et al., 2000; Gokturk et al., 2004; Oggier et al., 2004; Schwarte, 2004).

2.1.8 AMCW Heterodyne as used in the Waikato System

University of Waikato has developed a method similar to the phase-stepped homodyne configuration, instead using a heterodyne configuration (Dorrington, 2004; Carnegie et al., 2006). Rather than both the illumination and receiver modulation frequencies being identical, there is a small frequency difference ω_B between them. Applying this variation to Equation (2.11), with $\omega_T=\omega_R+\omega_B$, and again ignoring the high frequency terms due to the relatively long integration time, gives

$$I(t) = P_{av} G_{av} \left(1 + \frac{m_T m_R \cos(\omega_B t + \varphi)}{2} \right) \tag{2.22}$$

which can be rewritten in the simplified form

$$I(t) = A\cos(\omega_B t + \varphi) + B\,. \tag{2.23}$$

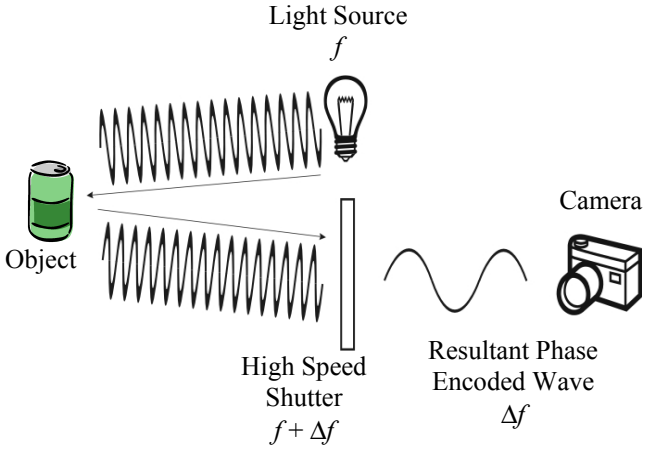

Figure 2.15: Heterodyne range imager (Payne et al., 2005).

Unlike the homodyne method, the resultant intensity *I(t)* now has a time dependent variation. The output intensity can be considered as a low frequency beat signal which has been down converted from the high frequency modulation signals. During this down conversion, both the phase and intensity relationship between the two modulation signals are preserved in the resultant beat signal. Measurement of the phase shift of the low frequency beat signal therefore allows the object range to be determined. The beat signal is captured using a digital video camera as illustrated in Figure 2.15. The maximum sampling rate of a video camera is limited, typically between 30 to 100 Hz; therefore to adequately sample the beat signal above the Nyquist rate the beat signal frequency must be in the range of single to tens of hertz. The resultant recorded video sequence appears to flash on and off, with objects at different distances flashing at different times.

The variables A, B, and φ in Equation (2.23) can be determined from a video sequence capturing at least one complete cycle of the beat signal. In this instance the modulation envelope phase is continuously changing during each sample with a difference between each image frame of

$$\Delta\theta = \frac{\omega_B}{f_s} \tag{2.24}$$

where f_s is the camera sampling rate in frames per second (fps). If an integer number of beat cycles are sampled, then the variables can be determined in a similar manner to the homodyne system

$$A = \frac{\sqrt{\left[\Sigma\left(I_i \cos(i.\Delta\theta)\right)\right]^2 + \left[\Sigma\left(I_i \sin(i.\Delta\theta)\right)\right]^2}}{2} \tag{2.25}$$

$$B = \frac{\Sigma I_i}{N} \tag{2.26}$$

$$\varphi = \arctan\left(\frac{\Sigma\left(I_i \cos(i.\Delta\theta)\right)}{\Sigma\left(I_i \sin(i.\Delta\theta)\right)}\right). \tag{2.27}$$

As the object range is again encoded as a continuous variable (phase), the measurement accuracy is not directly constrained by the dynamic range of the camera. The beneficial characteristic of the heterodyne system over an equivalent phase stepped homodyne system is that the number of samples per beat cycle can easily be modified by altering the frequency difference between the two modulation signals. Taking a larger number of

samples per beat cycle allows higher resolution determination of the phase, as illustrated in Figure 2.16. As before in Figure 2.14, limiting the signal amplitude to 2 bits, 96 unique phase angles can be determined from data sampled 16 times per beat period, compared with the earlier 24 phase angles from four samples.

For this simplistic noise free model where quantisation limits the system performance, taking 16 samples separated by $\pi/8$ rad (16 samples over a single beat period) provides improved phase measurement resolution compared to taking 16 samples separated by $\pi/2$ rad (four samples per beat period, repeated over four periods). This is an important aspect for low intensity signals such as the 2 bit example shown where quantisation limits measurement resolution; hence it is preferable to capture a single beat period with fine phase resolution. In a real system with random noise the captured value will vary between each beat cycle period and quantisation error can be reduced by averaging the measurements together. It is still preferable however to capture a single beat cycle as discussed in Section 3.2.2.

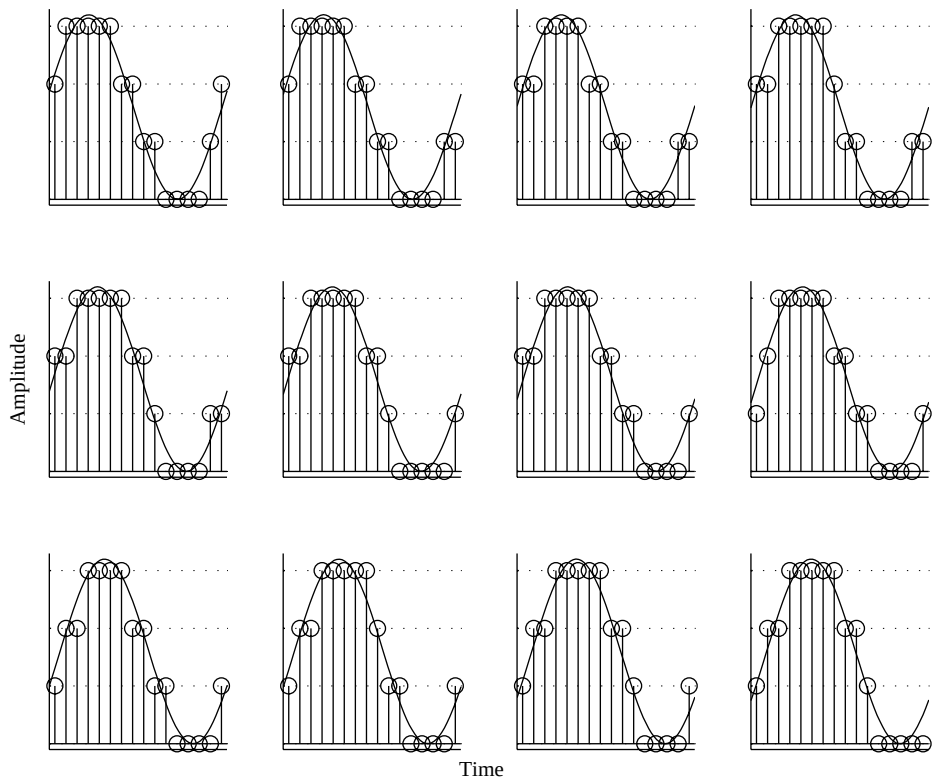

Figure 2.16: Capture of sixteen samples per beat period allows 96 unique phase angles to be determined from 2 bit data (the first 12 are shown spanning $\pi/4$).

Harmonics in the recorded waveform can also contribute an error in the measured range (for both the phase-stepped homodyne and heterodyne techniques) when taking a small number of samples per beat period, hence it is again preferable to take a large number of samples over a single beat period where possible. This is discussed in more detail in Section 3.2.2. However, since N is also proportional to the measurement acquisition time, a trade-off between speed and precision exists. The flexibility of the heterodyne configuration allows the resulting range measurement accuracy and precision, as well as acquisition time, to be controlled to meet individual application requirements by simply changing the heterodyne beat frequency. Hence dynamic scenes (such as a moving vehicle) can operate with as few as three captured image frames, while a static scene can capture hundreds (or thousands) of image frames to provide extremely high precision.

A prototype system, the 'Waikato Ranger' has been constructed at The University of Waikato using the AMCW heterodyne technique. The system is constructed from off-the-shelf components, and has demonstrated centimetre level range precision (Carnegie et al., 2005). It is anticipated that the system is capable of yielding measurements with millimetre precision through a number of improvements to the components and electronic signals used, exceeding the capability of other full field imaging devices available.

The prototype system is constructed using two Agilent 8648B bench top synthesised signal generators which produce the illumination and shutter modulation signals with 1 Hz tuning. By interconnecting the two devices, a single reference clock is used by both signal generators, ensuring that the two modulation signals are frequency locked. One signal is connected to an array of Agilent HLMP-EL series LEDs, which are modulated at 10 MHz and flood illuminate the scene. An image intensifier is used as the receiver and is modulated by the second signal, which has a small frequency difference compared to the LEDs (typically the beat frequency f_B is 1, 2, or 5 Hz). The resultant beat signal, or 'flashing image', is digitised by an 8-bit Pulnix TM9701 CCD camera and stored on a PC.

A zero padded FFT is performed on the recorded time series separately for each pixel, with the peak magnitude giving the beat frequency, and the corresponding phase angle measurement giving φ. The range for the pixel is then determined by rearranging Equation (2.8). The zero padded FFT is used to interpolate between samples to precisely measure the beat frequency as the camera frame rate and beat frequency are generated

independently (Carnegie et al., 2005); there is no mechanism available in the system to ensure that an integer number of samples are taken over the beat signal period, hence the beat signal may not fall precisely on a FFT frequency bin.

A second limitation of the independent camera frame rate is that the captured frames are not synchronised to the start of the heterodyne beat signal, resulting in an arbitrary offset in the resulting phase, hence range, measurement. To temporarily overcome this problem, a reference object is placed in the scene and measurements are made relative to the known object range. A preferable, permanent solution is to measure the phase difference between the two high frequency modulation signals, creating an electronic signal as a reference to determine absolute phase, negating the requirement of a known object in the scene.

2.1.9 FMCW / FSCW

Frequency modulated continuous wave (FMCW) and frequency stepped continuous wave (FSCW) techniques are often used in radar and sonar applications. The objective is to encode an object's range into the frequency of the received signal, rather than amplitude or phase.

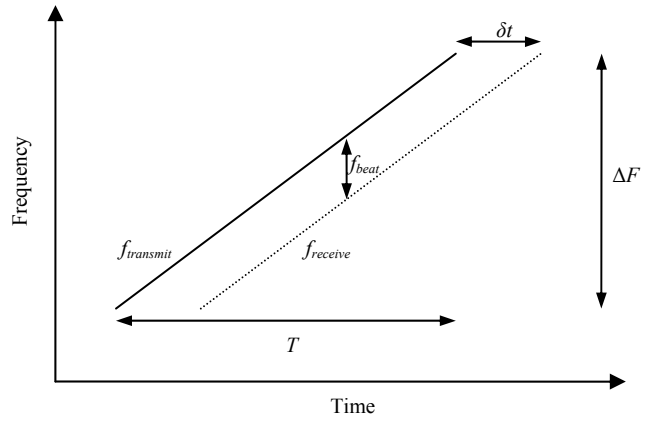

Figure 2.17: Frequency modulated continuous wave (FMCW) theory.

To construct a full-field imaging system, a configuration similar to that of the AMCW heterodyne system shown in Figure 2.15 can be used, where the illumination source and receiver gain are sinusoidally modulated (Stann et al., 1999). The modulation frequency is linearly increased over period T, from a low frequency to a high frequency, generating a chirp over a frequency range ΔF as illustrated in Figure 2.17. Technically this is not true frequency modulation (as the frequency/wavelength of the optical source is not changing), and the term 'chirped AMCW' is a more accurate description; however the term FMCW is

well understood in radar applications and the underlying operation remains the same in either case.

The received modulated signal (dashed line in Figure 2.17) experiences a time delay δt due to the round trip distance. Gain modulation within the receiver mixes the original and delayed signals, producing a low frequency beat signal f_{beat} which has a frequency proportional to the object range as given by

$$r = \frac{c f_{beat} T}{2\Delta F}. \tag{2.28}$$

The low frequency signal is sampled, and a FFT performed to determine the range to each object. The FSCW method employs the same technique, although the continuous frequency chirp is replaced by a frequency step ΔF_s between each sample, and the corresponding $T/\Delta F$ term in Equation (2.28) is replaced with $1/\Delta F_s$ (Gulden et al., 2002). An advantage of using a chirped frequency is the ability to detect and measure multiple objects within a single pixel. If an object is partially transparent, a signal will be measured from both the transparent and subsequent objects, and assuming the FFT is adequately sampled, the two beat frequencies can be resolved to simultaneously give the distance of both objects. The inherent resolution (minimum object separation), ΔR, is given by Equation (2.29) for FMCW operation (Stann et al., 1999), illustrating that a large frequency range, ΔF, is desirable.

$$\Delta R = \frac{c}{2\Delta F} \tag{2.29}$$

As the beat frequency varies with object range, care must be taken to ensure the Nyquist sampling criteria is met at all times when using the limited frame rate of a typical CCD camera, i.e. for a CCD frame rate of 30 fps the beat frequency must not exceed 15 Hz. This constrains the maximum measurable scene depth for a chosen ΔF and T.

There are a number of limitations which make the FMCW method less desirable than the previous heterodyne method for use in a full-field imaging system:

1) To attain high measurement resolution, the modulation frequency chirp must cover a wide frequency range. A high upper frequency is therefore required, necessitating an electronic bandwidth much higher than that typically required for the heterodyne method.

2) The limited CCD frame rate implies that a long chirp period is required, resulting in an acquisition time which is significantly longer than the heterodyne system (which is capable of operating with only three captured frames).

3) If multiple object measurement within a single pixel is not required (as is the case for most applications with opaque objects), beat frequency determination can be improved by interpolating between the FFT frequency bins. Interpolation can be performed by fitting a polynomial to the FFT data (Stann et al., 1999), or by simply padding the FFT, with the peak location determining the object range with high resolution. Although the measurement resolution increases significantly (potentially by orders of magnitude), the unique multi-object determination advantage of using the FMCW technique is lost. Accurate determination of the beat frequency is also reliant on the collection of an adequate number of samples during the chirp period.

4) Maintaining a constant modulation index as the high frequency signals are swept over the entire wideband frequency range can be difficult due to any impedance mismatch between the electronic driver and corresponding components. This causes the amplitude of the resulting beat signal to vary during the chirp period, making it difficult to accurately determine the beat frequency (Stann et al., 1999). The modulation frequency used by the heterodyne method may fall within the same frequency range spanned using FMCW, but as the frequency used is static, the modulation index remains constant during the acquisition period and the heterodyne method is therefore less susceptible to impedance matching problems.

5) If the illumination and shutter modulation waveforms are not sinusoidal, an effect which is discussed further in Section 3.2.2 for the heterodyne method, then the beat signal produced will contain harmonics. A single object will therefore contribute multiple peaks in the spectral analysis, which would appear as separate objects (spaced at regular distances). This problem could easily be identified and corrected for an opaque object; however for the instance where multiple objects occur within a single pixel, harmonic

frequencies from one object could fall on, or near, the frequency of a second object, contaminating the measurement. The FMCW technique is therefore suited to systems with linear electronic components which are capable of producing high spectral purity.

2.2 Demonstrated Systems

A number of full-field ranging systems have been developed, both commercially and within laboratory environments, utilising the various techniques described. A selection of direct and indirect full-field time of flight systems are included in Table 2.1, and various parameters of interest are given. A prudent comparison should be made between performance values listed as the test conditions are typically unknown (including factors such as object surface scattering properties and ambient lighting levels). Parameters such as the illumination intensity and acquisition time also affect each system's performance (hence are stated where known), and in most cases can easily be adapted allowing a given system to be tailored towards meeting the needs of a specific application.

Some considerations of interest which can be taken from Table 2.1 are:

- Direct time of flight systems have limited depth resolution, but are suitable for long distance (km) applications.
- In contrast, the indirect time of flight measurements are capable of significantly higher depth resolution for short range applications.
- The field of view (FOV) is determined by the receiver technology density, with an image intensifier providing the highest spatial resolution. Image intensifiers are discussed in further detail in Section 4.3.
- The lower bandwidth requirement for indirect time of flight techniques (AMCW) allow low cost LEDs to be used as the illumination source in a number of cases.
- Acquisition time is an important factor in achieving high range resolution measurements.

Table 2.1: Selection of full-field time of flight range imaging systems.

Author/ company[a]	Ranging method	FOV (pixels)	Range resolution (cm)	Range (m)	Illumination source	Receiver technology	Max update rate
MIT (Marino and Davis, 2005)	Direct TOF	32×32	40	150	Laser, 300 ps, 16 kHz	Avalanche photodiode	16 kHz
ASC (Stettner et al., 2004; Advanced Scientific Concepts, 2006)	Direct TOF	128×128	15	3–1500	Laser pulse	Avalanche photodiode	Single shot to 30 Hz
Christie et al. (Christie et al., 1995)	Shuttered light pulse	500×500	2	0.8–1.8	Laser pulse, 40 mW	Image intensifier	Not specified
3DV Systems (Iddan and Yahav, 2001; 3DV Systems, 2003)	Shuttered light pulse	Video camera dependent	1 0.5	0.5–7 1	Laser pulse	Image intensifier	60 Hz
Busck et al. (Busck and Heiselberg, 2004)	Range gating segmentation	582×752	0.1	1.5	Laser pulse, 200 ps, 140 mW avg.	Image intensifier	1 Hz
Kawakita et al. (Kawakita et al., 2004c)	AMCW	1280×720	1.7	2 7.5 max	LED, 1 W, 10–50 MHz	Image intensifier	29.97 Hz
Kawakita et al. (Kawakita et al., 2004b)	AMCW	768×493	0.24	0.075	Laser pulse, 68 ps, 9.7 mW	Image intensifier	15 Hz
Sandia National Lab. (Nellums et al., 2006)	Phase Stepped Homodyne	720×466	0.25	1.5–4.5	Laser, 12 W, 140 MHz	Image intensifier	7.5 Hz
Swissranger (Büttgen et al., 2006)	Phase Stepped Homodyne	176×144	1.5	0.5–2 7.5 max	LED, 1 W, 6.6–30 MHz	CCD/ CMOS	30 Hz
Canesta (Hsu et al., 2006)	Phase Stepped Homodyne	160×120	0.5	1 3.4 max	LED, 2.5 W[b], 44 MHz	CMOS	7.5 Hz
PMD Technologies (Schwarte, 2004) (PMD Technologies, 2006)	Phase Stepped Homodyne	120×160	>0.6	7.5	LED, 3 W, 20 MHz	CMOS	15 Hz
Stann et al. (Stann et al., 1999)	FMCW	256×256	375	1000	Laser, 2 W, 100–500 MHz	Image intensifier	0.25 Hz[c]
University of Waikato (Carnegie et al., 2005)	Heterodyne	768×484	1.2	5 7.5 max	LED, 10 MHz	Image intensifier	0.1 Hz

[a]MIT = Massachusetts Institute of Technology, ASC = Advanced Scientific Concepts Inc.
[b]Power specified as 5 W peak.
[c]Assumes 128 frames captured at 30 fps.

2.3 Fluorescence Lifetime Imaging

Although not a range measurement device, frequency-domain fluorescence lifetime imaging (FLIM) is a tool used in biomedical applications with hardware and computation that draws a very close resemblance to that of the homodyne/heterodyne range finder system (Lakowicz and Berndt, 1991).

A short wavelength source illuminates the sample, where fluorophores absorb the excitation photons, and in turn emit photons at a longer wavelength. The dynamics of the emission are used to fingerprint the molecule, and can often be characterised by the temporal (exponential) decay, or lifetime. By sinusoidally modulating the excitation source, using the "frequency domain method", the sample fluorescence will also be sinusoidally modulated with a phase delay proportional to the fluorescent lifetime. The phase shift can be measured with techniques identical to the homodyne/heterodyne methods discussed, allowing the molecules present in the sample to be identified. The advantage of this technique is that the molecules are identified by their physical properties; hence the measurement is largely independent of sample concentration and the illumination intensity.

The main differences between FLIM and range imaging systems are the optical configuration, the illumination wavelength, and the illumination power. The imaging area (and range) is typically relatively small for FLIM, with applications in optical microscopy (Becker and Bergman, 1993) and fingerprint enhancement (Dinish et al., 2005). In comparison, for range imaging applications a standard photographic lens is used to image a wide field of view and the target range may be tens or hundreds of metres. The excitation illumination wavelength must be shorter than the emission wavelength for fluorescence to occur, typically requiring use of UV or blue illumination sources. In contrast, range imaging systems often use NIR wavelengths which are invisible to the human eye, allowing non-disruptive operation within human environments. It also allows simultaneous operation of a standard colour CCD camera within the scene without interference from the range imaging systems, permitting data fusion between the two sensor types to produce colour depth images. High illumination intensity leads to rapid photobleaching of fluorescent samples, destroying the fluorophores being measured by the FLIM system. The illumination power should therefore be minimised where possible to

avoid this effect. In ranging applications, the illumination power should be maximised to increase the SNR, although it is constrained by eye safety requirements in some applications.

Despite these differences, enhancements in technology or processing techniques developed for FLIM can often be applied to the range imaging application and vice versa.

2.4 Summary

This chapter discussed the working principle of a number of full-field time of flight range finding systems, and the associated benefits and limitations of each method. The heterodyne method is of particular interest as it minimises bandwidth requirements (by not requiring pulsed lasers or extremely fast detectors and circuitry); hence the design complexity and cost are reduced compared with direct time of flight systems. The object range is encoded into the phase of the received modulation envelope, which avoids dynamic range limitations experienced by alternative systems where the object range is encoded directly into the image intensity. The heterodyne method can be configured to operate with high speed or high precision by simply changing the beat frequency: altering the number of samples captured over the beat period. This flexibility allows it to be adapted to meet the requirements of a wide range of applications.

A prototype range imaging system has been demonstrated using the heterodyne method and has revealed centimetre measurement precision, comparable to results from other full-field range finding systems. It is anticipated that this technique is capable of generating measurements with millimetre resolution through a number of enhancements to the existing system.

3. Heterodyne Measurement Accuracy and Precision

Parameters influencing the heterodyne range imaging system are investigated in this chapter. By understanding the characteristics of the beat signal phase measurement utilised in the heterodyne system, the hardware and processing software can be designed to optimise the measurement performance, and to anticipate and avoid the occurrence of systematic errors.

Two terms are used to quantify system performance: measurement precision and accuracy. Precision is is the reproducibility of the system, i.e. the degree of the similarity between multiple measured values. It is affected by random noise that produces a random error in the measurement result, and is quantified using the standard deviation of the measurements. In contrast, accuracy is the measure between the true value and the reported value, where measurements can be affected by a bias or systematic (repeatable) error.

3.1 Heterodyne Measurement Precision

The RMS range error σ_R of an amplitude modulated laser range finder system is given by (Jelalian, 1992)

$$\sigma_R = \frac{c}{4\pi f m \sqrt{SNR}}. \tag{3.1}$$

Applying this formula to the heterodyne configuration used by the Waikato Ranger, it is apparent that the measurement precision can be improved by increasing the modulation frequency f, modulation index m, signal to noise ratio SNR, or a combination of all three.

3.1.1 Frequency

Increasing the modulation frequency is conceptually straightforward. In practice the system components have limited bandwidth, hence constraining the maximum operating frequency achievable. Replacing the LEDs used with laser diodes (refer Section 5.1), and redesigning the electronic signal generation and driver stages allows an increase in modulation frequency. Increasing the modulation frequency also reduces the unambiguous measurement range, necessitating the need for multiple operating frequencies (refer Section 3.2.1).

3.1.2 Modulation Index

The modulation index (or modulation depth) is used to describe the ratio of the maximum to minimum levels of a modulated signal. In the heterodyne system, assuming a sinusoidal waveform, it can be measured as

$$m = \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}} \tag{3.2}$$

where I is the signal intensity. Figure 3.1 illustrates a sine wave with modulation depth values of 0.2 and 0.8 respectively, where the average power remains constant in both cases.

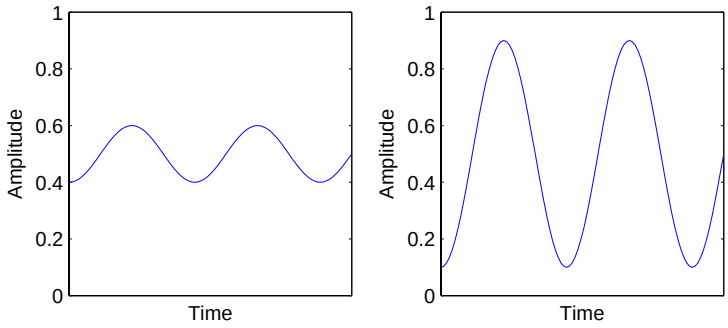

Figure 3.1: The modulation depth describes the ratio of the maximum to minimum levels. Left: a modulation depth of 0.2; Right: a modulation depth of 0.8.

From Figure 3.1 it is immediately obvious that in the presence of noise a large modulation index value is desirable to ensure minimal disturbance in phase determination. It can be seen from Equation (2.22) that the heterodyne beat signal modulation index is determined

by $m_T m_R/2$, which is the product of the high frequency illumination and receiver gain modulation indices. Hence maximum sinusoidal modulation of both the illumination and receiver gain ($m_T = m_R = 1$) produces a beat signal modulation index value of 0.5.

Analogous to the difference between integrating-bucket and phase-stepping interferometry (Sirohi and Chau, 1999), the heterodyne configuration reduces the recorded modulation index due to the CCD integration time. The beat signal amplitude is constantly changing with time; hence the recorded intensity can be considered an average of the beat signal amplitude over the integration period. The signal from the image intensifier was given in Equation (2.23) as

$$I(t) = A\cos(\omega_B t + \varphi) + B, \tag{3.3}$$

which can be rewritten as

$$I(t) = A\cos(\theta(t) + \varphi) + B \tag{3.4}$$

where $\theta(t)=\omega_B t$ is the instantaneous beat signal phase at zero distance. The average intensity for a single captured frame over the camera integration period can be written as

$$\begin{aligned} I_i &= \frac{1}{\Delta\theta}\int_{\theta_i-\frac{\Delta\theta}{2}}^{\theta_i+\frac{\Delta\theta}{2}} \left[A\cos(\theta_i(t)+\varphi)+B\right] d\theta(t) \\ &= \operatorname{sinc}\left(\frac{\Delta\theta}{2}\right) A\cos(\theta_i+\varphi)+B \end{aligned} \tag{3.5}$$

where $\Delta\theta=\omega_B/f_s$ (as given in Equation (2.24)) is the change in beat signal phase through the CCD integration period, and θ_i is the average beat signal phase for the ith measurement. To maximise the recorded modulation index, the camera integration period and/or signal beat frequency should be made as small possible to minimise $\Delta\theta$. Reducing the camera integration period typically corresponds to an increased frame rate; however a video camera is generally limited by the sensor readout electronics to approximately 30–100 fps. The integration period could also be reduced by introducing a dead time between captured frames (maintaining a low frame rate); however this will reduce the total light collected, thus reduce the SNR.

Given a camera frame rate of 30 fps, Table 3.1 illustrates the reduction in the modulation index value, hence measurement precision, for various beat frequencies (it is assumed that the integration period is the inverse of the frame rate, i.e. the sensor readout and reset time for each frame is instantaneous). Clearly it is preferable to use a low frequency beat signal to enhance the resultant modulation index value.

Table 3.1: Reduction in modulation index for an AMCW heterodyne ranging system for various selected beat frequencies using a camera operating at 30 fps.

Beat frequency (Hz)	Camera frame rate (fps)	Samples per beat signal period	Reduction in modulation index due to the camera integration period
1	30	30	0.18%
2	30	15	0.73%
3	30	10	1.6%
4	30	7.5	2.9%
5	30	6	4.5%
7.5	30	4	10%
10	30	3	17%

For configurations where ten or more samples are collected during each beat signal period, reduction of the modulation index is negligible, whereas for three or four samples significant reduction of the beat signal modulation index occurs. The number of samples collected during one beat signal period also typically constrains the minimum acquisition time (when the camera is operating at its maximum frame rate). This is an important consequence to consider when using the heterodyne configuration; the equivalent homodyne method does not reduce the modulation index value (as the signal is not changing during the integration period), hence is more suited to applications where a short acquisition time (three or four captured frames) is favoured over high measurement precision (resulting from tens to hundreds of captured frames). Reduction of the modulation index is discussed further in Section 3.2.2, where it is shown to be beneficial under certain circumstances to reduce the amplitude of signal harmonics.

3.1.3 Signal to Noise Ratio

The signal to noise ratio value in Equation (3.1) is comprised of a large number of factors, and is used to encompass the entire system. Some of the factors include:

- Received illumination intensity, dependent on
    - Emitter intensity.
    - Scene colour and surface scattering properties.
    - Distance.

- Receiver efficiency.
    - Quality of optics.
    - Optical to electrical efficiency.
- Receiver noise.
    - Level of ambient light.
    - Sensor dark current and readout noise.
    - Thermal noise.
    - Quantisation error.
    - Shot noise in the MCP.
- Acquisition time.

Some scene dependent parameters can be controlled in certain applications (e.g. controlling ambient lighting or limiting the maximum range to be measured), but in general a wide variation in target surface scattering properties may be encountered in a given scene. As these parameters cannot be controlled, the measurement uncertainty will be largely scene (application) dependent.

The remaining parameters however are dependent on the components and configuration of the system, and therefore an effort can be made to optimise the operation to enhance the SNR. Upgrading the 8 bit camera with a lower noise 12 bit model, or increasing the illumination power are two possible enhancements which could be used to improve the system SNR.

3.2 Heterodyne Measurement Accuracy

The range measurement accuracy is dependent on a number of factors that can lead to a systematic distortion of the measurements. Geometric distortions can be caused by the assumption that the projected and received illumination is a coincident point source; whereas in practice the illumination source is mounted beside (or around) the receiver, and is often *not* a point source. At long distances the illumination source approximates the ideal coaxial point source, but for short distances consequential measurement distortions can occur. Lens aberration can cause errors in the *x*-*y* location of measurements, and the measured range near the outer circumference can be elongated if the raw measurement values are not corrected for the sensor perspective projection. Through calibration and

system modelling, post processing of the captured data can be used to negate the effect of these disturbances.

Non-geometric factors also influence the measurement accuracy, and are related to measurement of the beat signal phase. There are two aspects to consider; the first is due to the phase measurement returning a cyclic value, and the second is due to systematic contamination of the phase measurement process.

3.2.1 Range Measurement Aliasing

Due to the use of a continuously modulated waveform an ambiguity occurs in the measured range. The object distance can be determined using Equation (3.6) (rearranged from Equation (2.8)) using the phase shift φ measured using Equation (2.21) and a modulation frequency f.

$$r = \frac{c\varphi}{4\pi f} \tag{3.6}$$

Although the arc-tangent function in Equation (2.21) only returns a value in the range $[-\pi/2, \pi/2]$, if the sign of the numerator and denominator of Equation (2.21) are taken into account, the phase angle can be determined over the range of $[-\pi, \pi]$. This corresponds to a range of $cT/2$ where T is the modulation period ($T = 1/f$). In a scene where object distances exceed this unambiguous range, the measured phase shift will wrap around 2π and the measured distance will be incorrectly offset by $ncT/2$ where n is an integer. Figure 3.2 shows an example of this occurring, where a number of object locations are indistinguishable due to the recovered phase angle of the beat signal being repeated at these distances. As an example, for a modulation frequency f =10 MHz ($cT/2$ = 15 m), the system is unable to uniquely determine if an object is at 2 m or 17 m (or 32, 47, 62,… metres).

One approach to overcome this problem is to select a modulation frequency which is low enough to encompass all possible object distances for a given scene. As the measurement uncertainty is proportional to the modulation frequency (as given by Equation (3.1)), this approach severely constrains the system performance. A preferable method is to use two modulation frequencies, with the simplest arrangement using a high frequency to achieve

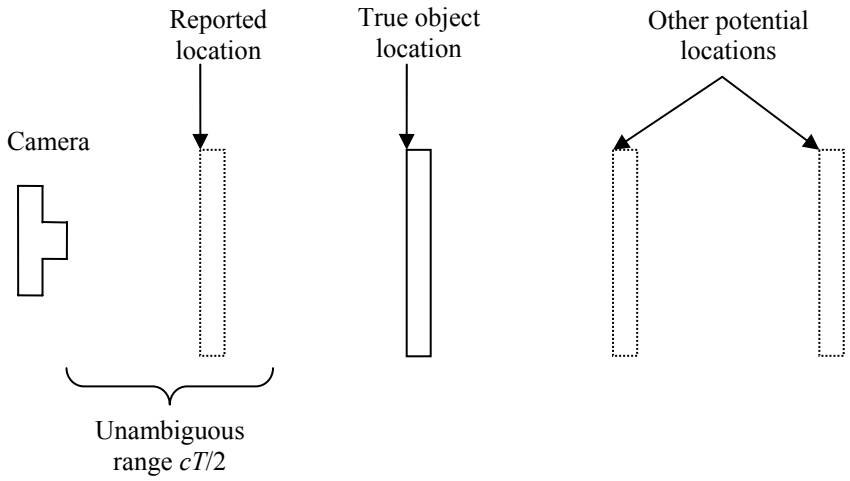

Figure 3.2: Continuous modulation leads to aliasing in the measurements. Objects that fall outside the unambiguous range are incorrectly reported to be closer to the imaging system.

high precision, and a low frequency to remove the ambiguity. A typical pair of frequencies may be f_1=3 and f_2=100 MHz, achieving unambiguous measurements over a 50 m range ($cT_1/2 = c/2f_1 = 50$ m), while maintaining high precision due to the high frequency measurement.

Assuming that operation with a modulation frequency of 100 MHz achieves measurements with 1 mm standard deviation, 3 MHz operation will achieve 33 mm standard deviation under the same operating conditions; hence the 3 MHz measurement is only useful to resolve ambiguity issues and does not contribute to the final measurement precision. In contrast, a technique can be used which is similar to two-wavelength interferometry (Creath, 1987). Using two high frequency signals, with a separation equal to the low frequency used above, e.g. f_1=97 and f_2=100 MHz, the unambiguous range of 50 m can still be achieved, and both measurements can contribute towards enhancing the measurement precision.

The long unambiguous range is determined by taking the difference between the two measured phase values at modulation frequencies f_1 and f_2

$$\Phi = \varphi_1 - \varphi_2 \tag{3.7}$$

generating Φ, the resultant phase of a so called synthetic frequency (two wavelength interferometry often refers to a synthetic wavelength Λ = 1/F (Hariharan, 1985))

$$F = |f_1 - f_2|. \tag{3.8}$$

The unambiguous range is then calculated by substituting Φ and F into Equation (3.6) in place of φ and f, and rearranging. The uncertainty in Φ, hence range, is

$$\delta\Phi = \sqrt{\delta\varphi_1^{\,2} + \delta\varphi_2^{\,2}} \tag{3.9}$$

therefore the final measurement precision is inferior to either of the two contributing measurements. This can be overcome by returning to the original phase measurements and calculating the two *ambiguous* range values. Each value will have an unknown offset of $n_x cT_x/2$, of which the integer n_x can now trivially be calculated from the determined low precision unambiguous range value (T_x is the modulation period). Adding the offset values to the two ambiguous measurements, the two range measurements are now unambiguous and independent of one another. The final range can then be calculated by averaging the two high precision measurements, resulting in a further reduction of the measurement uncertainty by $\sqrt{2}$. This is an important result, as all collected data contributes towards the measurement precision while the range ambiguity is simultaneously removed[†].

As an example, a measurement with a modulation frequency f_1=100 MHz gives a value of φ_1=3.770±0.019 rad, and f_2=97 MHz gives φ_2=3.280±0.019 rad. The synthetic frequency is therefore $F = |100-97| = 3$ MHz with a corresponding phase $\Phi = 3.770-3.280 = 0.490\pm0.027$ rad. The unambiguous object distance is

$$\begin{aligned} r_{synthetic} &= \frac{c\Phi}{4\pi F} \\ &= 3.900 \pm 0.214 \text{ m.} \end{aligned} \tag{3.10}$$

The value of n is calculated by dividing the object distance by the ambiguous depth and taking the integer quotient

[†] Range ambiguity is considered to be 'removed' when the unambiguous range exceeds the depth of a given scene.

$$\begin{aligned} n_1 &= \left\lfloor \frac{2r_{synthetic}}{Tc} \right\rfloor \\ &= \left\lfloor \frac{2f_1 \cdot r_{synthetic}}{c} \right\rfloor \\ &= 2. \end{aligned} \tag{3.11}$$

The 100 MHz modulation frequency range value can then be calculated unambiguously as

$$\begin{aligned} r_{100\,MHz} &= \frac{n_1 c T_1}{2} + \frac{c\varphi_1}{4\pi f_1} \\ &= \frac{c(2\pi n_1 + \varphi_1)}{4\pi f_1} \\ &= 3.900 \pm 0.005 \text{ m}. \end{aligned} \tag{3.12}$$

Repeating the procedure for the 97 MHz modulation frequency also gives $r_{97\,MHz} = 3.900 \pm 0.005$ m. The two measurements $r_{100\,MHz}$ and $r_{97\,MHz}$ are averaged to give a final range measurement of $r_{final} = 3.900 \pm 0.003$ m, providing the lowest measurement uncertainty.

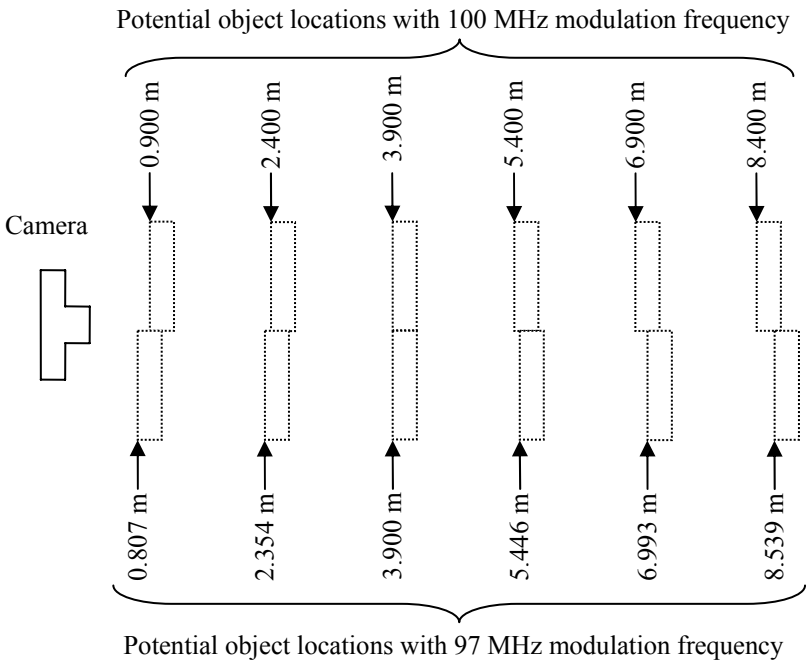

Figure 3.3: Potential object locations measured with two modulation frequencies (97 and 100 MHz). Only one common range value exists (3.900 m), hence is selected as the true object range.

An alternate approach to achieve the same result is to just compare the range values of potential object locations using n=0,1,2,… for each modulation frequency (Dorrington et

al., 2007b). Given that only one common range value will exist between the two modulation frequencies as illustrated in Figure 3.3, this value can be selected as the unambiguous range value. The two measurements can then be averaged together as before to reduce the measurement uncertainty by $\sqrt{2}$.

To reduce the acquisition time, it is possible to modulate the illumination source and receiver gain with the two high frequencies simultaneously during the measurement period (Carnegie et al., 2006), with the condition that the two beat signals generated are at different frequencies. For example 100 MHz modulation with a 1 Hz beat frequency and 97 MHz modulation with a 2 Hz beat frequency waveforms could be combined, and the amplitude and phase of each beat signal recovered independently by performing a DFT on the signal to resolve the two separate beat frequencies. As well as increasing the measurement speed, this simultaneous measurement also reduces artefacts due to motion within the scene as the object location is common to both measurements.

3.2.2 Nonlinearity due to Harmonics

Until now the illumination and receiver gain modulation signals for the heterodyne (and homodyne) techniques have been assumed to be sinusoidal – in practice they often contain higher order harmonics, which can be either intentional due to design parameters, or non-intentional due to use of non-linear components within the system. A method often used to describe the operation of the homodyne system is the use of a sinusoidal illumination

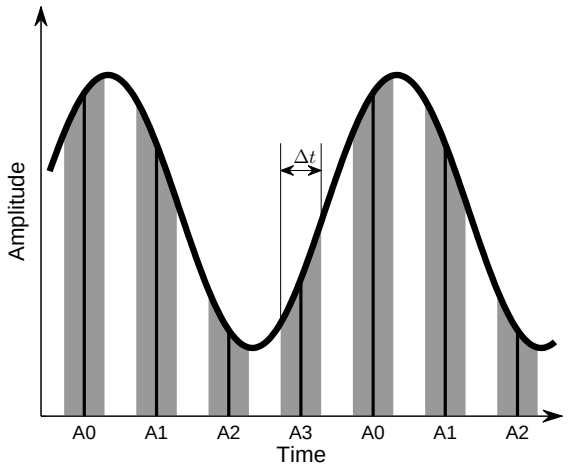

Figure 3.4: The sinusoidal illumination waveform is sampled by modulating the receiver gain with a rectangular profile, recording each sample over a time period Δt.

waveform which is sampled by a rectangular receiver waveform at the same frequency (Spirig et al., 1997; Lange and Seitz, 2001; Möller et al., 2005; Kahlmann and Ingensand, 2006). This approach is illustrated in Figure 3.4 where each sample is integrated over a time period Δt.

In this case the receiver gain modulation waveform is rectangular and can be written for a general case as

$$G(t) = G_{av}\left[1 + \sum_{i=1}^{n} m_{Ri} \sin(i\omega_R t)\right] \tag{3.13}$$

where m_{Ri} is the modulation index of the ith harmonic component, and ω_R is the fundamental frequency. Some implementations also use a square modulation waveform for the illumination source (Schwarte, 2004; Hsu et al., 2006), which again for a general case results in a reflected waveform given by

$$P(t) = P_{av}\left[1 + \sum_{i=1}^{n} m_{Ti} \sin(i\omega_T t + \varphi_i)\right]. \tag{3.14}$$

These equations are also useful for the situation where a sinusoidal waveform is intended, but due to a non-linear response or other electronic limitations the waveform becomes distorted. Through appropriate selection of the modulation index parameter for each harmonic, Equations (3.13) and (3.14) can be used to analyse any arbitrary illumination and receiver modulation waveforms. Multiplying these two equations in the presence of harmonics leads to a large number of terms as revealed by Equation (3.15), where the fundamental frequency component (i=1) and second order harmonic (i=2) have been expanded.

$$\begin{aligned} I(t) &= P_{av}\left[1 + \sum_{i=1}^{n} m_{Ti} \sin(i\omega_T t + \varphi_i)\right] \cdot G_{av}\left[1 + \sum_{i=1}^{n} m_{Ri} \sin(i\omega_R t)\right] \\ &= P_{av}G_{av}\{\ 1 + \tfrac{1}{2} m_{T1}m_{R1}[\cos((\omega_T - \omega_R)t + \varphi_1) - \cos((\omega_T + \omega_R)t + \varphi_1)] \\ &\quad + m_{T1}\sin(\omega_{T1}t + \varphi_1) + m_{R1}\sin(\omega_{R1}t) \\ &\quad + \tfrac{1}{2} m_{T1}m_{R2}[\cos((\omega_T - 2\omega_R)t + \varphi_1) - \cos((\omega_T + 2\omega_R)t + \varphi_1)] \\ &\quad + \tfrac{1}{2} m_{T2}m_{R1}[\cos((2\omega_T - \omega_R)t + \varphi_2) - \cos((2\omega_T + \omega_R)t + \varphi_2)] \\ &\quad + \tfrac{1}{2} m_{T2}m_{R2}[\cos((2\omega_T - 2\omega_R)t + \varphi_2) - \cos((2\omega_T + 2\omega_R)t + \varphi_2)] \\ &\quad + m_{T2}\sin(2\omega_T t + \varphi_2) + m_{R2}\sin(2\omega_R t) \\ &\quad + \ldots\ \} \end{aligned} \tag{3.15}$$

Again the high frequency terms can be discarded as they are filtered out during the long integration time of each captured image frame, resulting in

$$I(t) = P_{av} G_{av} \left(1 + \sum_{i=1}^{n} \frac{m_{Ti} m_{Ri} \cos(i\omega_B t + \varphi_i)}{2} \right) \tag{3.16}$$

where again the beat frequency is given by $\omega_B = \omega_T - \omega_R$.

This is a useful result as it shows that if *both* the illumination waveform and receiver gain waveforms contain harmonics *of the same order*, the output beat signal will also contain harmonics of that order. The use of two rectangular waveforms is one such occurrence of this, where the beat signal generated contains higher order harmonics, and is actually triangular in shape (Schwarte, 2004; Hsu et al., 2006). Alternatively, for the case where the illumination waveform is sinusoidal (hence contains a fundamental frequency component but no harmonics), yet the receiver modulation is rectangular (includes significant harmonic content), the beat signal only comprises the fundamental frequency component, thus is sinusoidal.

For a digitally sampled signal to satisfy the Nyquist-Shannon theorem, the sampling frequency must be greater than twice the highest frequency component within the signal. The presence of harmonics within the beat signal introduces higher frequency components; hence the sampling rate (i.e. the number of samples must be collected over the beat period) must be increased to adequately sample the signal. As the range measurement acquisition time is proportional to the number of collected samples, it is therefore preferable to ensure that harmonics within the beat signal are minimised where possible to allow a minimal number of samples to be used. A low pass filter (anti-aliasing filter) is used in sampled electronic systems to remove the higher order frequency components, but there is no such equivalent method for filtering the optical beat signal.

If the sampling conditions are not met then high frequency components are aliased, potentially onto the signal of interest. If the sampling rate f_{samp} is an exact integer multiple of the signal frequency f_{sig}, then the harmonic number of the frequencies aliased onto the signal of interest are given by (Dorrington et al., 2005b)

$$n = i\frac{f_{samp}}{f_{sig}} \pm 1, \tag{3.17}$$

where i is an integer. The aliased signals contaminate the phase measurement; hence decrease the range measurement accuracy. An example is given in Figures 3.5 and 3.6 for a homodyne case where both the illumination and receiver gain modulation waveforms are square, and four phase shifted samples are collected over the beat signal period (a typical value used to achieve a short acquisition time). From Equation (3.17), the beat frequency components aliased on to the signal of interest will occur at n=3,5,7,9,… as illustrated in Figure 3.5 where the beat signal frequency spectrum is shown.

A simulation of the resultant recovered phase is illustrated in Figure 3.6, and shows a cyclic systematic error (Lange, 2000). The visible error is predominantly due to the third and fifth harmonics, which have significant amplitude, creating an oscillation with four cycles over the 2π period. This error is visible in experimental results of demonstrated systems (Gokturk et al., 2004; Kahlmann and Ingensand, 2006), and has an amplitude and shape similar to that simulated here. This error can be calibrated for (Lange, 2000; Kahlmann and Ingensand, 2006), however this is a timely process and recalibration is likely to be required if any system parameters vary.

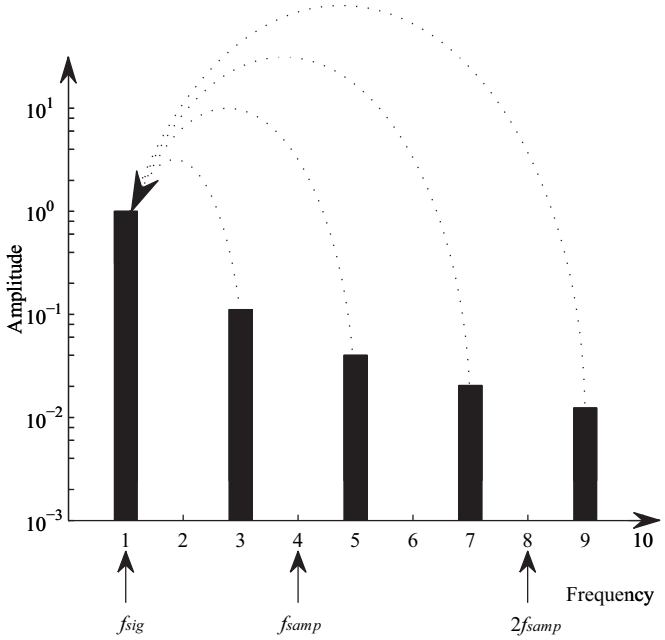

Figure 3.5: Normalised frequency spectrum of a beat signal generated from square illumination and shutter waveforms. The sampling rate is four times the signal frequency; hence all odd harmonics will be aliased onto the signal of interest during sampling.

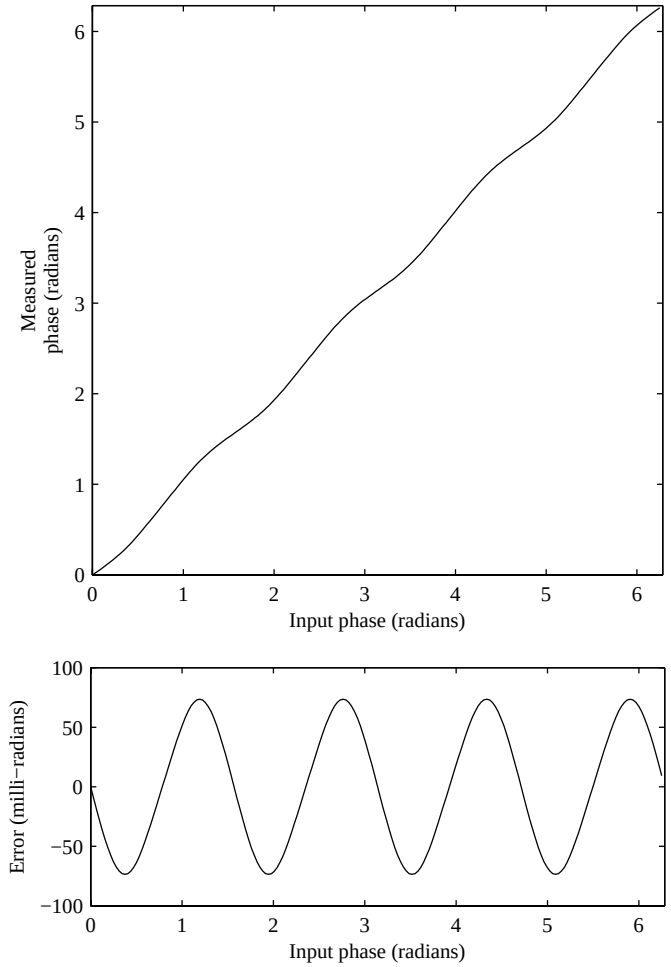

Figure 3.6: Phase error due to aliased harmonics for square illumination and receiver gain modulation waveforms with four phase shifted samples per beat signal period (homodyne method).

In interferometry, a reduction of the cyclic error can be achieved by using the integrating-bucket technique instead of the phase-stepping technique (Freischlad and Koliopoulos, 1990), analogous to using the heterodyne range method instead of the phase-stepped homodyne technique. Repeating the simulation using the heterodyne configuration, the beat signal modulation index value is reduced by a sinc function in Equation (3.5); this is also applied to the higher order harmonics. Given that the argument of the sinc function is frequency dependent, the corresponding change in magnitude is unique to each harmonic as given by

$$\begin{aligned} m_{Hi} &= \operatorname{sinc}\left(\frac{i\Delta\theta}{2}\right) \\ &= \operatorname{sinc}\left(\frac{i\pi f_B}{f_s}\right). \end{aligned} \tag{3.18}$$

Using the parameters from the previous example (i.e. four samples per beat period); the fundamental frequency component is reduced by approximately 10%. In comparison, the odd order harmonic components which are aliased onto the fundamental component during sampling, and hence cause the cyclic error, are reduced by ≥70% as illustrated in Figure 3.7.

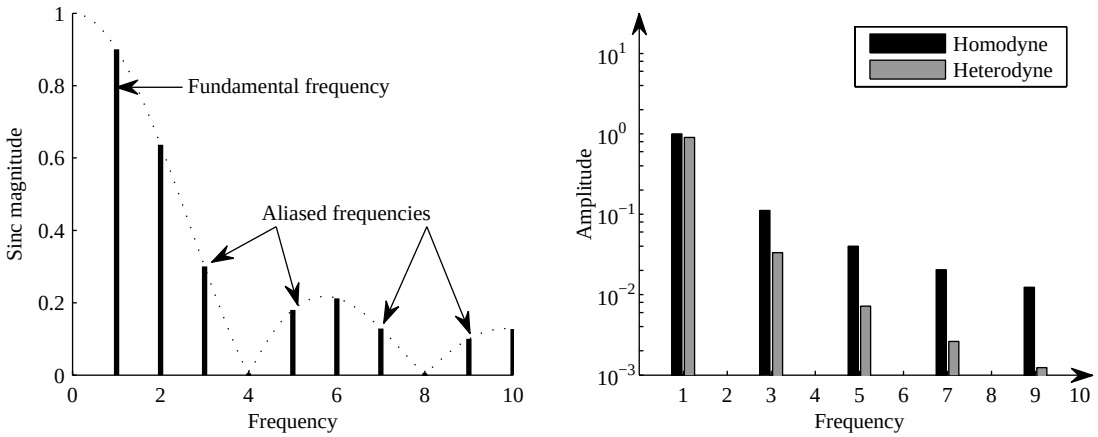

Figure 3.7: Left: The heterodyne method with four samples per beat period reduces the amplitude of the fundamental component by ~10%, while problematic harmonic components are reduced by >70%; Right: The resultant normalised amplitude of each beat signal harmonic component generated from square illumination and shutter waveforms, taking four samples per beat period. The heterodyne method provides significant attenuation of harmonic components compared to the homodyne method.

The resultant phase error is shown in Figure 3.8, where the cyclic error has been reduced from ±74 mrad for the homodyne method to ±46 mrad for the heterodyne method simulation, a 37% reduction in systematic error.

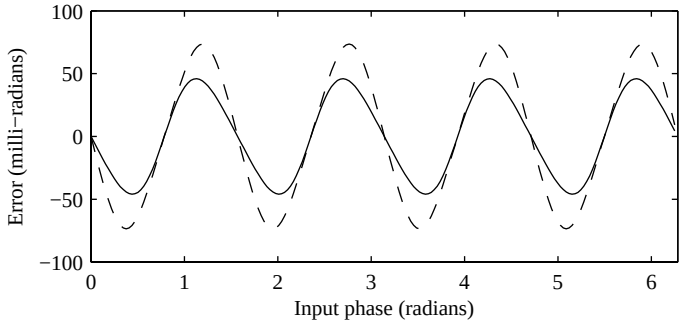

Figure 3.8: Repeat of simulation in Figure 3.6 using the heterodyne method. The amplitude of the cyclic error (solid line) has been reduced significantly over that of the homodyne method (dashed line).

This simple simulation shows that the heterodyne configuration is capable of reducing contamination of the phase measurement due to harmonics compared with the equivalent homodyne system. It must be remembered however that this gain in measurement accuracy comes at the expense of reducing the measurement precision (refer

Section 3.1.2), hence this trade off must be considered when selecting the heterodyne method over the equivalent homodyne method. The amplitude of the sinc function is dependent on the number of samples taken per beat period as illustrated in Figure 3.9 for the case of three (left) and thirty (right) samples.

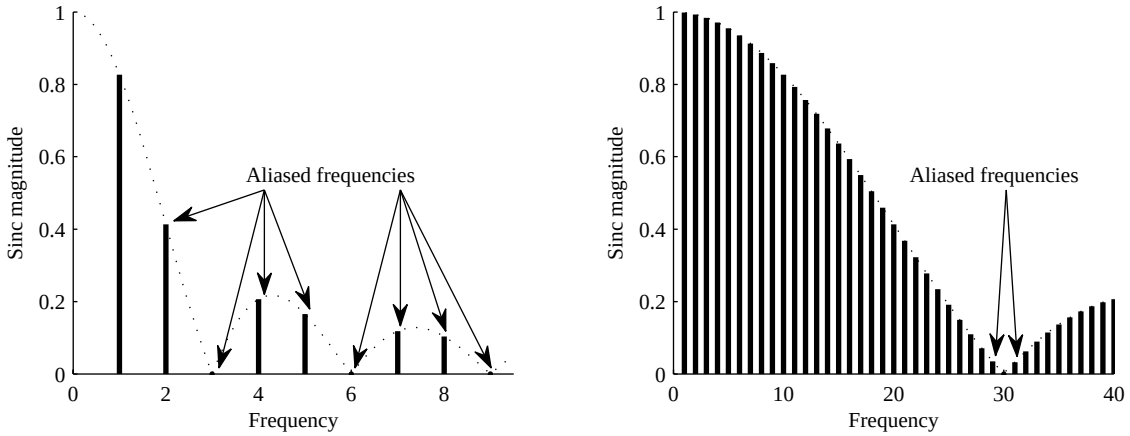

Figure 3.9: The number of samples collected per beat period has a significant impact on the amplitude of the fundamental and harmonic frequency components when using the heterodyne method. Left: 3 samples per beat period; right: 30 samples per beat period.

For three samples per beat period, the fundamental frequency component amplitude is reduced by approximately 17%, and aliased harmonics by ≥59%. The reduction of the fundamental component amplitude severely impacts the resulting measurement precision, although the reduction in harmonic amplitude reduces the cyclic error. In comparison, taking thirty samples per beat cycle reduces the fundamental component amplitude by 0.2%, and aliased harmonics by ≥97%. This is an ideal situation as the reduction of the fundamental component amplitude is negligible, yielding high precision, yet the aliased harmonics are almost completely removed, yielding high linearity. However, taking thirty samples per beat cycle has the disadvantage of extending the acquisition time by an order of magnitude over the minimum of three samples per cycle.

In situations where a short acquisition time is required, such as a moving scene, the number of samples collected is restricted. If the modulation waveforms contain significant harmonic content, it is likely that the influence of aliased harmonics will be reduced but not completely negated by use of the heterodyne method. However, there is often an unambiguous relationship between the true phase and the measured phase as shown in the top half of Figure 3.6, allowing the error to be corrected through calibration of the system (Lange, 2000; Kahlmann and Ingensand, 2006). The calibration is dependent on the shape

of the illumination and receiver gain modulation waveforms and is likely to change if the modulation frequency is varied, requiring timely recalibration of the system for each frequency used. Similarly, if an electronic component is substituted with a different part (e.g. substitution of LEDs used as the illumination source), recalibration would be required for the new response.

3.3 Summary

This chapter discussed the parameters affecting the measurement precision for the heterodyne range imaging, namely the modulation frequency, the modulation index (or depth), and the signal to noise ratio. To improve the measurement performance of the prototype range imaging system discussed in Section 2.1.8, these parameters should be optimised through development of the system hardware.

Object range is determined from the change in phase of the illumination modulation envelope; hence unambiguous measurements can only be made if the phase change is less than 2π. The use of two different modulation frequencies allows the corresponding unambiguous depth to be significantly extended. The method illustrated to combine the measurements from the two frequencies has been designed to maximise the measurement precision.

Variation of the illumination and receiver gain modulation waveform shape from a sinusoid is discussed; to avoid contamination of the phase measurement, one (or both) of the waveforms should ideally be sinusoidal and not contain higher order harmonics. If this requirement is not met, aliasing from harmonics can be avoided by using the heterodyne method with a large number of samples over the beat signal period. The higher order harmonics are then attenuated, while the fundamental frequency component of interest remains largely untouched. Alternatively a calibration routine can be used to correct the contaminated phase measurement, although recalibration is required every time the system configuration is changed (such as using an alternate modulation frequency).

4. Hardware Development I

The potential exists to upgrade a number of components within the prototype range imaging system discussed in Section 2.1.8, with the objective of reducing the measurement uncertainty from centimetres to millimetres. The principal gain is expected to be achieved by increasing the modulation frequency by an order of magnitude, with a corresponding improvement in measurement precision (refer Section 3.1.1). Enhancement of the signal to noise ratio is possible by increasing the illumination power, modulating the image intensifier with a higher voltage, and using a high-sensitivity, low noise CCD camera. The prototype system uses bench top signal generators to create the high frequency modulation signals (refer Section 2.1.8). It would be preferable to replace these signal generators with a small custom built circuit, allowing the ranging system to become a standalone unit, enhancing the prospects for commercialisation. This chapter describes the creation of a circuit to replace these signal generators. The operation of an image intensifier used as a high speed shutter is also discussed, and a high voltage, high frequency driver circuit is designed and constructed.

4.1 System Requirements

The upgraded hardware requirements are grouped into three main components:

- Signal generator:
    - Generation of two frequency locked signals.
    - Wideband operation up to 100 MHz.
    - Hertz (or sub-hertz) frequency tuning difference between the signals.

  - Reference measurement of the instantaneous phase difference between the two signals to allow absolute range measurements to be made.

- Image intensifier driver
  - Wideband operation up to 100 MHz.
  - Modulation of the photocathode with a higher voltage to enhance the image intensifier gain and spatial resolution.
- Illumination system
  - Wideband operation up to 100 MHz.
  - Wide field illumination of the entire scene.
  - High intensity to provide enhanced SNR.

4.2 Signal Generator (Direct Digital Synthesis)

Accurate signal generation and timing is critical to this project. Specifically, two wideband signals are required, operating at up to 100 MHz with single hertz, or sub-hertz, frequency difference. Typical temperature compensated crystal oscillators (TXCO) can provide stability of 0.1 ppm (parts per million), but this implies that the frequency variation between two such free running oscillators can drift by tens of hertz when operating over the required frequency range. The two signals therefore must be derived from a single reference oscillator to achieve the required stability, i.e. are "frequency locked", where any frequency drift is common to both signals.

Phase locked loops (PLLs) are capable of deriving a chosen frequency output from an input reference frequency. The tuning parameters multiply and divide the input clock frequency by integer values N and R to achieve the required output frequency F_{out} as given by Equation (4.1); however limited possible integer values constrain the output frequency tuning step size to f_{ref} (Vaucher, 2002). For example, a 10 MHz reference clock can output 100 MHz (N equal to 10, and R equal to 1), but the closest second output frequency would be 99 MHz (N equal to 99, and R equal to 10) assuming a reasonable limit for the maximum internal VCO operating frequency of ~1 GHz.

$$F_{out} = N \cdot f_{ref} = N \cdot \frac{f_{xtal}}{R} \tag{4.1}$$

Extended techniques such as adding additional division stages after the output, or fractional-N designs, where the integer divisor is dynamically changed by ±1 to synthesise fractional division (Vaucher, 2002) do not significantly alleviate the problem of requiring single hertz tuning. It should be emphasised that two output signals from two independent PLLs share the same long term drift as they can operate from the same reference oscillator; however the short term drift or phase noise is independent to each PLL. Therefore any theoretical PLL designs that can achieve the required frequency tuning are also constrained to meet strict spectral purity requirements; any phase noise, resulting in jitter in one output signal, is not common to the second signal.

Alternatively a technique known as direct digital synthesis (DDS) can be used to produce the required frequency outputs (Vankka and Halonen, 2001; Vaucher, 2002). Operating from a reference oscillator, a digital counter accumulates an input value, known as a frequency tuning word (FTW), on each clock cycle. The accumulated value is used to index a sine wave lookup table, with the output being transferred to a digital to analogue converter (DAC) to generate an output voltage. The output sine wave is quantised in both phase, due to the discrete time steps of the digital clock, and in amplitude due to the resolution of the ADC; and is therefore passed through an analogue low-pass "reconstruction" filter. The filter removes the high frequency components from the quantised output, producing a clean sinusoidal signal (Vankka and Halonen, 2001). Figure 4.1 illustrates the architecture of a DDS and the output after each stage. The accumulator overflows at the end of each cycle to automatically begin indexing the next cycle, producing a continuous output. By adjusting the FTW, the phase step size through the lookup table can be precisely adjusted in very small increments, altering the accumulator overflow rate and therefore the output frequency. The discrete points in the waveform generally change for each subsequent cycle, allowing the system to tune the output frequency with a resolution given by Equation (4.2), where j is the resolution, in bits, of the accumulator (Vankka and Halonen, 2001). By using a single reference oscillator to drive multiple synthesisers, a number of frequency locked outputs can be generated with extremely high relative stability as any drift in the reference oscillator, both short and long term, is experienced by all outputs. The output frequency can vary from DC to the Nyquist frequency (half the digital clock rate), although typically the maximum frequency used is limited to about 40% of the clock rate due to the limited performance of

the analogue filter which must remove the higher order spurious frequency components (Vankka and Halonen, 2001).

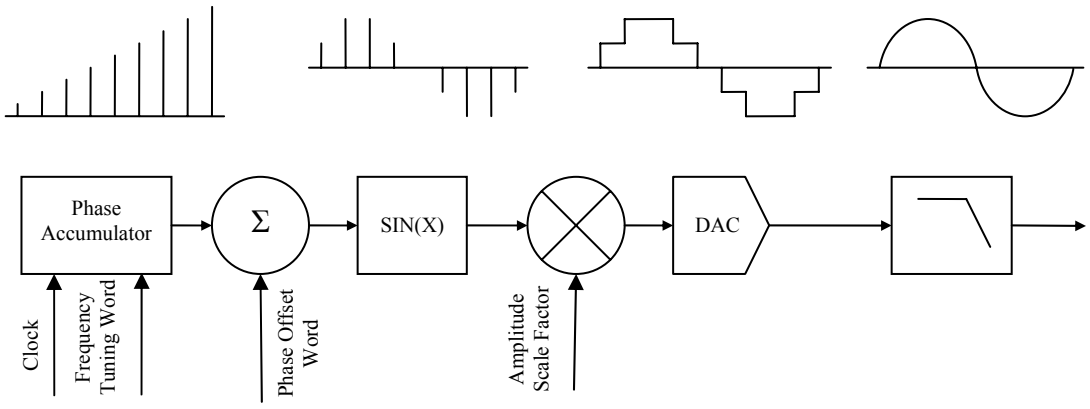

Figure 4.1: Direct digital synthesis architecture.

$$\Delta f = \frac{f_{clock}}{2^j} \tag{4.2}$$

The frequency tuning word in the DDS can be changed during operation, producing an instantaneous, phase continuous, frequency stepped output. Similarly, a phase offset or amplitude change can be made by adjusting the corresponding parameters shown in Figure 4.1; hence the DDS technology is extremely versatile and can be adapted to a wide range of applications. Despite being a conceptually simple method, the requirement of clocking the digital components at a minimum of twice the output frequency has limited the usefulness of the technique in the past due to limitations in the speed of available technology. Techniques such as mixing the DDS output with a local oscillator and filtering out the unwanted frequencies (Razavi, 1996), or multiplying the generated DDS output with a PLL (Vankka and Halonen, 2001) have increased the upper frequency limit at the expense of added complexity and compromised performance. Due to recent advancement within the commercial design and manufacturing process of these devices, low cost ICs are now available which can be clocked at hundreds of megahertz, allowing high precision frequency tuned signals to be generated over the required range (DC to ~100 MHz).

4.2.1 DDS Initial Design

A prototype DDS PCB was previously built within the Physics and Electronic Engineering Department, similar to that by Lacoste (Lacoste, 2001). The PCB was largely incomplete however, requiring changes to the hardware and software to produce a functional output. The DDS design can be divided into three sections; digital design, analogue design, and control software. A photo of the prototype DDS PCB is shown in Figure 4.2. The analogue and digital sections are separated to reduce digital switching noise contaminating the analogue outputs. A comparator within each DDS generates a square wave output which is available in addition to the sinusoidal output, hence the need for four output connectors allowing selection of sinusoidal or square output signals at each frequency.

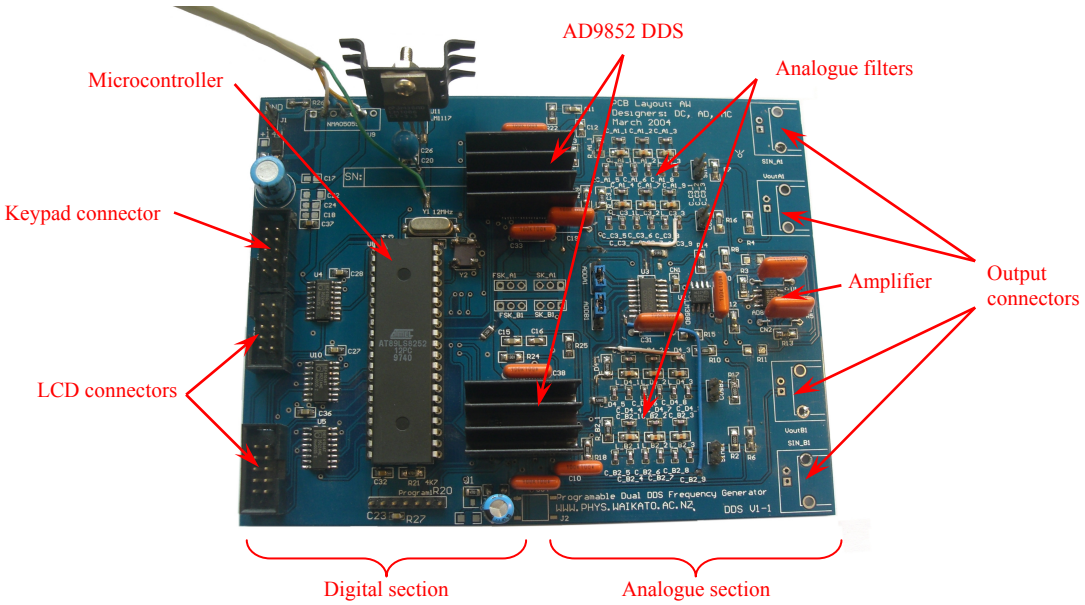

Figure 4.2: Prototype DDS PCB.

4.2.1.1 Hardware

The digital section consists of two DDS ICs, a digital clock, a microcontroller, and interface components as shown in Figure 4.3. The selected AD9852 DDS ICs (Analog Devices, MA, USA) provide the accumulator through to DAC stages shown in Figure 4.1 within a single package, and can operate with an internal clock frequency up to 300 MHz producing the required maximum 100 MHz output. The frequency tuning word is 48 bits long, which when operating at the maximum clock frequency allows output tuning with 1.07 µHz resolution (from Equation (4.2)). Both synthesisers are connected to a single 20 MHz temperature controlled crystal oscillator (TXCO), providing a common reference

clock. The clock is multiplied up to 300 MHz using an internal PLL within each DDS IC, allowing PCB routing constraints to be significantly relaxed due to use of the lower external frequency.

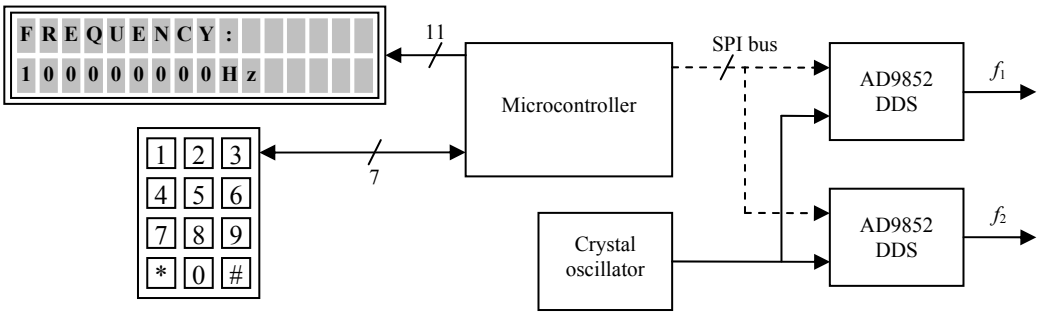

Figure 4.3: Overview of the digital section of DDS design.

The user interface is provided by a 3×4 numeric keypad, and a 16×2 LCD character display, allowing the user to type in the selected operating frequencies. A microcontroller (AT89LS8252, Atmel) interprets the users input, and is connected to each DDS by a serial peripheral interface (SPI) to set the operating parameters such as the frequency tuning word. Each DDS is programmed separately by the microcontroller, followed by a common update signal which initiates the instruction. The DDS ICs and microcontroller operate from 3.3 V, so voltage translator ICs are used to connect the signals to the 5 V LCD display. The microcontroller flash memory can be reprogrammed without removing it from the PCB by connecting a cable to the parallel port of a PC, allowing simple upgrades to the software.

The analogue section, shown in Figure 4.4, is responsible for filtering unwanted intermodulation components from the DDS output, and amplifying the output signal to a suitable level for the shutter and light source PCBs. As the DDS output is generated from discrete time samples, the spectrum contains the chosen output frequency f_{out}, but also replicas at $f_{\text{clock}} \pm f_{\text{out}}$, $2f_{\text{clock}} \pm f_{\text{out}}$,… with amplitudes given by $\text{sinc}(\pi f / f_{clock})$ (Vankka and Halonen, 2001)[‡] as illustrated in Figure 4.5. A low pass filter is used to remove these higher order components, although in specialised applications a band pass filter can be

[‡] Vankka and Halonen specify the amplitude using a *normalised* sinc function, commonly used in digital signal processing, hence the π term is not explicitly stated as shown here.

used to select one of these images for use when the required output frequency is above the Nyquist frequency (Negreiros et al., 2007).

Each AD9852 provides complementary current outputs from the DAC. The two outputs are terminated into a resistor, passed through a seventh order elliptic low pass filter with a cutoff frequency of 120 MHz, before being terminated again into a second resistor. This double termination ensures the reactive loading on the DAC is minimal, which could otherwise distort the output at high frequencies. The simulated frequency response of the low pass filter using available component values (Figure 4.6) shows a sharp cutoff and good attenuation in the stop band. Both halves of the complementary output are terminated in this manner, despite only one signal being used in the following amplification stage, in order to maximise the DAC output spurious free dynamic range (Analog Devices, 2004).

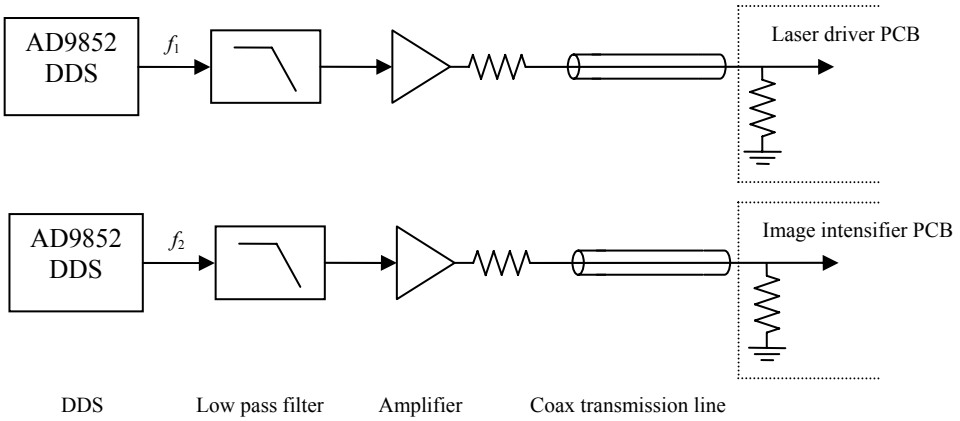

Figure 4.4: Overview of the analogue section of DDS design.

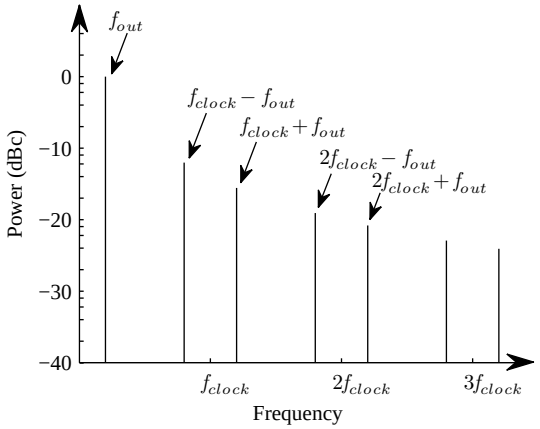

Figure 4.5: DDS Frequency spectrum before analogue filtering is applied.

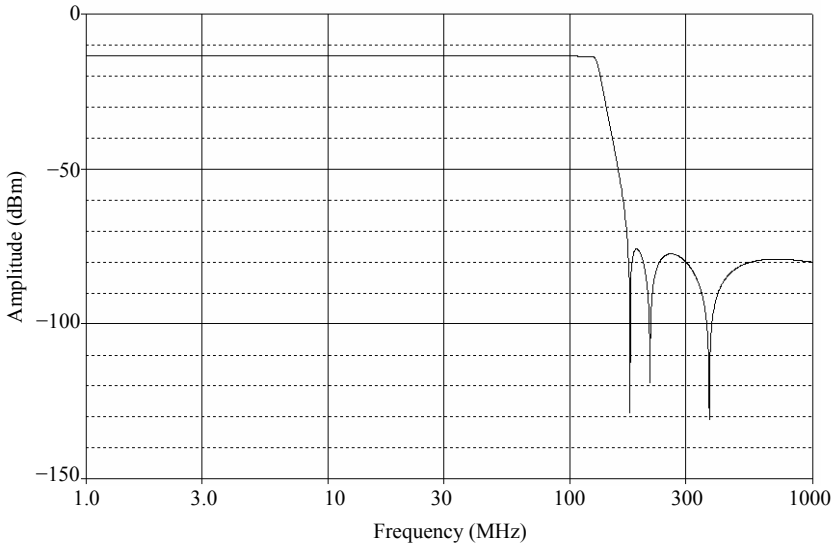

Figure 4.6: Simulated frequency response of the 120 MHz low pass filter.

After being filtered the signal is added to an adjustable DC bias and the two outputs are amplified by a current feedback amplifier (AD8002, Analog Devices, MA, USA). A 50 Ω series resistor is used to increase the source impedance to match the 50 Ω coax transmission line and termination resistances at the shutter and light source PCBs. Although this halves the amplitude of the signal, it prevents reflections from travelling through the transmission line and provides optimal signal integrity.

4.2.1.2 Software

The microcontroller software performs a number of tasks:

- Initialise the LCD display, microcontroller SPI bus, and DDS ICs.
- Read last used values from non-volatile memory.
- Interact with the user with the LCD display and keypad.
- Store the results into non-volatile memory.
- Translate a chosen user set frequency into a frequency tuning word.
- Rrogram each DDS IC and start the outputs.
- Wait for further user interaction.

Most of these tasks are self explanatory. Storing the entered data into non-volatile memory (EEPROM) makes it convenient to recall and continue using the same values after a power cycle, as typing each nine digit frequency (to specify single Hz resolution) on the keypad quickly becomes tedious.

Translating the desired frequency into a frequency tuning word (FTW) is not a trivial task for the 8 bit microprocessor. The calculation to be performed is given by

$$FTW = \frac{f_{out} \cdot 2^{j}}{f_{clock}} \tag{4.3}$$

where j is the number of accumulator bits. The system clock frequency is generally constant, and the number of accumulator bits is fixed for a specified device (the AD9852 accumulator resolution is 48 bit), so the calculation can be simplified to a single multiplication. To prevent rounding errors the constant value ($2^{j} / f_{clock}$) must be at least 54 bits long, which when multiplied by the desired output frequency (with 32 bit resolution) produces an 86 bit result. This is then rounded to the required 48 bit frequency tuning word length. Standard data types in the C compiler allow variables with up to 32

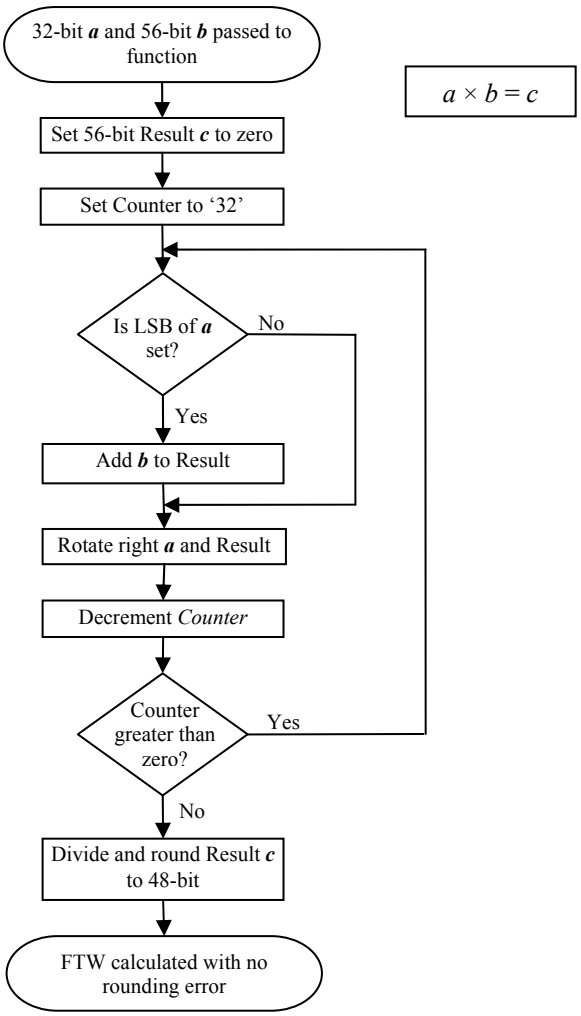

Figure 4.7: Calculating the 48 bit frequency tuning word with full precision.

bit precision (integer type unsigned long); hence it is apparent that it is not possible to perform the required calculation with this standard data type. A simple iterative solution given in Figure 4.7 was written in assembler to overcome this problem; an extension of that used by Hosking (Hosking, 2002) which uses additions and right shifts to perform the calculation without any rounding errors. A simplification was added, discarding the lowest bit of the intermediate result during each iteration, only requiring storage of 56 bits during the calculation without reducing the precision of the final resultant value. The long computation time (approximately 3 ms per FTW) does not limit the DDS performance for this application as there is no requirement to rapidly change the output frequency.

4.2.2 Limitations of the DDS Prototype

The DDS prototype design suffered from a number of issues varying from fundamental limitations to simple problems such as PCB routing mistakes. The most severe limitation arises from the need to know the phase difference between the two outputs, a requirement for performing absolute range measurements. Equation (4.4) gives the relationship between the distance r to the object to be measured, the modulation frequency f_{mod}, and the resultant phase angle φ. Using this equation, to achieve range measurements with millimetre accuracy necessitates determination of the phase with an accuracy of 4 mrad (assuming a modulation frequency of 100 MHz). The reference phase must therefore be known with better precision than the final measurement, implying that the reference phase need to be known to at least 12 bit resolution (1.5 mrad) (Payne et al., 2005).

$$\varphi = \frac{4\pi f_{\mathrm{mod}} r}{c} \tag{4.4}$$

Numerous methods are available to measure the phase of signals (Webster, 2004). Digital methods, which require an ADC to create digital samples for processing, such as a fast Fourier transform, digital phase locked loop, or digital down-conversion add significant cost and complexity due to the sampling speed required (hundreds of MHz). The zero-crossing method, measuring the time difference between the midpoints of each signal, is a suitable approach but requires relatively high speed components. Directly measuring the generated time delay requires timing resolution to the order of picoseconds, which is not practical in a low cost system. Using a flip-flop IC to generate a variable width square wave dependent on the phase of each signal and then low pass filtering (averaging) as

shown in (Webster, 2004) produces an analogue output which can be precisely measured. Webster claims that a typical accuracy of 0.02° can be achieved, although for the 100 MHz signals in this application typical CMOS or TTL logic gates are not fast enough to produce a true square wave output, resulting in a non-linear phase measurement output. Higher speed logic gate families (such as ECL or LVDS) produce a much smaller voltage swing at their outputs, and therefore induce tighter constraints on the analogue measurement stage with a corresponding increased susceptibility to noise.

Alternatively, instead of directly measuring the phase difference between the two signals, a third reference signal can be created. By configuring the initial starting phase of each output signal, and synchronously initiating the DDS outputs, the initial phase difference is known and continuously changes through 2π radians with a period equal to the frequency difference between the two high frequency signals. In practice, a third DDS operating from the same reference clock can generate this signal, with the three DDS outputs producing f_1, f_2, and f_2-f_1. The phase of the low frequency (f_2-f_1) signal, which is now in the order of hertz rather than tens or hundreds of megahertz, can be measured precisely using the simple zero crossing method (Webster, 2004).

Despite operating from the same reference clock, synchronising multiple AD9852 DDS ICs is not a trivial task. The difficulty arises from the fact that each DDS IC independently multiplies the 20 MHz oscillator clock to the system clock rate (300 MHz) using a PLL, creating independent jitter on each system clock. This jitter, or clock skew, between devices means that, potentially, the update signal may be received on a different clock edge to that of the neighbouring DDS, delaying one of the outputs. To overcome this limitation, the CMOS update signal must fall within a 1.8 ns window to ensure that the setup and hold times are met, and that all outputs start simultaneously (Brandon, 2003). Given that the manufacturer rated *variation* in propagation delay of a typical LVCMOS latch is of the order of 1.5–7 ns (Fairchild Semiconductor, 1997), meeting the timing requirements is a difficult challenge.

Analog Devices released an improved DDS IC, the AD9952, during development of the former design, with a critical feature “multichip synchronization” (Analog Devices, 2003). The AD9952 achieves this by using one master device to generate a synchronisation clock signal at one quarter of the system clock frequency, which used by the subsequent DDS

ICs to ensure the update command occurs on a common clock edge. This method does not overcome the problem of clock jitter between devices due to the use of independent PLLs, but the setup-and-hold times are inflated so that all devices process each command on the same clock edge. Jitter can be of the order of tens of picoseconds, discussed further in Sections 6.1 and 7.1, whereas clock skew generated by one DDS starting on the incorrect clock edge causes a constant delay of the order of nanoseconds. A comparison of important attributes between the former AD9852 and newer AD9952 is given in Table 4.1.

Table 4.1: Comparison between Analog Devices' DDS ICs.

Attribute	AD9852	AD9952
Maximum System Clock	300 MHz	400 MHz
FTW	48 bit	32 bit
DAC	12 bit	14 bit
Power consumption	2.70 W per device	0.16 W per device
Multi-chip Synchronisation	No	Yes

The length of the accumulator has been reduced from 48 to 32 bits, enlarging the frequency tuning step size from ~1 μHz to ~0.09 Hz, calculated from Equation (4.2) for the maximum system clock frequency of each device. This resolution still meets the single hertz tuning requirement for the ranging system. The large reduction in power consumption is significant, not only for efficiency, but also for practical constraints. The AD9852 DDS ICs are required to be soldered to a large ground plane, and the prototype design also used additional heat sinks to assist dissipation of the excess heat. A smaller power supply can be used to provide improved voltage regulation, reducing associated noise at the outputs. The improved DAC resolution enhances spectral purity by approximately 6 dB per additional bit (Goldberg, 1999) as quantisation of the produced sinusoid is reduced. Increasing the system clock rate has the obvious advantage, from the time domain perspective, that there are additional samples per output cycle for a particular chosen frequency, and as such the quantised output waveform more closely represents the ideal sine wave. From the frequency domain perspective, image frequencies (e.g. $f_{clock}-f_{out}$) have greater spectral separation from f_{out}, and therefore experience greater attenuation at the ensuing analogue filter. Harmonics produced by quantisation and DAC imperfections and nonlinearities are also aliased below the Nyquist frequency, and the increase in system clock frequency helps to reduce their influence by shifting the lower order harmonics (which generally have larger amplitude than higher order harmonics) into the stop band of the analogue low pass filter (Brandon, 2004).

4.2.3 DDS Design Revision

The second, and third, DDS PCB design revisions addressed the limitations described in Section 4.2.2. The full schematic diagram of the final design can be found in Appendix A. Three AD9952 DDS ICs are used with the interconnecting synchronisation clock to generate two frequencies, and the difference between them (equal to the beat frequency), all with a known initial phase (Payne et al., 2005). The revised digital configuration is illustrated in Figure 4.8. A number of other small improvements were also implemented in these PCB revisions. A buffer is used to increase the transition rate of the clock edges from the crystal oscillator, reducing clock skew between each DDS and reducing loading at the oscillator. A high stability temperature compensated crystal oscillator was used, providing stability of 1 ppm (part per million). An RS232 serial interface has been added to allow automated frequency control from a PC. The in-system-programming (ISP) cable, used to update the microcontroller firmware, (not shown in Figure 4.8) previously had to be disconnected from the PCB during normal operation; tri-state buffers are used in the later revisions to allow the cable to remain attached during development.

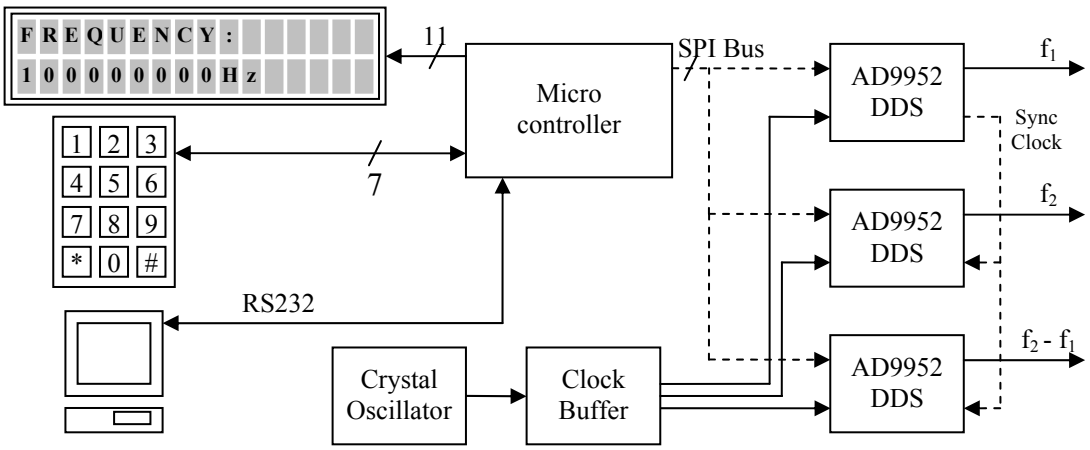

Figure 4.8: Overview of revised digital DDS layout.

During range data acquisition the low frequency optical beat signal is captured by the CCD and the signal phase is extracted from the data using a DFT as described in Section 2.1.8. Spectral estimations obtained using the DFT inherently assume that the signal is periodic over the sampling period (Webster, 2004); if this condition is not met, *spectral leakage* may occur where the spectral power is spread across multiple frequency bins. This is also sometimes referred to as a detuning error, for example in phase measuring interferometry

(Freischlad and Koliopoulos, 1990), where the phase shift between samples suffers from a constant offset error.

The camera frame rate is internally generated, operating at a standard NTSC frame rate of 29.97 frames per second (fps) (Carnegie et al., 2005), while the low frequency beat signal is generated by the DDS. The discrete frequency tuning steps of the DDS are constrained by Equation (4.2), which for the AD9952 corresponds to multiples of 0.0931 Hz (operating at the maximum 400 MHz system clock rate). As there is no common factor between the CCD frame rate and the DDS outputs, an *exact* integer number of samples cannot be taken over the beat signal period. The resulting spectral leakage can be reduced by; 1) increasing the observation period, and 2) applying a suitable window function to the measured signal (Webster, 2004). A long observation period is undesirable as it reduces the acquisition speed of the ranging system, as each sample must be acquired sequentially. Using a window function attenuates the samples near the beginning and end of the captured sequence, effectively reducing the signal duration (Webster, 2004), hence decreasing the SNR. This reduction can be determined for a particular window by calculating the processing gain, equal to the inverse of the normalised equivalent noise bandwidth (ENBW), with the ENBW values of typical windows ranging from 1.3–1.6 (Harris, 1978). A reduction of SNR is clearly undesirable as the measurement precision is reduced, as given by Equation (3.1). Another limitation of using a window is the impracticality of applying it to data captured over a short acquisition period, where the image sequence may only consist of three or four captured frames.

In order to minimise the phase measurement error due to spectral leakage, the existing processing software zero pads the FFT to eight times its original length. This results in a broadened peak in the FFT. An interpolation between frequency bins is then used to find the magnitude peak location with sub-bin accuracy, from an average taken over 4000 or so pixels with the strongest signal strength. This measured beat frequency is then used to determine the phase, hence range, for each pixel in the scene (Carnegie et al., 2005). This is a very computationally expensive task, requiring large amounts of memory and CPU time.

A preferred alternative is to ensure that the data are periodic during the observation time, so that the beat frequency lies exactly on one Fourier frequency bin. This can be achieved

by synchronously sampling the beat signal by ensuring the CCD frame rate is an integer multiple of the DDS generated beat frequency. To achieve this, a fourth output is generated from the DDS PCB operating at a frequency of $N{\cdot}(f_2{-}f_1)$, where N is an integer corresponding to the required number of samples per beat signal period. The CCD is then externally triggered by this signal. The computation required to recover the phase is now simplified as the beat frequency aligns precisely with one bin in the Fourier spectrum[§], allowing the inner product Equation (2.27) to be used instead of a full FFT.

An additional DDS IC could be added to produce the fourth output from the PCB, however the integer frequency relationship between the camera trigger signal and the reference phase output means that a simple digital counter can instead be used. The third DDS is reconfigured to produce N times its original frequency, and a comparator produces a square wave from the sine wave zero crossing. The square wave is connected to the CCD external trigger input, and also to a divide by N counter. Each transition at the counter output is therefore synchronised at a constant phase difference between the two high frequency DDS outputs. The counter must be initialised at the beginning of each acquisition, and the value of N may change, so to simplify the design the digital counter is implemented by driving an interrupt pin in the existing microcontroller, with a software counter controlling an output pin. Figure 4.9 shows the configuration of the third DDS output, while Figure 4.10 shows a timing diagram of all four outputs.

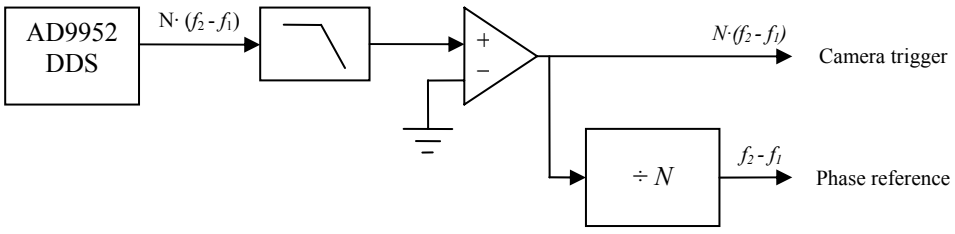

Figure 4.9: Synchronised camera trigger output.

A meticulous reader may realise that the comparator produces an asynchronous output, yet the phase reference output is now being sampled and updated by a microcontroller, with a

[§] The object and range finder must remain stationary during data acquisition. A Doppler shift occurs if the object is moving with a velocity component in the direction of ranging, again causing spectral leakage of the beat frequency. The Doppler shift can be measured to correct this range error, and to also produce a velocity measurement of the object (Carnegie et al., 2006), but is outside the scope of this thesis.

corresponding synchronisation to the microcontroller clock. The software processing time also introduces a variable delay between the true phase referenced time and the change in the output pin. However, as the camera frame rate is asynchronously generated by the third DDS, the start of each frame occurs at a consistent phase offset. For example, if N equals 8 then the camera trigger occurs when the phase difference of the first two outputs equals 0, 45, 90,…, 315°. The phase reference output pin only needs to indicate which captured frame started integration with the 0° phase offset, tagging the resultant video stream; hence the exact timing is not critical so long as the signal falls within the correct camera frame period. The phase reference signal is passed to the video acquisition computer where the frame tagging is performed by the video acquisition software, which is outside the scope of this thesis.

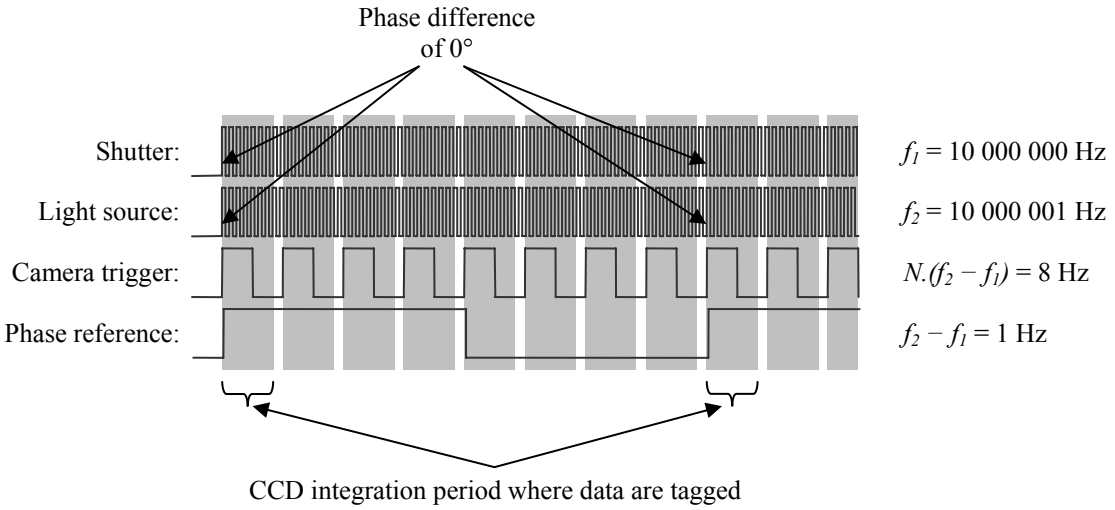

Figure 4.10: Example timing diagram of DDS PCB outputs.

The PCB layout was improved from the initial prototype by introducing split ground planes and independent power supplies for the digital and analogue sections within the design. This was done to reduce cross talk between the digital reference oscillator and the analogue outputs. In the earlier design, the output is driven by a dual current feedback amplifier (AN8002, Analog Devices, MA, USA) which amplifies both high frequency DDS outputs within one device. Noticeable crosstalk between the two outputs exists, presumably within the internal power supply rails of the amplifier. This part was replaced with independent discrete amplifiers for each DDS output, each with associated decoupling capacitors to minimise crosstalk between the high frequency outputs. A differential amplifier configuration was also implemented to utilise the voltage swing from both complementary DDS outputs, and is shown in the schematic in Appendix A. This reduces

the gain required compared with the single ended amplifier, and also produces an output centred at zero volts. A RS-232 serial port was included to allow control from either the keypad/LCD display or the PC.

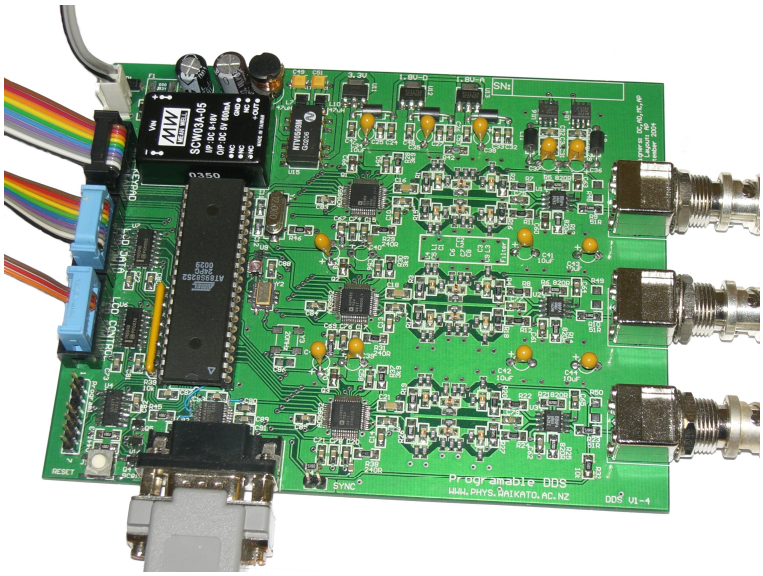

Figure 4.11: Photo of final DDS PCB.

Figure 4.11 shows the final DDS PCB. The layout and size is similar to the original prototype, although the power supply is integrated on the PCB, a third DDS to generate the phase reference (and camera trigger) signal is included, and the serial interface has been added. The PCB dimensions are approximately 142×112×20 mm, and the power consumption is approximately 3.5 W.

4.3 Image Intensifier

The high speed shutter required for image modulation is implemented using an image intensifier. The image intensifier was chosen because it is available as an off-the-shelf component, can be modulated at frequencies in excess of 100 MHz, and can be connected to a CCD to provide high spatial resolution (Bosch, 1995). An image intensifier works by focusing light collected from a scene onto a transparent window which is coated with a photosensitive compound, called a photocathode. The incoming light strikes this photocathode, and an electron is emitted due to the compound's low work function (the photo-electric effect). An electric field accelerates the electron towards a microchannel plate (MCP), where collision with the walls produces secondary electrons as illustrated in Figure 4.12. Electron multiplication by up to four orders of magnitude occurs within the MCP, before the electrons are again accelerated by an electric field, impinging on a

luminescent (phosphor) screen, converting the image back into light (Payne et al., 2006a). The output image can then be measured using a standard CCD camera.

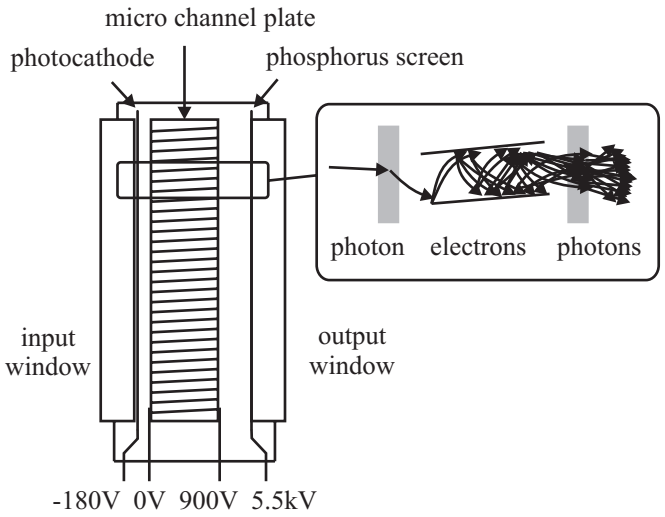

Figure 4.12: Operation of an image intensifier (Payne et al., 2006a).

Some typical operating voltages are shown in Figure 4.12 from the photocathode, across the MCP, to the luminescent screen. Two methods are available to manipulate the image intensifier gain; i) adjust the voltage across the MCP, or ii) adjust the voltage between the photocathode and the MCP input (Bosch, 1995). Adjusting the MCP voltage changes the amplification, however continuous high frequency modulation of the MCP voltage is a difficult task to accomplish due to the high voltage used (hundreds of volts). In comparison, the photocathode voltage is much lower and is a more suitable candidate for modulation up to hundreds of megahertz (Bosch, 1995). A negative photocathode voltage creates an electric field which accelerates the electrons towards the MCP, creating an image at the output, whereas a positive photocathode voltage deflects the electrons away from the MCP and the image intensifier is turned off.

The output image resolution is dependent on the photocathode voltage (Bosch, 1995), as the resultant electric field is responsible for focusing the image onto the MCP. Each electron released from the photocathode has an initial random velocity, including a transverse component parallel to the image plane. This transverse component causes the electron to follow a parabolic trajectory as it is accelerated towards the MCP (Thomas et al., 1993) as illustrated in Figure 4.13; hence a point source at the input results in an enlarged spot at the output. By increasing the voltage between the photocathode and MCP, the electric field strength, hence acceleration experienced by each electron, increases; and

the time required for the electron to travel between the photocathode and MCP decreases. The image quality is therefore dependent on the photocathode voltage, with the radial displacement of a point source proportional to $1/\sqrt{V}$ (Thomas et al., 1993). The existing prototype system uses a relatively low modulation voltage of 4 V_{pp} (Carnegie et al., 2005), hence it is desirable to substantially increase this voltage to improve the resultant image resolution.

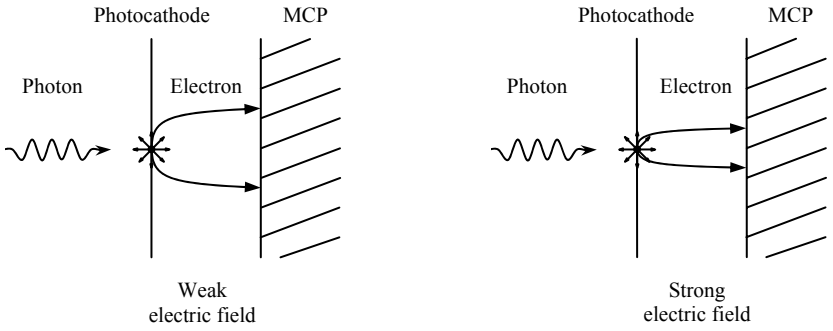

Figure 4.13: The electric field strength between the photocathode and MCP determines the image focus within the image intensifier.

The image intensifier transfer function of optical gain versus photocathode voltage is known to be non-linear (Bosch, 1995), especially when crossing from a negative to positive voltage which causes an abrupt cutoff in gain. The (static) transfer function of the image intensifier used, MCP125 (Photek, East Sussex, UK), was measured by illuminating the input window with a uniform light source, applying a DC voltage to the photocathode,

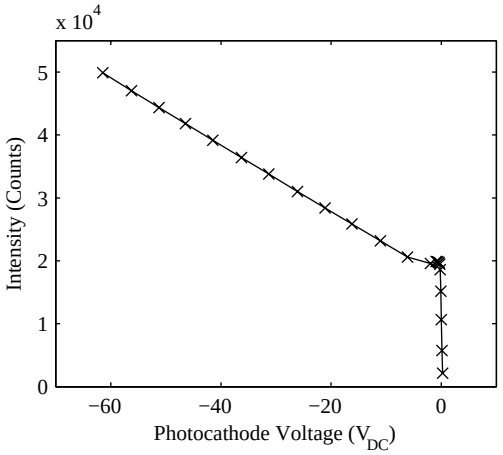

Figure 4.14: Measured image intensifier gain over a range of photocathode voltages (Payne et al., 2006b).

and measuring the output intensity with the CCD camera (Payne et al., 2006b). The response is shown in Figure 4.14. A positive voltage causes the output to drop to zero (slightly above zero in the graph due to the CCD dark current). A small negative voltage

allows the photocathode electrons to reach the MCP, providing an abrupt increase in gain. Larger negative voltages increase the kinetic energy of each electron, producing higher gain within the MCP.

To produce the required sinusoidal gain modulation, as given by Equation (2.10), requires a modulation voltage operating exclusively within one of the linear regions of Figure 4.14. The region of 0 to −2 V produces low spatial image quality due to the poor focus between the photocathode and MCP. Alternatively, the −5 to −60 V (or greater) voltage region provides improved spatial resolution, at the expense of a poor modulation index value as the gain cannot approach zero. A high modulation index is vital to produce range measurements with high precision (refer Section 3.1); therefore exclusive operation within either of the linear regions produces undesirable results.

In gated applications, the image intensifier is pulsed on for a short period with a voltage spanning +50 V to −200 V. In these applications, the FWHM is typically the most important attribute, rather than attempting to produce sinusoidal optical gain modulation. A similar voltage regime can be used for continuous modulation; however the optical gain modulation waveform will contain additional harmonics such as given in Equation (3.13). Section 3.2.2 discusses the effect of these harmonics, and shows that their influence can be minimised by using the heterodyne method with a large number of samples per beat period, allowing the range imaging system to function correctly despite the presence of a non-sinusoidal modulation waveform.

Image intensifiers are seen primarily as a capacitive load, with a small series resistance due to the photocathode material. The particular image intensifier used here was measured with a network analyser providing values of 60 pF and 12 Ω respectively. From these values, the predicted power dissipation within the photocathode for various modulation voltages is shown in Figure 4.15. For 250 V_{pp} (+50 V to −200 V), as used in pulsed applications, continuous modulation at 100 MHz would dissipate over 100 W within the image intensifier, likely overheating and damaging the photocathode material.

To balance the requirements of maximising the image focus and modulation depth, while minimising the power dissipation, a voltage level of 50 V_{pp} (+10 V to −40 V) was selected; the dynamic per-volt improvement in gain and spatial resolution rapidly diminishes for

larger modulated voltages (Bosch, 1995). The asymmetry around zero volts allows the maximum negative voltage to be applied, maximising the image resolution, while also ensuring the image intensifier is turned off completely during the positive part of the cycle. To further improve the spatial resolution the waveform should rapidly transverse the poor focus, low voltage region, hence a square electrical modulation waveform is preferred.

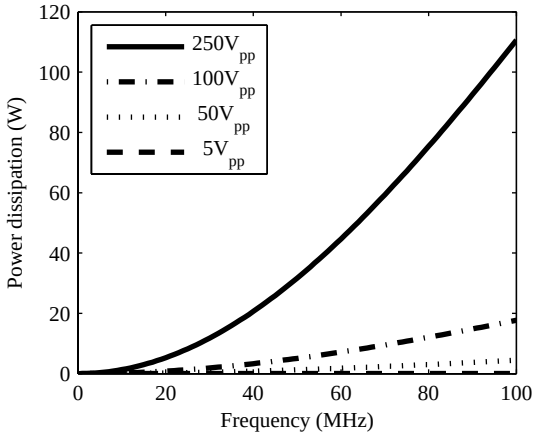

Figure 4.15: Predicted power dissipation within the image intensifier for various modulation voltages.

4.3.1 MOSFET Driver

Adapted from a pulsed image intensifier driver design (Behrend et al., 2004), a switching MOSFET driver was designed for continuous modulation with a voltage up to 100 V_{pp}. A switching (class D) amplifier is desirable due to the high efficiency provided (Whitaker, 2002) and inherent suitability of generating a square wave output. A simplified overview of the circuit constructed is shown in Figure 4.16, and a photo shown in Figure 4.17. The comparator creates two complementary square wave signals from the DDS signal to drive the high and low side MOSFETs respectively. The diode-resistor-capacitor network adds an adjustable dead time by delaying the turn-on time for each MOSFET without delaying the turn-off time in a process similar to that used by Bernius and Chutjian (Bernius and Chutjian, 1990). This prevents cross conductance from occurring where both transistors are turned on simultaneously. Both MOSFETs are N-channel enhancement type, so the high-side driver and high-side MOSFET must be floating at the output voltage; a digital isolator (ADuM1100, Analog Devices, MA, USA) is therefore used to provide an isolated signal to the high-side MOSFET driver. The low-side signal is also passed through a digital isolator in order to match the propagation delays to the two transistors. The digital isolator and MOSFET driver (DEIC420, Directed Energy Inc., CO, USA) can operate at

frequencies up to 50 MHz, and the MOSFETs (DE150-101N09A, IXYSRF, CO, USA) limit the output voltage due to the maximum V_{DS} of 100 V. Although this frequency is below the ideal 100 MHz, it is significantly higher than the 10 MHz frequency used by the prior prototype system described in Section 2.1.8 (Carnegie et al., 2005).

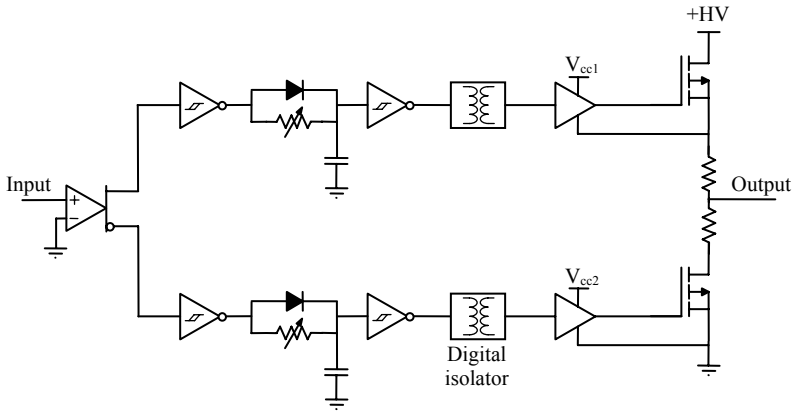

Figure 4.16: Overview of the MOSFET based image intensifier driver.

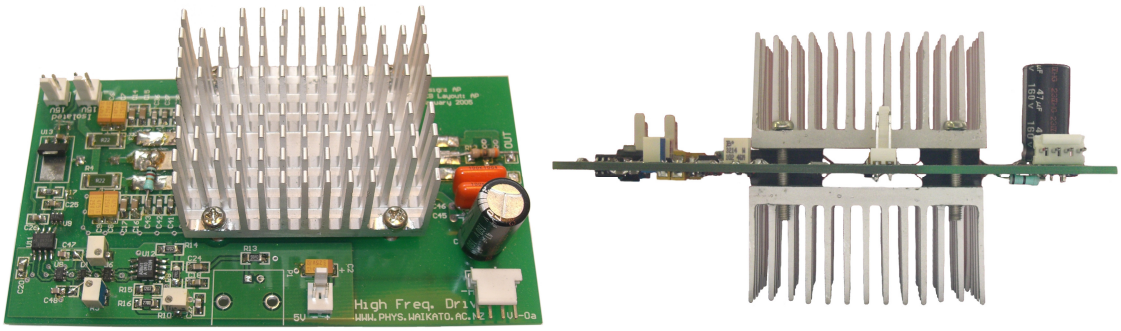

Figure 4.17: Constructed MOSFET image intensifier driver. The two transistors are mounted symmetrically on either side of the PCB, and sandwiched between two heatsinks.

The major limitation experienced by this design is due to the isolation of the high-side MOSFET driver power supply V_{cc1}. The MOSFET driver requires an isolated power supply to provide 4.5 A at 15 V for continuous modulation at 40 MHz. Typical power supplies capable of providing this output have an isolation capacitance (between their input and output) in excess of 1 nF. As the power supply floats at the circuit output voltage, this capacitance substantially increases the loading at the output, as it is seen in parallel with the image intensifier photocathode capacitance. The increased capacitive loading lowers the circuit efficiency and also encourages oscillation at the output during switching. The power consumption within the circuit is also significant: the two MOSFET drivers dissipate approximately 70 W each (estimated from the voltage and current specifications in the device datasheet), and the MOSFETs dissipate approximately 100 W (50 W each) solely due to the capacitive load, calculated using Equation (4.5) (assuming perfect

switching of a 50 V_{pp} output at a frequency f of 40 MHz, with a total load capacitance C of 1 nF).

$$P = CV_{pp}^{2} f \tag{4.5}$$

The isolated power supply used during testing (PT4583, Texas Instruments, TX, USA) was only capable of providing 2.5 A to the high-side MOSFET driver, limiting the modulation frequency during testing to approximately 20 MHz. The circuit output rise and fall times for a 60 V supply rail are approximately 5.2 and 5.6 ns respectively as shown in Figure 4.18, illustrating that the circuit is primarily limited by the maximum switching frequency of the MOSFET drivers. With additional improvements to the power supplies, inclusion of forced air cooling, and damping of the output voltage oscillation, the circuit would be potentially capable of operating near the 50 MHz maximum frequency.

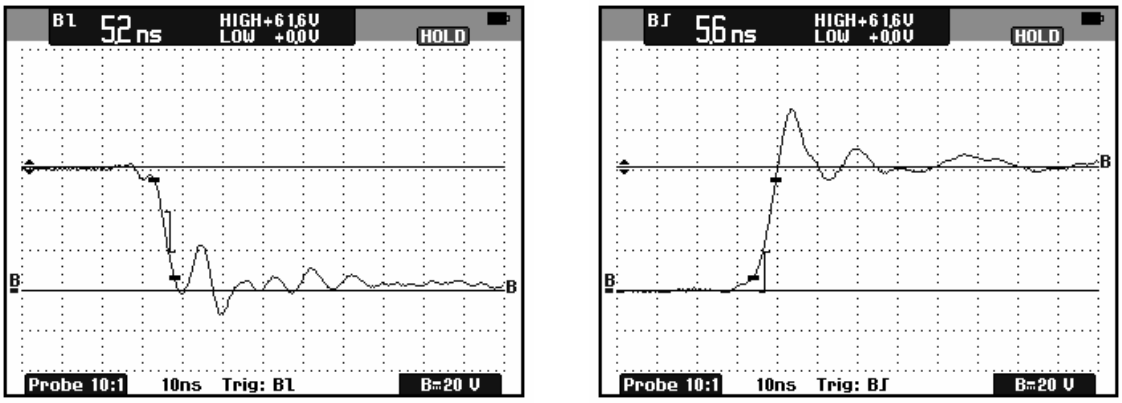

Figure 4.18: Rise and fall times of the image intensifier MOSFET driver circuit.

4.3.2 CRT Driver

A second image intensifier driver was developed concurrently with the MOSFET circuit, based on a CRT driver. The CRT driver (LM2412, National Semiconductor, CA, USA) is a monolithic class AB amplifier designed to operate with an output rise/fall time of 2.8 ns, and an output swing up to 50 V_{pp}. Three amplifier channels are provided within the single device to drive the red, green, and blue channels within a CRT monitor, with each channel rated for stability with a 0–20 pF capacitive load. To drive the image intensifier (with capacitance of approximately 60 pF), all three channels were combined in parallel as shown in Figure 4.19.

The input is passed through a comparator to create a 5 V square wave output. Before reaching the CRT amplifier, the square wave is buffered by an inverting logic gate (NL17SZ14, ON Semiconductor, AZ, USA). The logic gate can operate over a voltage range of 1.65 to 5.5 V, and provides an over-voltage tolerant input, allowing the square wave amplitude to be adjusted by changing the logic gate supply voltage. This provides a method of controlling the gain and hence final output peak to peak voltage. The signal is simultaneously amplified by all three channels of the CRT driver, and combined through small value resistors (<10 Ω). The resistors prevent any large current from flowing between the separate CRT channels if any small mismatches exist in the timing of the three outputs.

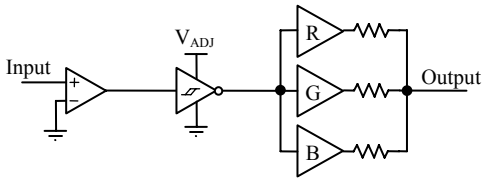

Figure 4.19: Overview of the CRT amplifier based image intensifier driver.

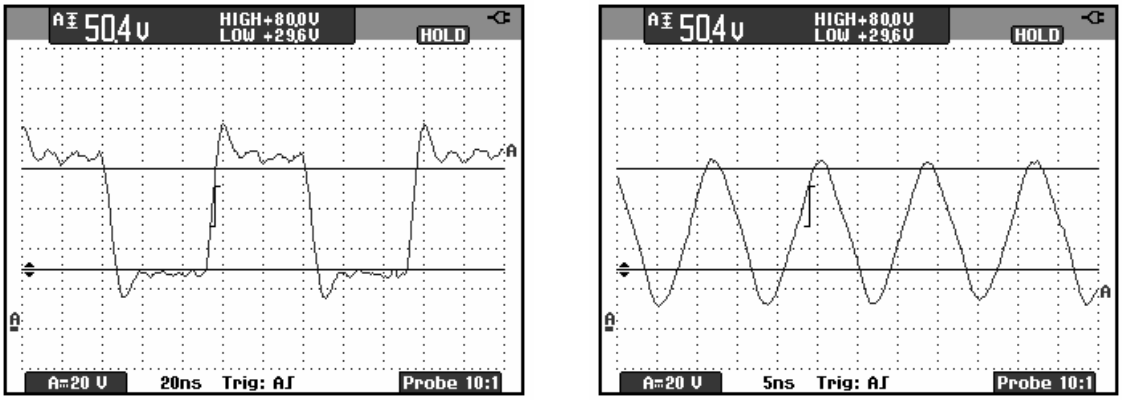

Figure 4.20: Modulation waveforms for the CRT based image intensifier driver at 10 MHz (left) and 75 MHz (right) when connected to a 60 pF load. The target of 50 V_{pp} is exceeded over the operating frequency range tested (up to 75 MHz).

The output waveform is shown in Figure 4.20 at 10 MHz and 75 MHz modulation frequencies, connected to a 60 pF dummy load (a 47 pF capacitor in parallel with the 12 pF oscilloscope probe). The Figure shows that the target of 50 V_{pp} is met over the entire operating frequency range tested, and the rise and fall times are approximately 3.5 ns (corresponding to a bandwidth of approximately 100 MHz). The modulation frequency was limited during testing due to the 80 V power supply used, which was only capable of sourcing approximately 300 mA (25 W). The power supply has since been replaced, removing this restriction. It is clear that the very simple CRT circuit out performs the

MOSFET based circuit with regards to the maximum achievable modulation frequency, and operates with much lower power dissipation, hence the MOSFET circuit was not further developed.

The power dissipation limit within the image intensifier is largely unknown, and although it is significantly reduced by only using 50 V_{pp} compared with 250 V_{pp}, the calculated dissipation is approximately 4 W for a modulation frequency of 100 MHz (as shown in Figure 4.15). As a comparison, for typical pulsed operation using 250 V_{pp}, and 5 ns FWHM at a repetition rate of 200 kHz, the power dissipation within the image intensifier is approximately 0.2 W, an order of magnitude lower than that for continuous modulation with 50 V_{pp}. The immediate concern is that heating of the photocathode material may be sufficient to damage the device. Rather than further reducing the modulation voltage, which would reduce the spatial image resolution, the circuit was modified to operate in a similar fashion to a pulsed system, ignoring some of the square pulses from the comparator. In this manner the power dissipation within the image intensifier can selectively be reduced proportionally with the ratio of ignored pulses to those kept. An overview of the circuit to achieve the pulsing is shown in Figure 4.21, and the detailed final schematic can be found in Appendix A.

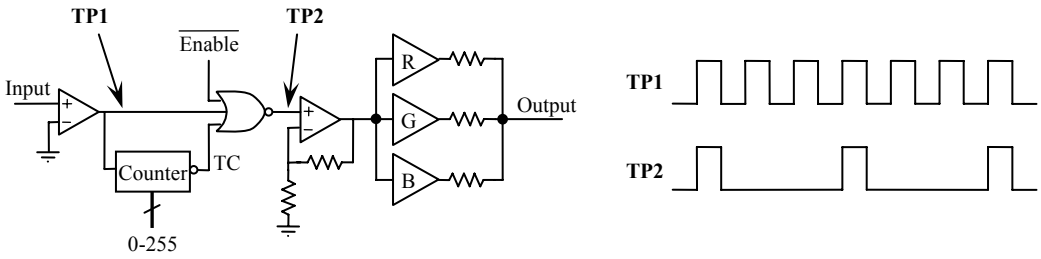

Figure 4.21: Left: The modified CRT circuit which allows the power dissipation to be reduced by selectively ignoring the input pulses. Right: The voltage at test point 1 and test point 2 for a counter value of 253, where the output signal is reduced by 2/3 to reduce power dissipation.

A digital counter is preloaded with a value from a parallel input, and is incremented by the output signal from the comparator. When the counter reaches the value 256, a terminal count (TC) signal becomes active, and the counter is reloaded with the parallel input value. The terminal count signal is NOR'ed with the comparator signal, reducing the pulse repetition rate as shown at test point two (TP2) in Figure 4.21. The number of "missing pulses" is equal to 255 minus the input value, allowing normal continuous operation or a reduction of the output repetition rate, hence power dissipation, by up to 255. ECL logic is

required for the comparator, counter and NOR gate, as the propagation delay for CMOS/TTL devices is inadequate for operation at 100 MHz. The circuit is configured so that signal pulse width is not altered by inclusion of the counter propagation delay. The NOR gate output signal is converted to TTL level (not shown in Figure 4.21), before being buffered by a high speed opamp (THS3201D, Texas Instruments, TX, USA). The opamp provides greater control over the signal amplitude and offset than the earlier method of changing the logic gate supply voltage, which also suffers from a slight speed versus supply voltage dependency. An output waveform is shown in Figure 4.22, where a single pulse has been selected from a 100 MHz input from the DDS. The output voltage swing is close to 50 Vpp even when operating at this high frequency.

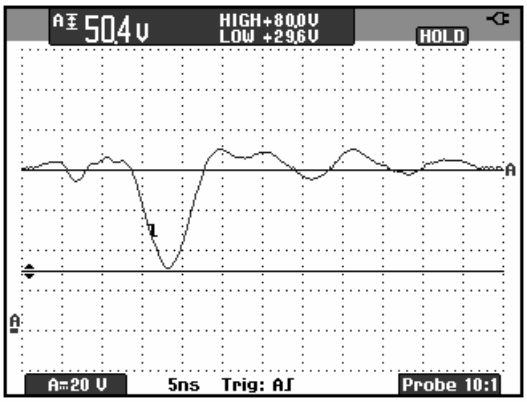

Figure 4.22: Output from the CRT amplifier based image intensifier driver where a single pulse is generated from a 100 MHz input.

A photograph of the driver circuit is shown in Figure 4.23, where it connects to a second PCB to add the required MCP and phosphor screen bias voltages before connection to the image intensifier. The photocathode modulation voltage is required to be in the range of −40 V to +10 V, yet the CRT based driver output is in the range +30 V to +80 V. As the required photocathode voltage is relative to the MCP input voltage, the bias PCB shifts the MCP input voltage from 0 V to +70 V. The required MCP output and screen voltages (refer Figure 4.12) are also added at the bias PCB. An enable input included in the circuit design allows the image intensifier to be turned off, i.e. the photocathode voltage to remain positive, while the high voltage MCP and screen bias voltages stay turned on. This provides the ability to disable the output modulation voltage and allow the driver IC to cool down during any periods of inactivity (as heat is primarily produced during switching), without having to shutdown the high voltage regulators which have a significantly longer start up period.

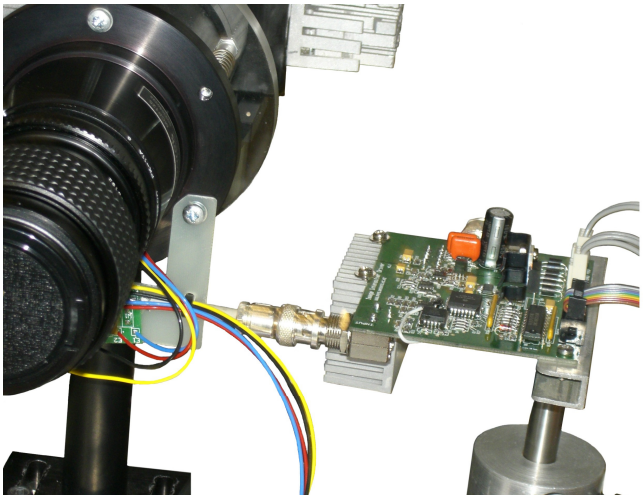

Figure 4.23: Constructed image intensifier driver based on a CRT amplifier, and connection to the bias PCB (beneath the camera lens).

4.4 Summary

In this chapter a direct digital synthesiser (DDS) circuit was designed and constructed to generate two high frequency signals in excess of 100 MHz, with tuning of 0.09 Hz, and relative frequency stability in excess of nine orders of magnitude. This extremely high stability is required for the illumination and the image intensifier modulation waveforms to accurately produce the low frequency beat signal. The final circuit uses newly available DDS ICs that are capable of being synchronised, allowing a third output to be generated at a low frequency. The low frequency output is used to trigger the camera frame capture, and to derive a phase reference signal, allowing the range imaging system to perform absolute distance measurements.

The operating mechanism of an image intensifier is discussed, where modulation of the photocathode voltage allows the device to operate as a high speed optical shutter. The optimal operating voltage was selected to balance the requirements of image resolution, modulation depth, and power dissipation. A MOSFET based driver circuit was constructed, but failed to meet the ideal system requirements. A second circuit was also built using a CRT driver, providing an output of 50 V_{pp} at frequencies up to 100 MHz. The circuit was also modified to selectively remove pulses from the modulation waveform, reducing the power dissipation within the image intensifier to avoid damage due to overheating.

5. Hardware Development II

In the previous chapter, a DDS circuit and an image intensifier driver circuit were designed to operate at frequencies up to 100 MHz; hence a modulated illumination source is also required to operate at a corresponding frequency. In addition to speed, the illumination wavelength and eye safety considerations are important factors in the system design. A number of auxiliary signals are also required to control the various electronic components within the system, suited to generation using a FPGA.

5.1 Illumination Source

The illumination source is required to be intensity modulated over the same frequency range as that of the image intensifier, namely at frequencies up to 100 MHz. Flood illumination is necessary in order to provide imaging of all pixels in the entire scene simultaneously; therefore the illumination source must be capable of providing a high radiant intensity over the measurement volume. Wavelength selection is predominantly limited by the responsivity of the image intensifier photocathode material, which is shown in Figure 5.1. Peak sensitivity occurs at a blue wavelength (450 nm), although the image intensifier has useful sensitivity from UV through visible wavelengths (300–700 nm). The quantum efficiency is poor for infrared wavelengths (>800 nm), and as such infrared light is not suitable for use with the image intensifier used. The illumination source should ideally be low cost, small and efficient to allow the range imaging camera to become a small and portable device.

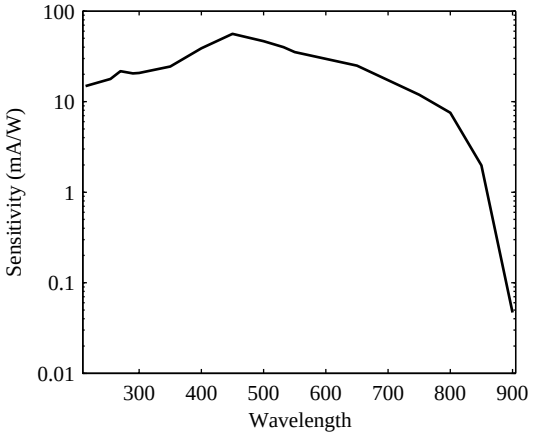

Figure 5.1: Spectral response of the image intensifier utilised (S20 photocathode) (Photek, 2002).

5.1.1 LEDs, Lasers and Optical Modulators

Light Emitting Diodes (LEDs) are an extremely cheap illumination source which are available in a wide selection of wavelengths and can easily be modulated by varying the junction current. Their efficiency is good, although the output power can be relatively low unless an array of devices is used. The existing range imaging system uses a bank of Agilent HLMP EL series LEDs that are sinusoidally modulated up to 20 MHz (Carnegie et al., 2005). To improve the measurement precision, modulation at a frequency equal to that of the image intensifier is required.

From the datasheet, the LEDs used have an exponential time constant τ of 20 ns (Agilent Technologies, 2004). Modulating them with a sinusoidal waveform gives a cutoff frequency calculated as

$$f_{c-\sin} = \frac{\sqrt{3}}{2\pi\tau} = 13.78\,\text{MHz}. \tag{5.1}$$

Equation (5.1) is analogous to that used to calculate the cutoff frequency of a RC voltage divider circuit (where τ = RC), where the cutoff frequency f_c is defined as the frequency where the output power decreases by 50%. To account for the fact that the LED output is power, rather than voltage (which must be squared to attain proportional power), Equation (5.1) is multiplied by an additional factor of $\sqrt{3}$ compared to the RC voltage divider circuit (Schubert, 2006).

Rather than driving the LEDs with a sinusoidal wave, a square wave can instead be used to achieve higher frequency operation. Using the exponential time constant τ of 20 ns and assuming equal rise and fall times, the rise time from 25% to 75% output intensity is calculated as

$$t_{0.25-0.75} = \ln(0.75) - \ln(0.25) \times \tau = 21.97 \text{ ns}. \tag{5.2}$$

The 75% to 25% decay is also 22 ns (since equal rise and fall times are assumed), leading to a cutoff frequency of

$$f_{c-square} = 1/(2 \times 22 \text{ ns}) = 22.76 \text{ MHz}. \tag{5.3}$$

The modulation frequency of the LED circuit can therefore be extended by approximately 65% by using a square wave electrical signal over a sinusoidal waveform, albeit at the expense of producing a non-sinusoidal intensity output. To further extend the operating frequency, it is possible to change the waveform to exhibit a high current at the edge transitions (preemphasis) as shown in Figure 5.2, reducing the LED turn on and turn off times. Dawson achieved approximately 50% higher modulation bandwidth by applying this method to the turn off transition (Dawson, 1980). In practice a current limiting resistor is often used in the circuit; by placing an appropriately valued capacitor in parallel with this resistor the desired waveform can be achieved with little effort.

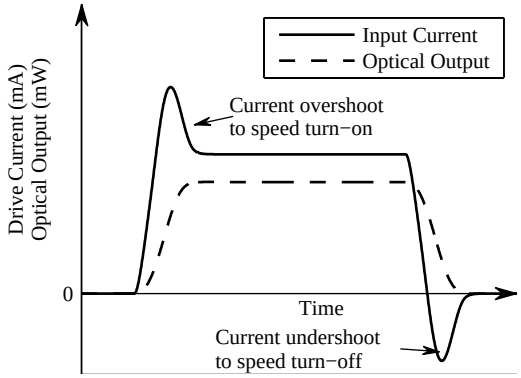

Figure 5.2: Square waveform with high current peaks, used to enhance the cutoff frequency of modulated LEDs.

Although this modification decreases the transition time, the LEDs used are still not capable of being modulated near the required 100 MHz frequency. This is a limitation of LED manufacturing design. There are two fundamental issues constraining LED modulation bandwidth: carrier recombination time and junction capacitance. The carrier

lifetime can be reduced, for example by decreasing the active layer thickness or increasing the doping density, although the gain in speed is offset by a reduction in output power (Lee and Dentai, 1978; Saul et al., 1985). For high intensity LEDs, rather than carrier lifetime, it is likely that the device capacitance limits the modulation frequency. By decreasing the active area of the LED, the capacitance can be reduced, although again at the expense of output power (Elder et al., 2006). The power-bandwidth product appears to be approximately constant (Mynbaev and Scheiner, 2001), hence it is not possible to achieve high output power and modulation bandwidth simultaneously (for a given semiconductor material). Ideally a single high power 3 W LED (e.g. Luxeon III Star, Philips Lumileds Lighting Company, CA, USA), where 3 W is the input electrical power producing approximately 300 mW optical power, could be used as the illumination source, but these LEDs have been demonstrated to have a limited cutoff frequency of 40 MHz (Elder et al., 2006).

Resonant cavity LEDs (RCLED) are capable of producing higher bandwidth than a standard LED, and are primarily designed for use in plastic optical fibre with high bit rates. Table 5.1 shows a few suitable RCLEDs which meet the bandwidth requirement and have an appropriate wavelength. The optical power of an individual device listed in Table 5.1 is insufficient for the illumination source, so an array (containing on the order of 100 LEDs) would be required. A figure of merit is given for each device by multiplying the optical power (in mW) by the image intensifier quantum efficiency (QE) for the given wavelength, and dividing by the unit price where known. Prices shown are only intended as an estimate, and are expected to vary significantly dependent on the quantity ordered.

Table 5.1: Commercially available RCLEDs which have suitable bandwidth and wavelength.

Model	Manufacturer	Wavelength	Optical power	Bandwidth	Approx. Price	Figure of merit P×QE/$
ZL60003	Zarlink	650 nm	2 mW	125 MHz	US$9.50	1.0
1A466	Mitel	650 nm	0.7 mW	200 MHz		
L9534	Hamamatsu	650 nm	0.5 mW*	100 MHz	US$5.34	0.4
XR65-R5P0U	Optowell	650 nm	1.6 mW	155 MHz	US$5.50	1.4
* Maximum power coupled through 1mm diameter plastic fibre						

Utilising a large array of LEDs can cause an error in the measurement due to differing path lengths from the illumination source to the target as illustrated in Figure 5.3. The LED intensity modulation waveforms are superimposed, and for small values of d or large values of x the resultant wavefront becomes distorted leading to a measurement error. The

influence this effect has is dependent on the system geometry, although should be avoided (or at least minimised) where possible.

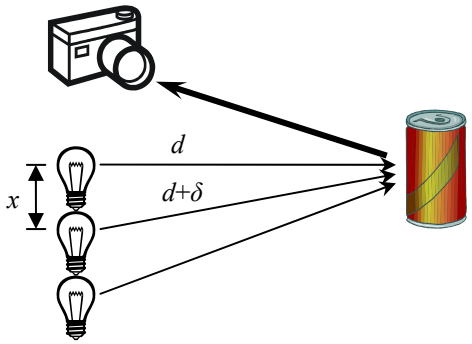

Figure 5.3: Multiple illumination sources in an array can cause an error in the measured range due to the differing path length from each source to the target.

Laser diodes offer much higher intensity and modulation speed than LEDs, simultaneously providing GHz bandwidth and up to hundreds of milliwatts of optical power. As such they are suited to being used in the illumination source design, while also being compact and efficient. Laser diodes do however have a number of disadvantages over LEDs:

- Higher cost per unit.
- Complex drive requirements.
- Thermal sensitivity.
- Interference patterns/speckle.
- Eye safety concerns.

The cost of laser diodes heavily depends on the wavelength required, driven by the current level of technology and applications utilising that wavelength. The latest generation of high definition DVD players (HD DVD and Blu-ray) use laser diodes with a wavelength of 405 nm, which is near the peak sensitivity of the image intensifier; however they are in the early stages of the product lifecycle and therefore are currently expensive. The previous generation of DVD players use a wavelength of 650 nm, and as the technology is mature, high power laser diodes are available at much lower cost. Infrared wavelengths are not appropriate as they produce poor responsivity when used with the image intensifier. Green DPSS (diode pumped solid-state) lasers, despite being near the image intensifier peak responsivity, are also not suitable as they cannot be directly modulated at high frequencies due to the long fluorescent lifetime of the gain material (Hu et al., 2006). A comparison is

shown in Table 5.2, where the figure of merit has again been calculated for various diodes purchased in single unit quantities.

Table 5.2: Comparison of a selection of laser diodes.

Model	Manufacturer	Wavelength	Continuous optical power	Approx. Price	Figure of merit P×QE/$
DL-LS5017†	Sanyo	405 nm	60 mW	US$2400	0.3
HL6501MG	Hitachi	658 nm	35 mW	US$49	3.4
ML120G21	Mitsubishi	658 nm	80 mW	US$68	5.6
HL6535MG†	Hitachi	658 nm	90 mW	US$66	6.5
ML101J27†	Mitsubishi	660 nm	130 mW	US$169	3.7
† Diodes which are now available but were not on the market at the time of initial design					

Comparing Table 5.1 with Table 5.2 it is quickly apparent that the figure of merit calculated is much higher for (red) laser diodes than RCLEDs, - despite having a much higher cost per unit, this is compensated for by the high optical power provided. The quantum efficiency of the image intensifier is significantly higher near the 405 nm wavelength; however the current cost of these laser diodes makes their use uneconomical.

The electronic drive circuit required to operate a laser diode is more complex than that for an LED. Figure 5.4 shows the output power for a laser diode plotted against the input current for two different temperatures. A threshold current is required to begin lasing, after which point the output power is approximately linear with an increase in current. A variation in temperature alters both the threshold point and the efficiency slope; hence the driver circuit must take these effects into consideration.

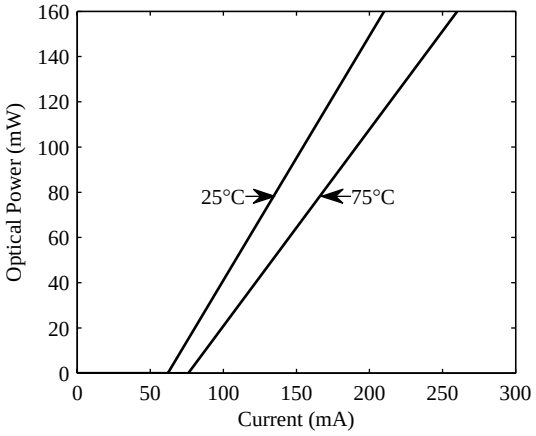

Figure 5.4: Response of ML120G21 laser diode pulsed with a 50% duty cycle at 25 and 75°C (Mitsubishi Electric, 2004).

The illumination from a laser source is coherent, where photons exhibit the same frequency phase and direction. When this coherent source is reflected from a rough surface

(compared to the wavelength of the light), interference causes fluctuations in intensity, producing an image containing a seemingly random pattern of dark and bright areas known as speckle (Briers, 1996). Laser diodes also operate with multiple traverse modes within the lasing cavity, which can result in spatial intensity patterns in the projected illumination (Koechner, 1999). Uniform illumination of the scene is required to acquire an accurate range measurement; hence these effects due to the use of a laser are undesirable. Laser diodes are capable of producing much higher irradiance than LEDs, and therefore eye safety can be an issue. Safety is dependent on the power, duration, wavelength and geometry of the illumination system, and is discussed further in Section 5.1.3.

Optical modulators are an alternate method for varying the intensity of an illumination source. Rather than directly modulating the illumination source itself, the modulator is an external device with a controllable refractive index, where the change in refractive index can be configured to vary the illumination intensity from a continuous wave (CW) source. An acoustic-optic modulator uses a piezoelectric transducer to generate a sound wave through an optical medium. The sound wave generates periodic compression within the medium which, when a light beam enters at an angle, causes Bragg diffraction to occur (Webb and Jones, 2004). The output beam intensity can be altered by varying the amplitude of the sound wave within the medium. Alternately an electro-optic modulator can be used where the original beam passes through a linear polariser, and the modulator rotates the angle of the polarisation plane (Webb and Jones, 2004). If the plane of polarisation of the beam aligns with that of a second polariser placed at the exit of the modulator, the light is able to pass through the system. The rotation of polarisation is controlled by applying an electric field to a crystalline material; hence the illumination intensity can be controlled by varying the applied voltage.

Some optical modulators can be used with high power CW illumination sources, producing an optical output power in the order of many watts. Acoustic and electro-optic modulators have a number of disadvantages however, including requiring precise alignment, high operating voltages, low efficiency, and they are much larger and more expensive than LEDs and laser diodes.

From this brief overview, it is apparent that laser diodes are the optimum illumination source for this project, providing high bandwidth, efficiency and optical power with a relatively low cost.

5.1.2 Illumination Driver Design

Two methods are used to control CW laser diodes, constant current and constant power. A constant current source, as the name implies, is a precise current source which does not change, whereas a constant power driver utilises a feedback loop, monitoring a photodiode to adjust the current through the laser diode to maintain a constant average optical power (Webb and Jones, 2004). A change in temperature alters a laser diode's threshold current and efficiency as illustrated in Figure 5.4; hence the output power of a laser diode operating with a constant current source is temperature dependent. A constant power driver is therefore preferred, which for a modulated system must monitor and control the average power rather than the instantaneous level.

The requirements of the laser driver circuit are:

- Sinusoidal output illumination waveform.
- Constant power control.
- Wideband operation from 1–100 MHz.
- Small size, low cost and high efficiency.

A number of discrete IC laser diode drivers exist, however these are typically designed for either CW operation or high speed switching (fibre optic communication, CD/DVD drives, etc.) and it is difficult to produce a sinusoidal output from these devices.

Generating a sinusoidal intensity waveform from a laser diode requires operating in the linear region of the laser diode response above the threshold current. This requires a DC bias current to maintain the diode current above the threshold, coupled with a modulation current to transverse the linear operating region. To achieve this, a circuit was designed using two separate ICs; the first to control the average power, and the second to inject a modulation waveform as shown in Figure 5.5. The iC-WK is a CW laser diode driver which controls the current through the laser by monitoring an integrated photodiode within the laser diode to achieve a constant output power, with the power level determined by setting a variable resistor. External transistors are used to boost the current capabilities of

the iC-WK to that required by the laser diode used (HL6501MG, Hitachi, Tokyo, Japan), and in this configuration the iC-WK is at a floating potential. A high speed opamp is configured as a Howland current pump (Fraden, 2004), which when driven with a ±1 V input from the DDS produces a ±40 mA output. These two currents are summed to produce the required modulation current.

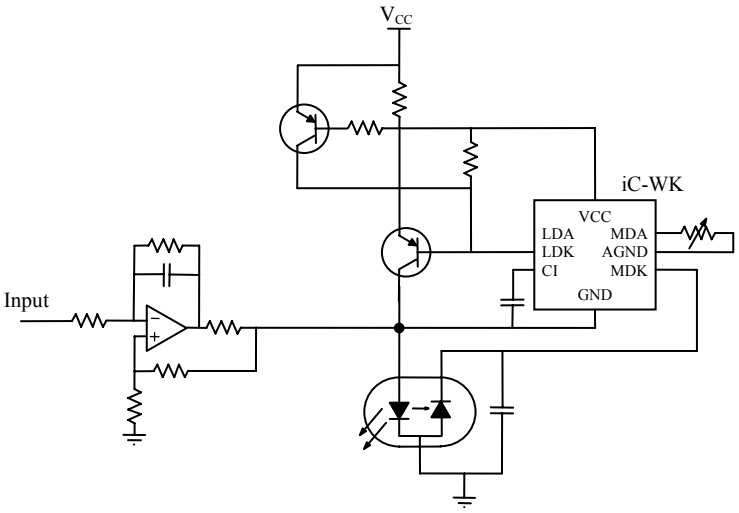

Figure 5.5: Sinusoidal laser modulation circuit.

The circuit suffers from a number of limitations. The laser diode selection is constrained to devices containing an integrated photodiode, as an external monitor diode is impractical for a simple system design. This restricts the output power available from a single laser diode, as high power (>50 mW) visible laser diodes typically do not contain an integrated photodiode. In order to compensate for the loss of output power, a number of laser diodes, and associated control circuits are therefore required.

The modulation index or depth, i.e. the difference in intensity between the high and low peaks, is difficult to control. Ideally the modulation current (during the negative peak of the cycle) will reduce the laser diode current to the threshold value so that the output intensity reduces to zero. Similarly, the positive peak will increase the output intensity to reach the maximum power rating specified by the manufacturer. Operating outside this range will cause a nonlinearity in the output intensity (if the current reduces below the threshold), or damage the laser diode if the maximum power limit is exceeded. As the range measurement precision is directly proportional to the modulation index (refer Section 3.1), it is desirable to operate with a depth as close as possible to 100%. As illustrated in Figure 5.4, a laser diode threshold current and slope efficiency have a

temperature dependency; hence the amplitude of both the DC and AC currents need to be controlled, which the circuit is not capable of doing. Stability problems were also encountered in the constant power control circuit. The feedback loop inherently attempts to compensate for the intentional AC changes in output intensity. To prevent this from occurring, a capacitor is used as a low pass filter in the feedback loop. In the circuit shown, it is possible for the modulation current to cause contamination of the power control feedback signal at the filter capacitor, and, in practice, instabilities lead to oscillation of the output power. Although it is possible to isolate the DC current supplied to the laser diode, for example by using an inductor in series with the source, an alternate circuit was designed to enhance the system precision by using square wave illumination and higher power laser diodes.

The laser diodes used (HL6501MG, Hitachi Ltd., Toyko, Japan) were selected as they offered the highest power available with an integrated photodiode, rated at 35 mW continuous and 50 mW peak output power. For a sinusoidal output, the average output power is constrained by the lesser of the continuous power rating, or half the peak power rating, hence for the HL6501MG the maximum average power is 25 mW. From Table 5.2 it is obvious that alternate higher power diodes are available which have a superior figure of merit, although these devices do not include an integrated photodiode. The Mitsubishi ML120G21 is rated at 80 mW continuous and 160 mW peak output power, allowing modulation with an 80 mW average power – over three times that of the HL6501MG. Driving the laser diode without a monitor photodiode requires the use of a constant current circuit, with the associated detrimental temperature dependency, or the inclusion of an external photodiode. As the system will be operated in a laboratory environment, the ambient temperature can be assumed to be stable and therefore only have a limited influence on a laser diode operating in the constant current mode.

Rather than modulating the laser output with a sinusoidal waveform, it was changed to square wave modulation. Square wave modulation allows the driver circuit to be simplified, requiring a digital switch rather than analogue current control, which also improves electrical efficiency. The inclusion of harmonics in the modulation waveforms, such as is the case for a square waveform, is discussed in Section 3.2.2 where it is shown to cause a linearity error in the phase measurement. It is also shown that this error can be

overcome by taking a large number of samples of the signal to avoid aliasing. Note that the beat signal modulation index can be increased by using square modulation compared to that of sinusoidal modulation, resulting in an increase in measurement precision for identical average and peak illumination power.

This increase in modulation index can best be understood by considering the fact that the ranging system determines the object distance by measuring the phase of the fundamental Fourier component of the beat signal. The sinusoidal waveform consists of only the fundamental Fourier component, illustrated in the left of Figure 5.6 as a solid line. On the right hand side of Figure 5.6, a square wave of the same frequency is shown as a solid line, with identical average and peak power levels. The dashed line drawn over the square wave illustrates the corresponding fundamental Fourier component, which is $4/\pi$ (or approximately 27%) greater than that for the sinusoidal illumination waveform. The increase of the fundamental Fourier component increases the modulation index value of the beat signal, resulting in an improvement in the final range measurement precision. The square waveform is therefore preferred over the equivalent sinusoidal waveform for situations where a large number of samples can be collected, corresponding to an extended acquisition time, avoiding the non-linearity discussed in Section 3.2.2.

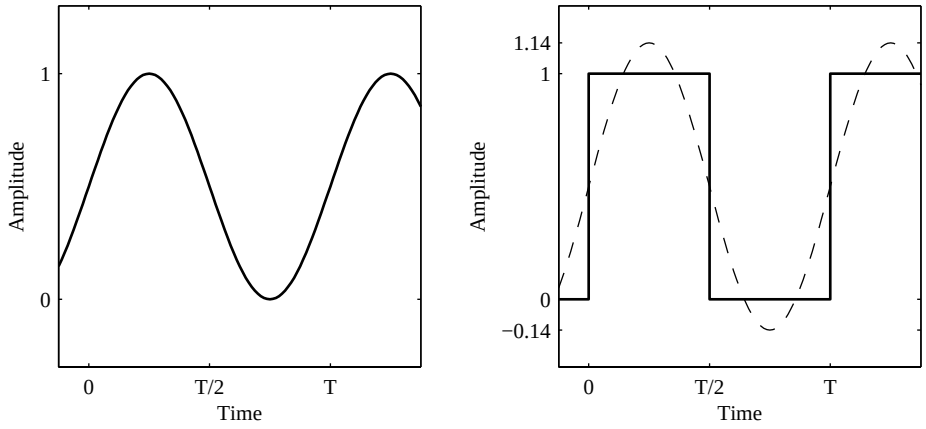

Figure 5.6: Square wave illumination allows the amplitude of the fundamental Fourier component to exceed that of equivalent sine wave illumination, resulting in improved measurement precision despite the use of an identical modulation frequency, peak and average power level.

A revised circuit, designed to generate square wave modulation, is shown in Figure 5.7 and in detail in Appendix A. The sinusoidal input signal from the DDS is converted into a square CMOS level signal through a comparator. The comparator threshold level can be adjusted to vary the duty cycle, and includes feedback to add hysteresis to prevent

unintentional switching due to noise near the threshold level. The CMOS signal is buffered and input to a laser switch (iC-HK, iC-Haus GmbH, Bodenheim, Germany), which acts as a voltage controlled current switch. The IC supports laser modulation frequencies from CW up to 155 MHz, and has two channels. The first channel is permanently enabled (EN1) which provides the DC laser threshold current, while the second channel (EN2) provides the square wave modulation current. The ratio of threshold and modulation currents is set by a resistor for each channel.

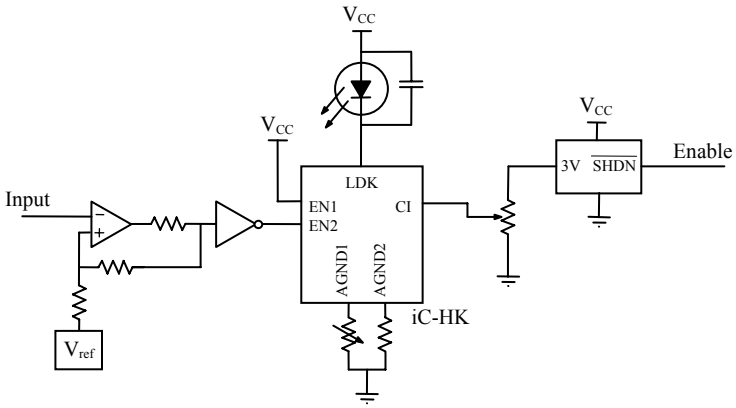

Figure 5.7: Square wave laser modulation circuit.

The laser current is also dependent on an analogue voltage on pin CI. A linear regulator is used to provide this voltage with minimal noise, and provides a simple method to enable and disable the laser output through the shutdown pin on the regulator. A soft-start circuit has been added to the regulator input to prevent any voltage spikes from the main power supply causing a large current to flow through the laser diode on start up, and is shown in the final schematic in Appendix A. The laser diode from the prior circuit is replaced with the higher power ML120G21, as the circuit does not require feedback from a photodiode.

As multiple laser diodes with the same model number have a similar operating current, it is relatively easy to extend the circuit to operate with an array of laser diodes. The input signal conditioning and reference voltage elements of the circuit do not need to be replicated, so each additional laser diode only requires the addition of a laser switch (iC-HK), resistors and capacitors. The final PCB, shown in Figure 5.8, utilises four ML120G21 laser diodes which are all modulated with the same signal. The light from each diode is coupled through a multimode optical fibre, which acts as a mode scrambler to avoid spatial intensity patterns, and produce a more uniform circular intensity profile over the scene than would be achieved directly from the laser diode. Figure 5.9 illustrates the

mounting of each fibre equidistantly around the camera lens, providing (on average) coaxial projection of the light with a lens to alter the divergence angle to match the camera lens configuration. The effects of shadowing are reduced using this arrangement, as light is projected from above/below and left/right of the camera optical axis. Superposition of the illumination from the four laser diodes also reduces the resultant speckle pattern as each laser diode operates at a slightly different wavelength; therefore the coherence of the combined illumination is significantly reduced. The four laser diodes provide a combined average output power rating of 320 mW (640 mW peak) into the optical fibres.

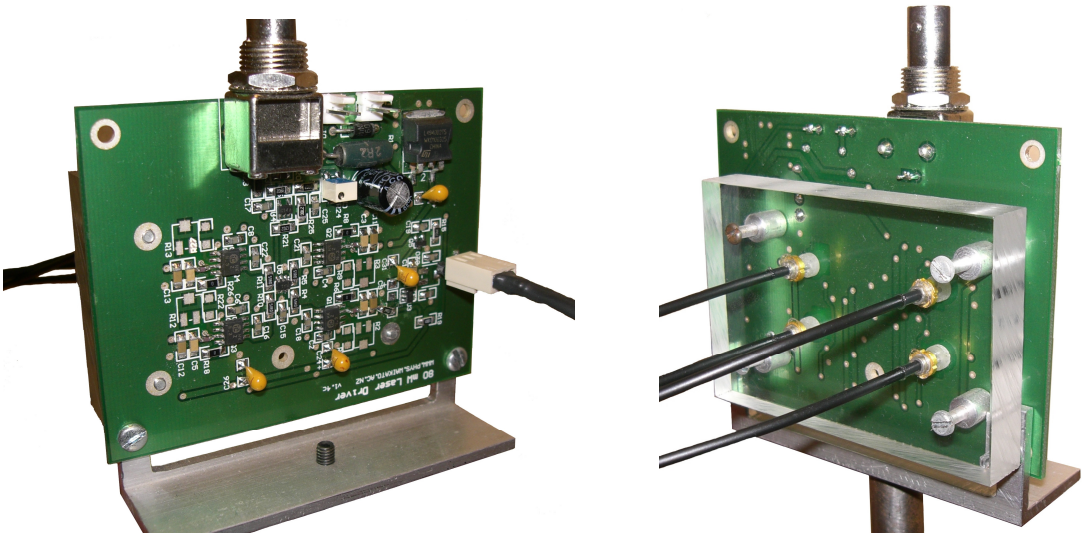

Figure 5.8: Laser driver PCB with four ML120G21 diodes coupled into optical fibres with a combined average output of 320 mW (peak output of 640 mW).

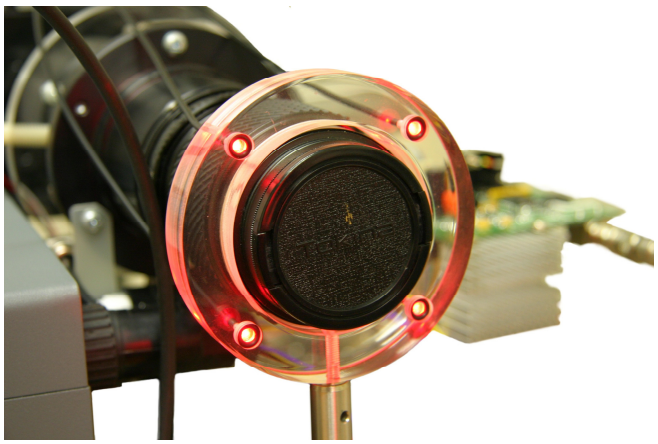

Figure 5.9: The optical head provides equidistant spacing of the four output fibres around the camera lens to generate a coaxial average illumination source.

5.1.3 Eye Safety

Eye safety is a particular concern for any application using lasers, as the energy density can be extremely high in a collimated beam. Ideally the illumination source should use

infrared light with a wavelength exceeding 1500 nm, often considered as an "eye safe" wavelength, as the radiation levels required to cause injury are much higher compared to visible wavelengths (Henderson and Schulmeister, 2004). At these wavelengths the eye is unable to focus the light onto the retina, and the maximum permissible exposure (MPE) is 100 times higher than that of the 658 nm laser diodes used (Standards New Zealand, 2004). The spectral response of the image intensifier decreases sharply for infrared wavelengths above 800 nm as illustrated in Figure 5.1, hence the use of infrared wavelengths is impractical. It is therefore important to calculate the system constraints required to be considered eye safe for the wavelength used. Of particular interest is the nominal ocular hazard distance (NOHD), which is the minimum distance from the diverging laser source to an observer which can be considered eye safe.

The eye safety calculations are based on the New Zealand laser safety standard AS/NZS 2211.1:2004 (Standards New Zealand, 2004). There are two situations of interest; accidental viewing and intentional viewing. As the wavelength being used is visible, accidental viewing is defined as the time required for a blink reflex to occur, and is specified as 0.25 s. Intentional viewing, such as would occur in an application where a human face is being scanned, is defined as an exposure time greater than 10 s. The maximum permissible exposure (MPE) for these two cases is given in Table 5.3, along with the MPE for a single pulse (which is required as the waveform can be considered a series of short pulses).

Table 5.3: Maximum permissible exposure (MPE) at the cornea for direct exposure to laser radiation for selected conditions applicable to the range imaging system (Standards New Zealand, 2004).

Single Pulse ($10^{-9} < t < 1.8\times10^{-5}$ s)	**Accidental Exposure** ($t < 0.25$ s)	**Intentional Exposure** ($t > 10$ s)
$MPE = 5\times10^{-3}\ C_6\ Jm^{-2}$	$MPE = 18\ t^{0.75}\ C_6\ Jm^{-2}$	$MPE = 18\ C_6\ T_2^{-0.25}\ Wm^{-2}$

$C_6 = \alpha/\alpha_{min}$ $\quad \alpha_{min} = 1.5$ mrad $\quad 1.5 < \alpha < 100$ mrad
$T_2 = 10\times10^{[(\alpha-\alpha min)/98.5]}$ s $\quad \lambda = 658$ nm

The NOHD is the minimum distance at which the exposure is below the maximum permissible exposure limit, and can be calculated using Equation (5.4), where φ is the beam divergence angle (defined by the points where the irradiance falls to 1/e of its central peak value), P is the radiant power, and a is the diameter of the beam at the exit of the laser. The geometric variables are illustrated in Figure 5.10.

$$NOHD = \frac{1}{\varphi}\left[\frac{4\times P}{\pi\times MPE}\right]^{0.5} - \frac{a}{\varphi} \tag{5.4}$$

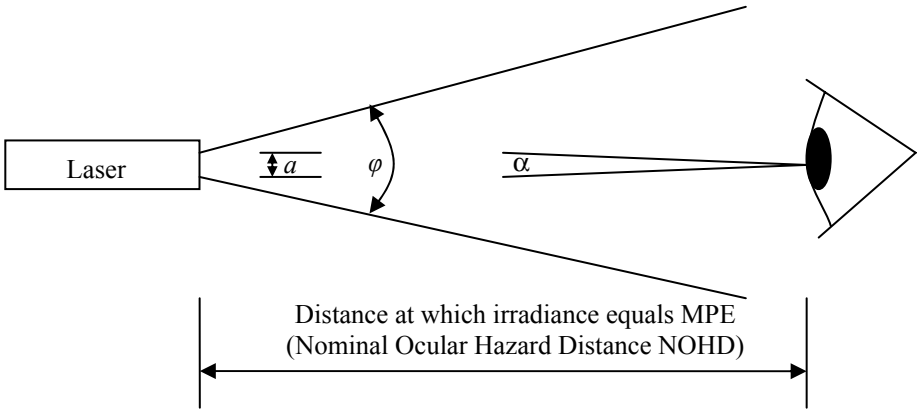

Figure 5.10: Diagram illustrating geometric parameters used to calculate the NOHD.

Appendix B.1 shows that eye safety for this system is constrained by the average optical power rather than the peak power of each pulse. As the light source is created from an array of laser diodes, the MPE must be considered for different groupings of 1, 2, or 4 diodes, with the most restrictive case being used. A MATLAB® program was written using Equation (5.4) to calculate the nominal ocular hazard distance (NOHD) for each grouping, refer Appendix B.2, with results given in Table 5.4 for expected typical operating conditions. The illumination source is considered as four separate sources in a square pattern as shown in Figure 5.9, with a linear separation of 80 mm. The calculated distance, even for intentional exposure, is less than the minimum distance of 100 mm required by the standard to calculate the angular subtense.

At large viewing distances, the eye is capable of focusing a small-area divergent source onto a small spot on the retina. As the distance is decreased, the optical power collected by the eye increases, while the spot size on the retina does not change significantly due to the eye maintaining focus on the source. This implies that the source is more hazardous for smaller distances. However, for very short distances (below 100 mm) the eye is no longer able to focus the source into a small spot, instead distributing the energy over a larger area of the retina (Henderson, 1997). Despite the fact that the total radiant power collected by the eye continues to increase for shorter distances, the power required to cause damage to the retina also increases. It is therefore unlikely that direct intrabeam viewing will be hazardous for the range imaging system using the conditions listed.

Table 5.4: Calculated Nominal Ocular Hazard Distance for the illumination system developed.

Accidental Exposure (t < 0.25 s)	Intentional Exposure (t > 10 s)
Calculated NOHD: 35 mm†	Calculated NOHD: 87 mm†

P_{av} = 50 mW (per source) φ = 15°
λ = 658 nm separation between sources = 80 mm
† From Section 3.7 of the laser safety standard (Standards New Zealand, 2004), the angular subtense (α) must be determined at a viewing distance of not less than 100 mm

It is worth noting that as the optical system design is flexible, the values shown are only for one specific case. For example, if the beam divergence angle φ is decreased then the hazard posed to an observer will increase. Similarly if the configuration were changed so that the same radiant power was generated by a single laser diode (instead of the four laser diode 'extended source'), then the calculated NOHD for intentional exposure is increased to 178 mm. Multiple laser sources therefore not only reduce shadowing and interference problems such as speckle, but also reduce eye safety constraints. Although the current configuration within the laboratory environment is unlikely to pose a hazard to an observer, for a commercial product the NOHD calculations would need to be repeated for the particular system parameters, and a protective housing may be required to prevent an observer from being exposed to radiation exceeding the MPE.

5.2 Co-processor

The range imaging system was constructed with the hardware presented, and although functional, a number of possible improvements were identified. A field programmable gate array (FPGA) was chosen to act as a co-processor to the microprocessor (responsible for controlling the DDS ICs) to implement the improvements, including:

- Reducing jitter in the camera timing signal.
- Adding a camera shutter signal to remove image smear.
- Control of the illumination (laser) intensity.
- Control of the image intensifier gain.

5.2.1 Reducing Camera Trigger Jitter

Noise on the low frequency (camera trigger) DDS output is likely to create jitter, potentially generating an error in the final range measurement. The DDS output at the camera frame rate (typically ~30 Hz) is converted to a square wave through an onboard

comparator. The comparator combats noise by using hysteresis to avoid multiple oscillations at each transition; however an uncertainty exists in the exact time the comparator transition occurs in the presence of noise. The phase difference between the illumination source and the image intensifier modulation signals is constantly changing; therefore the average phase captured in each frame is dependent on the integration start and end times. Under certain conditions, such as acquisition of a small number of image frames, jitter can potentially lead to an error in the recorded pixel intensity values and hence a corresponding error in the final range measurement. Figure 5.11 shows how noise leads to an uncertainty in the exact timing of each transition, and that the uncertainty is related to the frequency of the input signal for a given SNR.

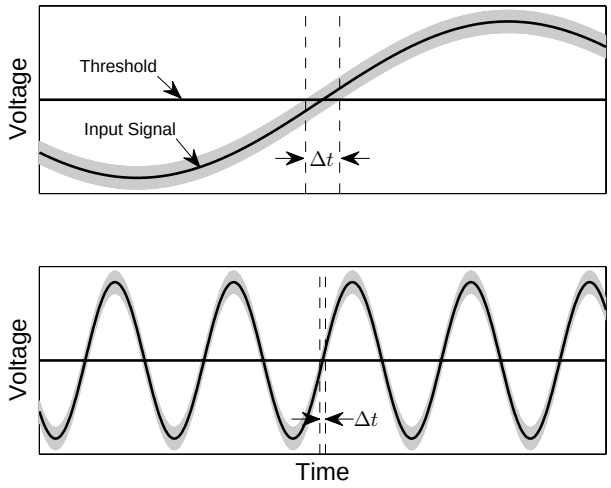

Figure 5.11: A comparator converts the sinusoidal input into a square wave output by comparing the input with a fixed threshold voltage. Noise on the input signal (shown as a grey shaded area) creates uncertainty in the time the transition occurs (Δt), leading to jitter in the output signal. By increasing the input frequency (bottom) the uncertainty decreases proportionally.

The proposed solution is to operate the DDS at a higher multiple of the desired frequency, convert the signal to a digital square wave with minimal jitter, then use a digital counter to divide the output frequency back down to the desired frequency. The resultant jitter will be considerably lower as the digital stages do not contribute any significant error to the signal. To ensure that the digital output has the correct phase relationship with the high frequency signals, at start up, all three DDS signals must be stopped and the counter reset before simultaneously restarting all three DDS signals.

Previously the microcontroller was configured to divide the camera trigger signal to a lower frequency (refer Section 4.2.3) by generating an interrupt to increment a counter, and toggle an output when the counter reaches a certain value. The time required for this

process to complete can fluctuate slightly given that the microcontroller operates from an independent clock. However, precise timing of the output signal was not previously important as it was only required to tag a frame in the recorded sequence.

Rather than using the microcontroller, a FPGA was selected as it can operate asynchronously and therefore generate minimal jitter when dividing the DDS signal frequency, while simultaneously performing a number of other tasks.

5.2.2 Image Smear

To extract an image from a standard CCD device, the sensor shifts the image down into a readout area. While the image is being shifted out of the sensor, it continues to collect light, resulting in a smearing effect down the image. Figure 5.12 illustrates this effect on an artificial image, where the smearing has been exaggerated for clarity (equivalent to shifting the image out at a slow rate).

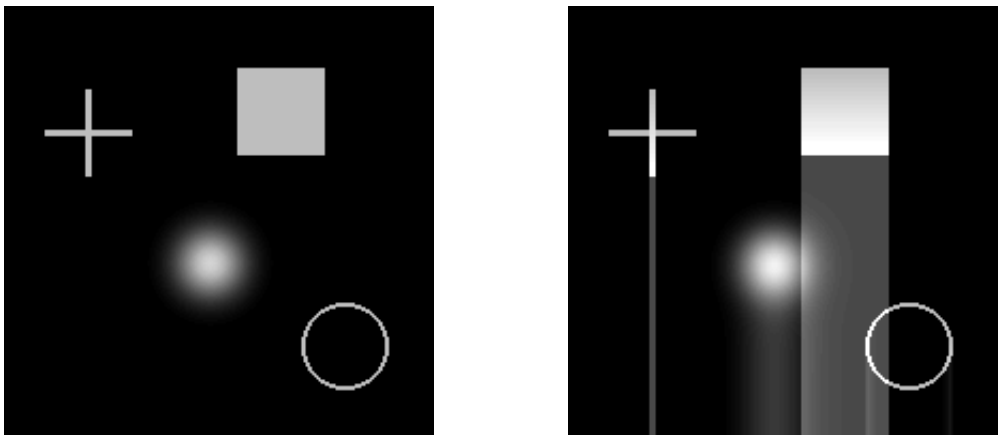

Figure 5.12: Left: An artificial scene; Right: The same scene where the imaged has been smeared due to the CCD sensor still collecting light during the frame transfer period. The smear has been exaggerated to clearly show the effect.

The level of image smear is proportional to the time required to move the image out of the sensor, which is limited by the speed of the readout electronics. To reduce smearing, a frame transfer CCD is used (Pantera 1M60, Dalsa, ON, Canada). The CCD sensor consists of two sections; the first active area captures the image, and the second storage area, hidden from incoming light, stores it during the slow readout time. The time required to shift the image from the active area to the storage area is significantly shorter than the time required by the readout electronics, therefore reduces the image smear. It also allows integration of a new image to begin while the previous image is being read out.

While the smear effect is not noticeable for most standard imaging applications, it still exists using the frame transfer CCD, and causes contamination of the range measurement (Dorrington et al., 2005a). The intensity of the smear can be predicted by the ratio of the integration time and transfer time, both of which are known, allowing post processing software to correct the image, however quantisation and rounding errors limit the effectiveness of this technique (Dorrington et al., 2005a). To completely remove the effect, an alternate camera technology could be implemented (CMOS or interline transfer CCD), or a shutter can be used to block the incoming light during the image transfer stage. The image intensifier is already being used as a high speed shutter; hence the preferred solution is to disable the image intensifier during the image readout period.

The image intensifier driver has the capability to turn off the image intensifier by biasing the photocathode with a positive voltage, as discussed in Section 4.3.2; hence with the addition of an appropriate timing signal, the image intensifier can be disabled during the sensor readout period. The signal can be derived from the same FPGA that generates the camera trigger for the start of each image frame (refer Section 5.2.1), enabling the image intensifier driver at the beginning of the frame integration, and from an internal timer, disabling the image intensifier at the end of the exposure time as shown in Figure 5.13. The image frame is then read out while the next frame is being exposed.

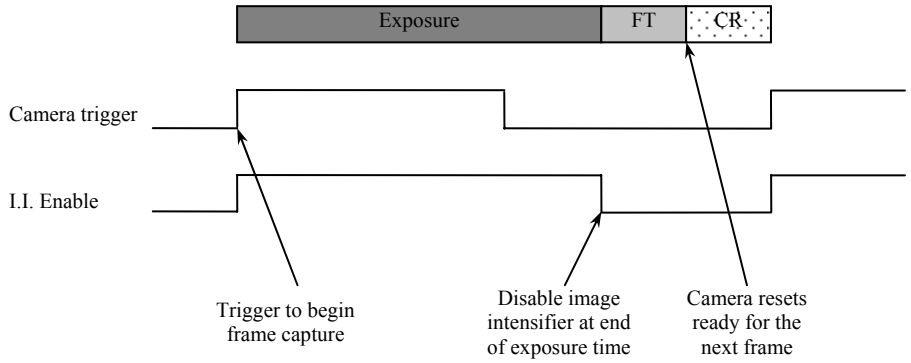

Figure 5.13: Each image frame is broken into three segments; exposure, frame transfer (FT) and charge reset (CR). The image intensifier is only enabled during the exposure time.

In practice it is preferable to add a small fixed delay after the camera trigger signal before enabling the image intensifier. This prevents any propagation delays in the camera trigger signal, which passes through a PC CameraLink interface before reaching the camera itself, from generating an error in the range measurements. Similarly, the image intensifier should be disabled slightly before the end of the camera integration period to ensure the

exposure time is precisely controlled for each frame and image smear is avoided. Potentially the laser driver could also be disabled during the frame transfer and charge reset periods to provide a small reduction in power consumption and reduce eye-safety constraints, although the improvement in either case would be minimal and would likely result in an irritating flicker to a human observer.

5.2.3 Illumination Intensity Control

As the laser diodes have a temperature dependency (refer Section 5.1.2), and they are not temperature controlled, they must reach an equilibrium temperature before generating a stable intensity output. Equilibrium is reached when the heat generated within the laser diode equals the heat dissipated into the surrounding environment, and typically takes tens of seconds from when the circuit is first started. The temperature itself is not an issue as the laser remains well within its rated operating range (−10 to 75 °C); however the laser diode is more efficient at the initial room temperature than at the higher equilibrium temperature. This causes the output intensity to peak when the laser diode is first turned on, followed by decay as the diode temperature reaches steady state. If the laser is set to operate at its maximum rated output intensity during the steady state, then this limit will be exceeded during start up. Alternatively, if the output power remains within the maximum ratings during the peak, the steady state illumination power will be significantly below the capability of the diode. Neither option is ideal; instead a soft start circuit is preferred to slowly increase the laser diode current until the normal operating temperature is reached. A digital to analogue converter (DAC) was added to the FPGA, allowing precise control of the laser diode current. A profile was generated in the data acquisition software to slowly increase the laser diode current on start up over a period of tens of seconds to operate the laser near its maximum rating during steady state, without exceeding the device limits at the initial lower temperature.

5.2.4 Image Intensifier Gain Control

High speed modulation of the image intensifier gain is achieved by modulating the photocathode voltage as discussed in Section 4.3; however the optical gain is also dependent on the microchannel plate (MCP) voltage. A DC voltage is applied to the MCP, manually controlled over a typical range of 400–700 V, to set an appropriate gain. Increasing the gain of the image intensifier by raising the MCP voltage generally results in

higher noise generation within the MCP amplification stage; although in the limiting case where photon noise is dominant, the image intensifier SNR is for the most part independent of gain (Frenkel et al., 1997). Hence high gain is preferable and has the additional benefit of being able to utilise the full dynamic range of the CCD.

Quantisation error in the recorded intensity image can be significant for low level signals. Assuming the signal intensity measured by the CCD covers 3 bits (8 discrete levels), the quantisation error in each intensity measurement can be up to 6.25% (half of one bit). In comparison, if the signal spans 8 bits (256 discrete levels) the maximum error due to quantisation reduces to 0.20%. By maximising the intensity of the signal reaching the CCD, i.e. increasing the image intensifier gain, this source of error can be avoided. This effect has been demonstrated by evaluating the range measurement precision of objects at various distances, comparing a fixed image intensifier gain value to an optimally selected gain for each target (Dorrington et al., 2007b), with a significant improvement shown in the latter case. The dynamic range of a CCD is fairly limited however; therefore saturation limits the gain achievable for a given scene (the image intensifier is also capable of saturating at high output intensities). Selection of the optimal gain value is therefore scene dependent in order to maximise the output level while avoiding saturation. The MCP voltage is controlled by a second DAC in the FPGA design, allowing a feedback loop to operate where the intensity from a given scene is assessed by the processing software on the PC, and the image intensifier gain adjusted accordingly.

In the situation where two objects in a single scene have a large variation in their reflected intensity, a single high precision range measurement of each object is difficult to achieve. In this case it could be possible to perform two measurements, one with a low gain and the other high gain (intentionally saturating part of the image), with the two measurements used to determine the range to the bright and dull objects respectively.

5.3 System Overview

An overview of the complete electronic system configuration is shown in Figure 5.14, and a photo is shown in Figure 5.15. The PC software is responsible for configuring and controlling each component within the system, through either the serial interface to the microcontroller, or the CameraLink interface to the CCD, allowing automation of the range

data acquisition process. The flexibility afforded by the software controlled system also simplifies experimental investigation of the effect of varying key parameters, such as the modulation frequency or CCD frame rate.

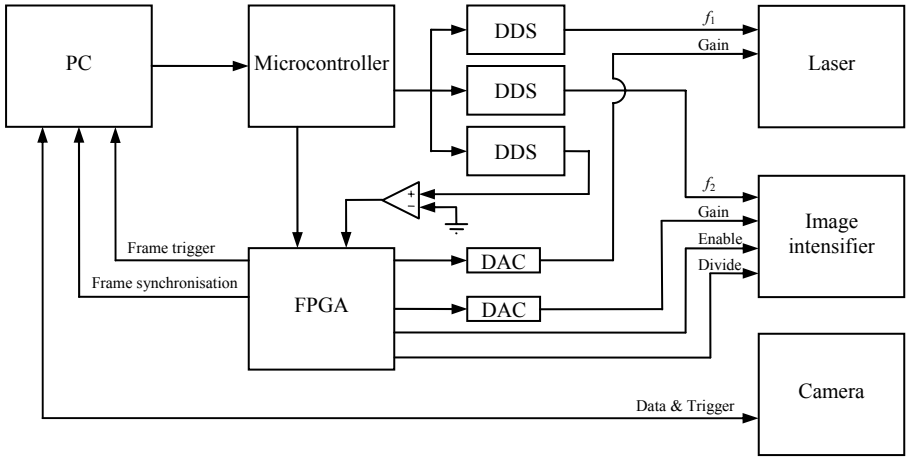

Figure 5.14: Overview of the electronic system showing interconnections between each component in the system.

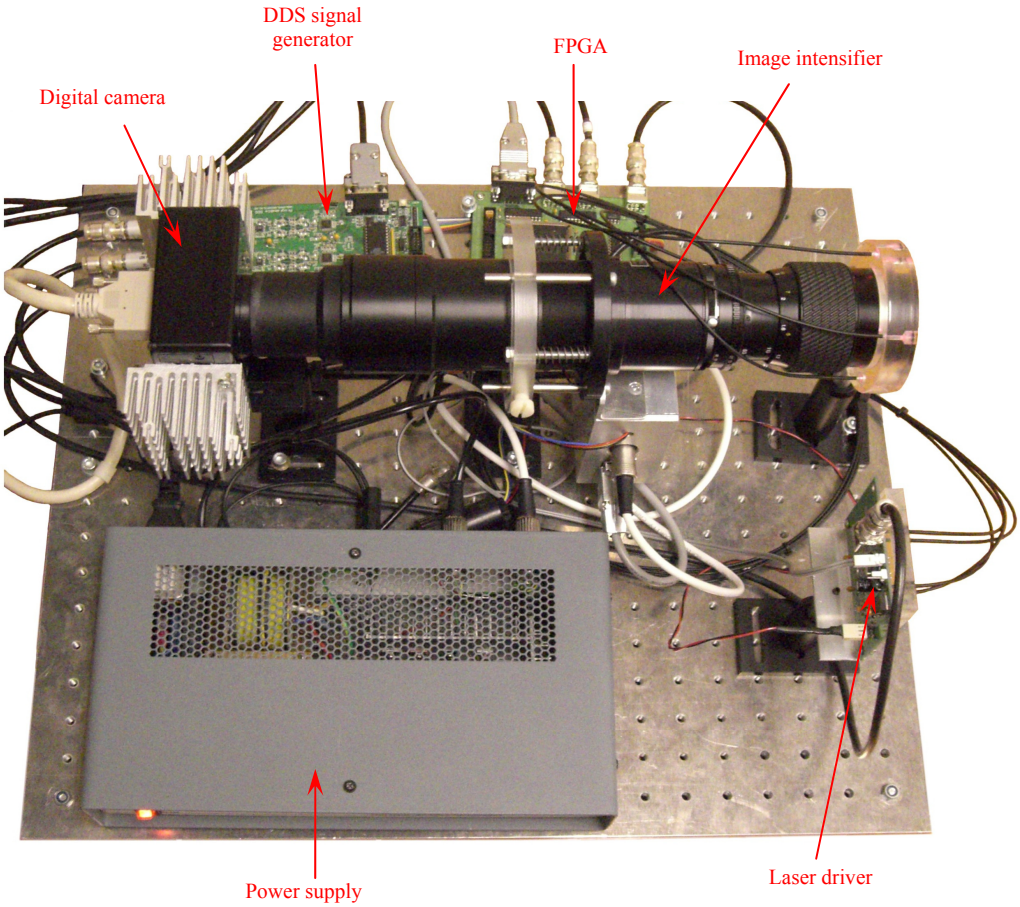

Figure 5.15: The constructed range imaging system.

The responsibilities for each component are individually described in the preceding chapters, but are summarised again here:

- PC
    - Control of all components in the system, either through a serial interface to the microcontroller, or through a CameraLink interface to the CCD.
    - Receive a frame trigger signal from the FPGA, which is relayed to the camera to begin integration of each frame.
    - Receive a frame synchronisation signal, used as a phase reference to mark a single frame in the image sequence to allow absolute range measurements to be made.
    - Record and process data from the CCD to extract range and intensity information.
- Microcontroller
    - Configure each DDS with parameters received from the PC.
    - Relay instructions to the FPGA.
- DDS
    - Generate a configurable output frequency, where the three DDS outputs are frequency locked with a known phase relationship.
- FPGA
    - Divide the third DDS output to generate the frame trigger and frame synchronisation signals with minimal jitter.
    - Enable the image intensifier only during the camera exposure time to remove image smear.
    - Control the illumination intensity, allowing the lasers to perform a soft start to avoid exceeding the maximum rated output power at low temperature.
    - Control the image intensifier gain, allowing the camera's entire dynamic range can be used without saturation.
    - Set the image intensifier divide ratio to reduce power dissipation at high modulation frequencies.

Operation of the ranging system consists of four stages: start up, preview, auto-gaining, and acquisition. During the start up phase the laser power is slowly increased to allow the temperature of the laser diodes to reach steady state, and is only performed once during

each session. The preview mode is used to view the scene through the imaging system before capture commences, allowing the camera to be aligned with the scene and the imaging lens in front of the image intensifier to be manually adjusted (including zooming and focusing). During this process the image intensifier is operated at a low (~1 MHz) frequency to minimise heat generated by the driver circuit, allowing extended periods of thermally safe operation. During auto-gaining, the image intensifier is modulated at a high frequency and the PC software adjusts the image intensifier gain to maximise the intensity at the CCD for the given scene. The final stage acquires the image sequence, allowing the PC to calculate the intensity and range data for each CCD pixel, either in real time, or by saving the data to disk for off-line processing.

6. System Characterisation

Characterisation is required to validate operation and performance characteristics of each component within the system, to identify undesirable traits which may constrain the final measurement precision, and to remove or reduce these traits and optimise the system performance where possible. Ideally the signals from the DDS should demonstrate minimal jitter, and the laser and image intensifier should both provide a high modulation index value. Evaluation of the range measurement precision of the complete system is also important, providing verification between experimental results and predicted outcomes for various parameters, such as changes in modulation frequency and the intensity of the received signal.

6.1 DDS Characterisation

One important performance attribute of the constructed DDS is the output spectral purity. Non-harmonic spurious frequency components contribute jitter during the conversion from a sine wave to a square wave, with the amplitude given by Equation (6.1), where $JitterUI_{pp}$ is the peak-to-peak jitter per unit interval (one output clock period), and dBc_n is the magnitude of the n^{th} non-harmonic spurious frequency component relative to the fundamental component in dB (Brandon and Gentile, 2006).

$$JitterUI_{pp} = \frac{1}{\pi}\arctan\left(\sum_{n=1}^{N} 10^{\frac{dBc_n}{20}}\right) \quad \textbf{(6.1)}$$

Spurious frequency components can originate from higher replicas of the output frequency at $nf_{clock} \pm f_{out}$ as previously illustrated in Figure 4.5, from harmonics of the output signal, and also from the oscillator and other digital clocks in the design. Ideally the analogue low pass filters in the output stage of the PCB design attenuate these spurious frequency components so that they are not present in the output signals, thereby preventing jitter from occurring in the laser and image intensifier modulation waveforms. An HP8560A (Agilent Technologies Inc., CA, USA) spectrum analyser is used to measure the spectral content of the DDS output signal, illustrated in Figure 6.1, sweeping from DC to 500 MHz with a generated output frequency of 55 MHz. The 50 Ω input impedance of the spectrum analyser ensures that the DDS signal is terminated correctly. The origin of each spectral component is identified, where f_{out} is the output frequency, f_{osc} is the oscillator frequency (20 MHz), f_{sync} is the synchronisation clock frequency between each DDS IC (100 MHz), and f_{clk} is the DDS system clock frequency (400 MHz). Components also exist outside the frequency span shown in Figure 6.1 at 745 and 855 MHz ($2f_{clock} \pm f_{out}$), with amplitudes of −45 and −52 dBm respectively. From Figure 6.1, the output spurious free dynamic range (SFDR) can be seen to be 34.5 dBc.

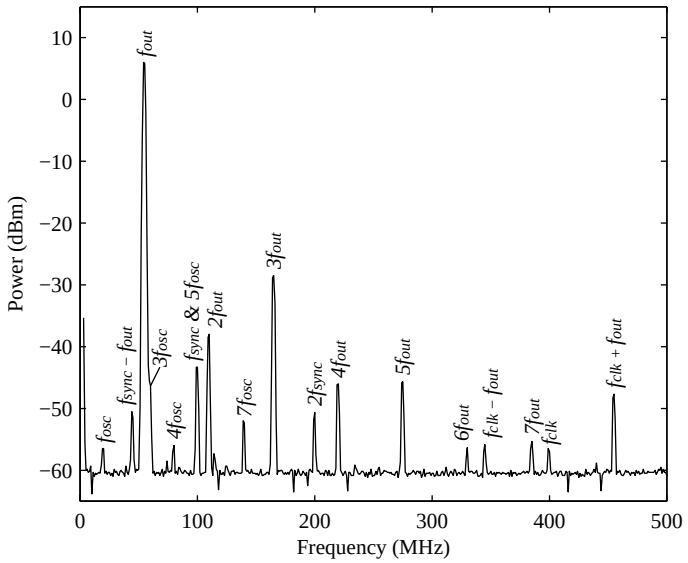

Figure 6.1: DDS spectral response for 55 MHz output.

The harmonic components ($2f_{out}$, $3f_{out}$,…) in Figure 6.1 are due to non-ideal matching of the low pass filter electronic components and PCB layout of the complementary signals, and any DDS DAC or amplifier non-linearity (e.g. due to the limited slew rate of the amplifier output). Harmonic components distort the ideal DDS sine wave output, where the distortion is identical during each period of the waveform, creating a constant phase

shift of the generated square wave. As there is no cycle-by-cycle timing variation due to these harmonic components, they are excluded from Equation (6.1), and can therefore be ignored during analysis.

The remaining spurious components do contribute jitter in the square wave, hence should ideally be attenuated as much as possible. The spectral purity is limited by the performance of the low pass filter and PCB layout; in particular due to the digital oscillator clock harmonics being coupled to the output (nf_{osc}), and limited attenuation of the high frequency replicas of the output signal ($nf_{clock} \pm f_{out}$). Rather than redesigning the PCB, a second (inline) low pass filter (BLP series, Minicircuits, NY, USA) was added to each output after the amplifier. Selection of an appropriate cutoff frequency is dependent on the maximum modulation frequency being used by the system, and the filters can easily be substituted as required. The resultant spectral response is shown in Figure 6.2 for the same 55 MHz output with the addition of a BLP-70 low pass filter (with a cutoff frequency of 67 MHz), where the higher frequency components are no longer visible in the output signal and the SFDR has increased by 16 dB to approximately 51 dBc.

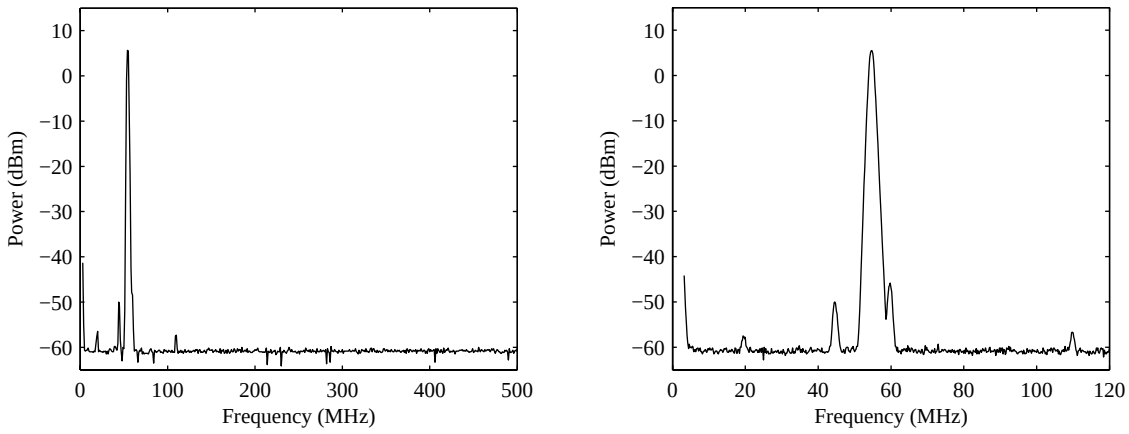

Figure 6.2: Improved DDS spectral response for a 55 MHz output with inclusion of a BLP-70 low pass filter. Left: Spectrum from DC to 500 MHz; Right: Spectrum from DC to 120 MHz.

Another potential source of spectral contamination is crosstalk between the two high frequency outputs. Figure 6.3 shows the spectral output where the second frequency is operating with a 50 kHz difference to that of the signal being measured (i.e. outputs at 54.95 and 55 MHz). Crosstalk measurements with single hertz differences between the two output frequencies is difficult due to the limited spectral resolution of the spectrum analyser, hence a 50 kHz difference was selected to represent the likely crosstalk seen at

small frequency differences. The crosstalk from the second output in Figure 6.3 is approximately −50 dBc.

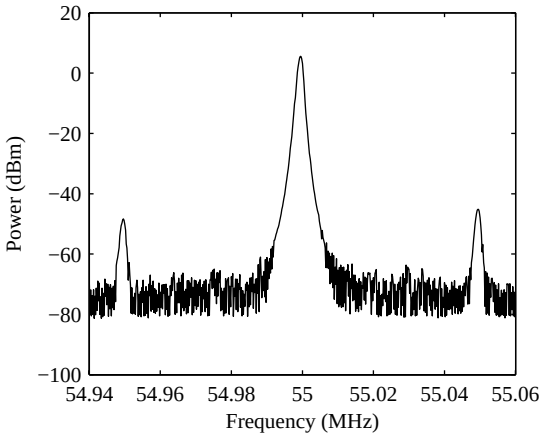

Figure 6.3: Spectral response showing crosstalk between the two high frequency outputs.

From Equation (6.1), the theoretical peak to peak jitter due to the (non-harmonic) spurious spectral components before addition of the external filter was approximately 0.0076 times the period of the output signal, or 140 ps peak to peak. The addition of the external filter significantly reduces this value to approximately 60 ps, although these figures do not include the effect of random jitter due to noise, e.g. within the comparator itself. The significance of jitter towards range measurement precision is discussed further in Section 7.1, where it is shown that small values such as those measured here contribute negligible range measurement error, and in some instances can be beneficial.

6.2 Laser Modulation Characterisation

The illumination waveform from the laser should maximise the modulation index value to achieve high measurement precision, while providing operation up to 100 MHz. By maintaining a laser diode modulation current above the threshold level, the optical output will closely track the input current waveform, with modulation frequencies over 20 GHz achievable (Suhara, 2004). The threshold current is temperature dependent as illustrated in Figure 5.4, so, in practice, the threshold current level must be set above the minimum value to ensure that the diode continues to lase at higher operating temperatures. This reduces the modulation depth as the laser cannot turn off completely, and therefore reduces the measurement precision (refer Section 3.1). A thermal electric cooler, e.g. a Peltier cooler, could be used to maintain a constant temperature, although this adds complexity and cost

to the design, as do other methods of compensating for changes in temperature (such as monitoring the optical output of each laser diode with an external photodiode).

For this application, the maximum modulation frequency required is limited by the bandwidth of the image intensifier and driver; hence there is little benefit in being able to modulate the laser diode at frequencies above 100 MHz. With this in mind, it is possible to operate the laser diode with the bias current set below the threshold current, despite the fact that the laser diode frequency response will be reduced. The advantage to this is that: i) the laser output turns off completely, maximising the modulation depth and ii) during the "off" state the current can be reduced significantly, minimising the heat generated within the laser diode to allow operation at a reduced temperature.

There are however, a number of associated issues with operating the laser diode below the threshold current. A response similar to that illustrated in Figure 6.4 is expected over the required operating frequency range of 10−100 MHz, where a delay can be seen between the input current rising edge and the optical output rising edge, while the falling edges occur (almost) simultaneously. The turn on delay is caused by the time required for the carrier density to reach the lasing threshold (Suhara, 2004). Using the laser driver circuit from Figure 5.8, the delay between the electrical and optical rise times was measured to be approximately 2 ns. The on-time of the output illumination waveform is therefore reduced, with the most significant influence at higher modulation frequencies (e.g. at 100 MHz, a 2 ns delay changes the duty cycle by 20%). Partial compensation for this effect can be made by increasing the duty cycle of the input current waveform; however in Section 7.6 it is shown that a reduction of the illumination duty cycle, whilst maintaining constant average output power, can actually be beneficial for the image ranger application.

From Figure 6.4, it is also apparent that a large overshoot occurs on the rising optical power transition, followed by a "relaxation oscillation" before reaching a steady state value (Suhara, 2004). The oscillation is due to the abrupt change in the injection current, which increases the carrier (electron) density, bringing the number of carriers above the lasing threshold. The photon density increases, although the rate is limited by the photon lifetime (Suhara, 2004). The increase in photon density is accompanied by a reduction of the carrier density, which in turn reduces the photon density and again increases the carrier

density, leading to an oscillation. Due to damping within the system, the output eventually reaches a steady state value after a period of a few nanoseconds.

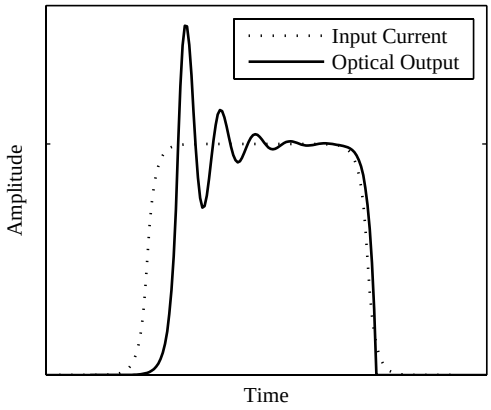

Figure 6.4: Anticipated optical output from a laser diode operating below the threshold current level while being modulated with a frequency of 10–100 MHz. The turn on transition is delayed, and experiences a relaxation oscillation.

Laser diodes are constrained by both a maximum average power where failure occurs due to the diode overheating, and also a maximum peak power where failure occurs due to catastrophic optical damage (COD), where localised heating melts the facet. The relaxation oscillation generates a peak much higher than the steady state intensity; therefore the output power must be reduced to maintain operation within the peak power rating. To overcome this limitation, the oscillation can be damped by reducing the transition speed of the input current, at the expense of slightly reducing the modulation bandwidth. Figure 6.5 shows the measured relaxation oscillation of a laser diode modulated at 15 MHz. The peak amplitude shown is likely to be limited by a combination of the 1 ns rise time of the photodiode (Thorlabs DET210) and 1.7 ns rise time of the oscilloscope (Fluke ScopeMeter 199C); however the oscillation is clearly visible. Figure 6.5 also shows the response where a 300 pF capacitor has been placed in parallel with the laser diode to reduce the current transition speed, critically damping the oscillation so that it is no longer visible. This capacitor is shown in the laser driver schematic in Figure 5.7, and was determined experimentally to achieve adequate damping with minimal reduction to the modulation bandwidth of the circuit.

Illumination waveforms at 10 and 100 MHz respectively are shown in Figure 6.6, illustrating operation of the laser output spanning the required modulation frequency range. The modulation depth and harmonic content are largest at low frequencies where the waveform is square, with a reduction in harmonic amplitudes occurring at higher

frequencies. The modulation index remains high over the operating frequency range as required, with the output approaching zero during each cycle period.

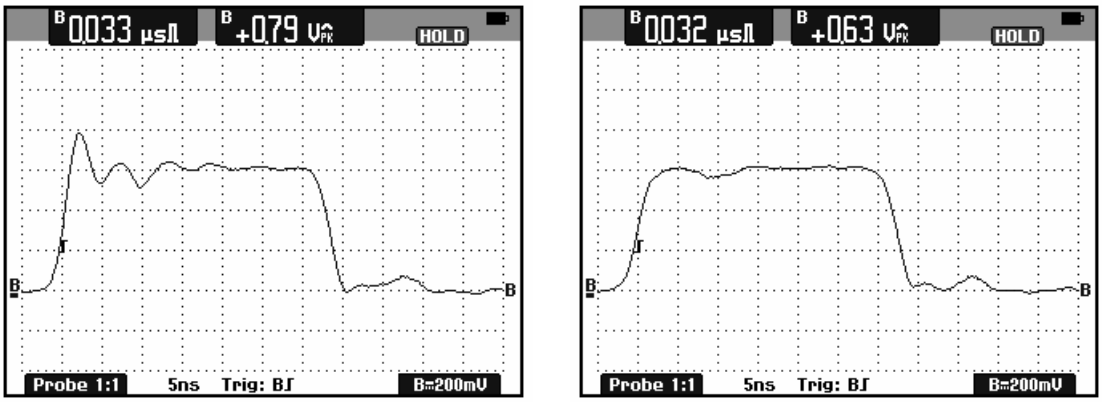

Figure 6.5: Laser diode modulated without a bias current. Left: The relaxation oscillation can clearly be identified on the rising transition. Right: The oscillation has been damped by placing a small capacitor in parallel with the laser diode to reduce the current rise time.

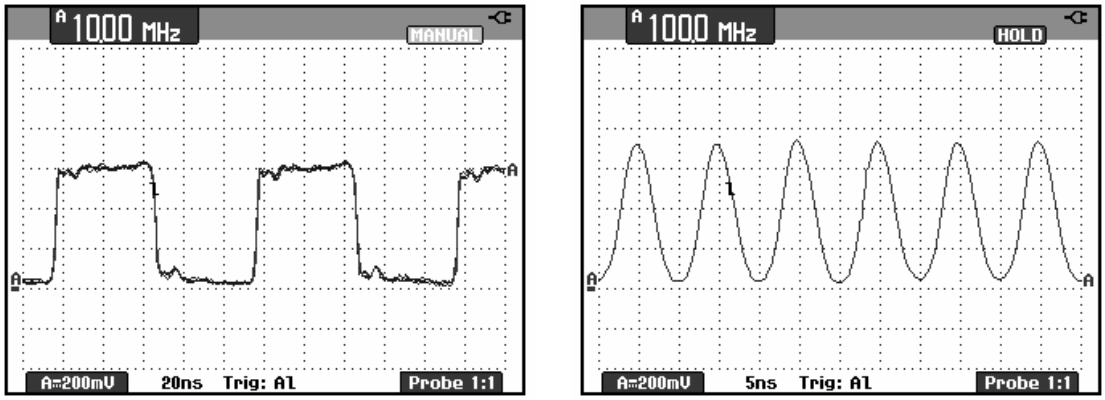

Figure 6.6: Illumination modulation waveforms at 10 MHz (left) and 100 MHz (right).

6.3 Image Intensifier Modulation Characterisation

Characterisation of the image intensifier is a complex task due to the inclusion of both temporal modulation, and spatial variations such as irising, both of which affect the range measurement result. As with the laser modulation, it is important to provide a high modulation index to maximise the measurement precision.

6.3.1 Inferred from DC Response

The measured image intensifier DC response is illustrated Figure 4.14, where a uniform light source was applied to the image intensifier input window, and the output intensity measured for various photocathode voltages. By combining this response with the measured photocathode modulation voltage from Figure 4.20, the high frequency optical

modulation waveform can be determined. This approach, while simple, has some significant shortcomings. The electrical waveform measured used a dummy load, rather than being connected to the image intensifier, and therefore may deviate from that shown during actual operation. Simply measuring the voltage with the image intensifier connected is also a non-trivial task; the capacitance of a typical oscilloscope probe is approximately 12 pF, which when placed in parallel with the image intensifier photocathode capacitance of approximately 60 pF will contribute significant load to the RF amplifier, altering the electrical modulation waveform. It is also not possible to measure any spatial variation in the gain modulation across the image intensifier surface, e.g. due to irising.

6.3.2 Optical Temporal Probe

A preferred alternative is to characterise the image intensifier optically. This can be achieved by modifying the configuration of the ranging system shown in Figure 2.15, replacing the square (or sinusoidal) modulation illumination source with a picosecond pulsed laser as illustrated in Figure 6.7 (Payne et al., 2008a) to determine the optical modulation response.

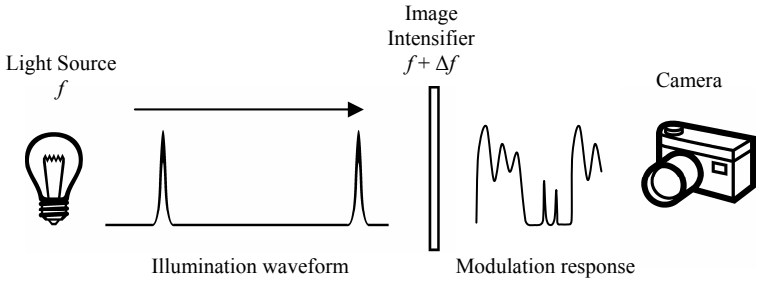

Figure 6.7: System configuration to characterise the dynamic image intensifier gain.

The narrow illumination pulse temporally probes the image intensifier gain at one instant during each modulation period as illustrated in Figure 6.8. As with the ranging configuration, the image intensifier output is integrated by the CCD over a large number of pulses to provide high SNR for each captured frame. The frequency difference Δf between the light source repetition rate and the image intensifier modulation frequency creates a continuous relative phase shift between the laser pulse and the intensifier gating waveform, as illustrated by the right-facing dashed arrow in Figure 6.8. Each captured image frame is therefore a sample of the image intensifier gain at a different phase offset.

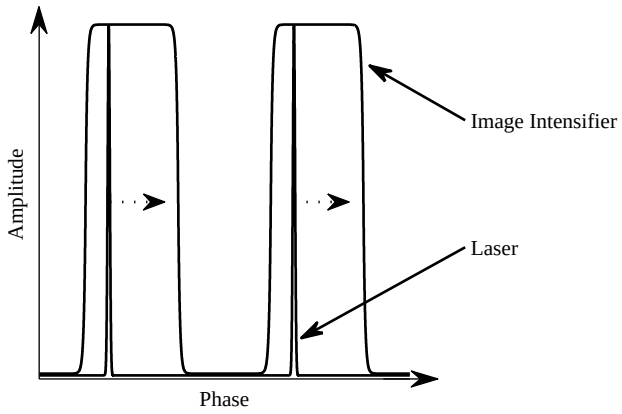

Figure 6.8: A pulsed laser is used to temporally sample the image intensifier gain (Payne et al., 2006b).

By configuring the DDS to output the laser and image intensifier signals with a very small frequency difference Δf, e.g. 0.1 Hz, it takes a long time (10 seconds) for the relative phase difference between the two signals to cycle through 2π. Selecting a high camera frame rate, e.g. 100 fps, allows a large number of samples to be collected during this period. The result is an image sequence, in this instance of 1000 frames, where the image intensifier gain waveform is recorded as the intensity of each frame.

A similar technique has previously been used where the picosecond laser and image intensifier repetition rates are identical, and either an electrical delay element (Thomas, 1994; Höß et al., 1998) or an optical delay (Thomas et al., 1991) is added to alter the laser probe timing with respect to the image intensifier waveform. Adjustment of these delays is a cumbersome process, particularly when using the optical delay method which requires altering the optical path length between captured images. The authors of these papers characterise the response of gated (or pulsed) image intensifiers, where high speed gating of the image intensifier occurs at a low repetition rate and the pulse width and repetition rate are unlikely to be significantly altered.

The significance between the heterodyne technique and previously used methods is the incorporation of a phase shift rather than a time delay between samples. This is beneficial when characterising a continuously modulated image intensifier, such as that used in the range imaging system, which is capable of operating at modulation frequencies spanning several orders of magnitude. For example, if 100 ps electrical delay elements were utilised using prior techniques, a 1 MHz modulation frequency would consist of 10000 samples, while a 100 MHz modulation frequency would consist of only 100 samples. In contrast,

from the earlier example with a camera frame rate of 100 Hz and Δf equal to 0.1 Hz, 1000 samples are captured regardless of the modulation frequency being tested. This makes it trivial to characterise and compare the image intensifier response over a wide range of modulation frequencies.

The ability to use the existing signal generation hardware, i.e. the DDS, is obviously also advantageous from the view that the only new hardware which needs to be introduced into the system to characterise the image intensifier is a pulsed laser source to replace the existing modulated lasers.

6.3.3 Picosecond Laser Pulser

A number of commercial products are available which are capable of producing a picosecond laser pulse, such as the PDL 800-B (PicoQuant GmbH, Berlin, Germany), although the requirements for this application make constructing a low cost, gain switched, laser pulser relatively simple.

Gain switching is a technique used in semiconductor lasers to produce a short optical pulse from a longer input electrical pulse (Lau, 1988). Making use of the earlier (problematic) relaxation oscillation described in Section 6.2, gain switching excites the first period of the oscillation and terminates the electrical pulse before the onset of the second period. Using this method, a short optical pulse can be generated from a longer electrical pulse or an appropriately biased sinusoidal waveform (Lau, 1988).

The simplified circuit shown in Figure 6.9 makes use of this technique (Payne et al., 2006b), and the full schematic diagram can be found in Appendix A. The sinusoidal input from the DDS is first converted to a CMOS logic level using a comparator. The output is passed to both inputs of an AND gate, although one input is first inverted. The ~3 ns propagation delay of the inverting logic gate generates a corresponding period where both inputs to the AND gate are high each time the comparator toggles on a positive edge, resulting in an output pulse from the AND gate of similar duration to the inverter propagation delay. The pulse controls an iC-HK laser switch, the same part used for the modulated laser driver in Section 5.1.2. The gain of the laser switch IC is adjusted by changing the voltage V_{ADJ} to generate the optimum optical pulse width. The final PCB is also shown in Figure 6.9.

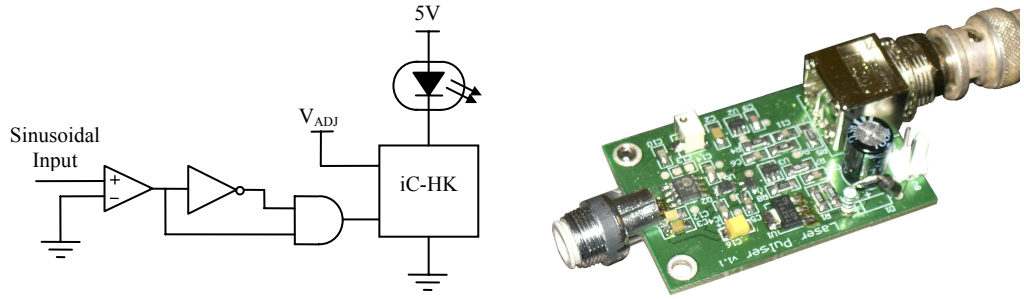

Figure 6.9: Left: Picosecond laser pulser circuit, Right: Picosecond pulser PCB.

The measured optical pulse is shown in Figure 6.10 using a SV2-FC photodiode (Thorlabs, NJ, USA). The left hand side shows a pulse with a FWHM of 266 ps, which is limited by the ~2 GHz bandwidth of the photodiode; hence the optical pulse duration may be shorter than that shown. The pulse in the right hand side shows the affect of setting V_{ADJ} (the laser switch gain) too high; the electrical pulse does not turn off immediately after the first period of the relaxation oscillation, and the second period of the oscillation is evident in the optical output waveform resulting in extended pulse duration. The optimal gain voltage V_{ADJ} is dependent on the pulse repetition rate (i.e. the DDS output frequency), and therefore must be manually adjusted to suit the desired image intensifier modulation frequency being characterised. If V_{ADJ} is too high, the broadening of the pulse width will reduce the resolution of the measurement, in contrast if V_{ADJ} is too low a reduction of the output intensity occurs, decreasing the measurement SNR. Note that the negative values in Figure 6.10 are an artefact of the photodiode impulse response due to electrical ringing, and are not part of the true optical pulse.

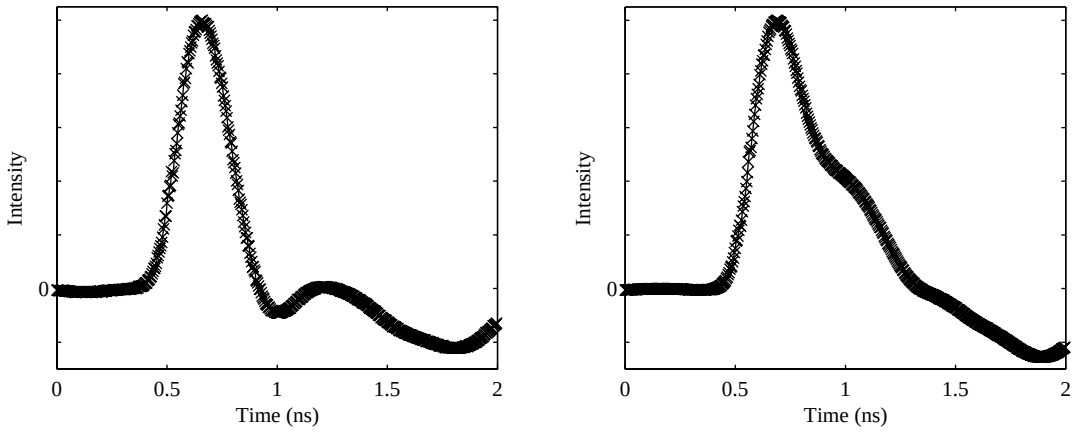

Figure 6.10: Left: Ideal laser pulse with FWHM 266 ps. Right: The laser switch gain is set too high resulting in the second period of the relaxation oscillation contributing to the output.

The pulse width of a gain switched laser diode can be reduced to tens of picoseconds (Suhara, 2004), although the bandwidth of the SV2-FC photodiode used limits the observable pulse duration. The measured 266 ps FWHM pulse duration is significantly shorter than the image intensifier modulation period (typically 10–100 ns) and hence is adequate for characterising the image intensifier gain without the need for further pulse width reduction.

6.3.4 Dynamic Characterisation Results

The system was setup as illustrated in Figure 6.7, with the pulsed laser source facing the image intensifier. The laser beam was expanded to illuminate the entire image intensifier input window, with a large separation (1.5 m) between the laser and image intensifier to minimise geometric variations in the pulse arrival time for different areas across the face of the image intensifier. Ground glass was placed in the optical path to remove interference patterns generated by the laser diode, and to ensure a uniform illumination intensity pattern. The imaging lens, which normally focuses the scene onto the image intensifier input window during normal operation, is removed during the characterisation.

The CCD was configured to operate at 100 fps, and the frequency difference Δf from the DDS set to 0.1 Hz, providing capture of 1000 frames over the modulation period. Each frame captures a large number (tens to hundreds of thousands) of pulses during the integration time, which is set to 8.2 ms, providing high SNR. To measure the temporal response, 49 pixels (7×7) in the centre of each frame are averaged together to further improve the SNR. The (normalised) intensity of each frame is illustrated in Figure 6.11 for modulation frequencies of 10 and 65 MHz respectively.

The response, particularly for the lower 10 MHz modulation frequency, is clearly far removed from the desired square wave profile. During each transition, significant ringing is apparent, with the ‘on’ state gain varying by up to 60%, while during the ‘off’ state the electrical ringing is sufficiently large to bias the image intensifier on for short bursts.

The electrical ringing is largely due the capacitance of the photocathode, combined with the inductance of the interconnecting wires and driver output. The capacitance is defined by the size of the photocathode, and the separation to the MCP, hence is fixed for a given image intensifier. In contrast, the connection between the driver and the photocathode can

easily be modified. The driver PCB is connected to a second PCB which is responsible for adding the required (high voltage) bias signals, before the combined voltages are connected to the image intensifier itself (refer Figure 4.23). To reduce the inductance of this path, the driver PCB was redesigned to combine the required bias voltages and the high frequency modulation output, removing the previous interconnection between the two PCBs. The wiring to the photocathode was also significantly reduced in length and twisted together, to reduce the inductance. The combination PCB is shown in Figure 6.12, and the (short) connection to the image intensifier can be seen in Figure 6.13. The PCB is mounted in a solid aluminium box which acts as a heatsink for the driver and also adds shielding to reduce electromagnetic interference (EMI) between other devices.

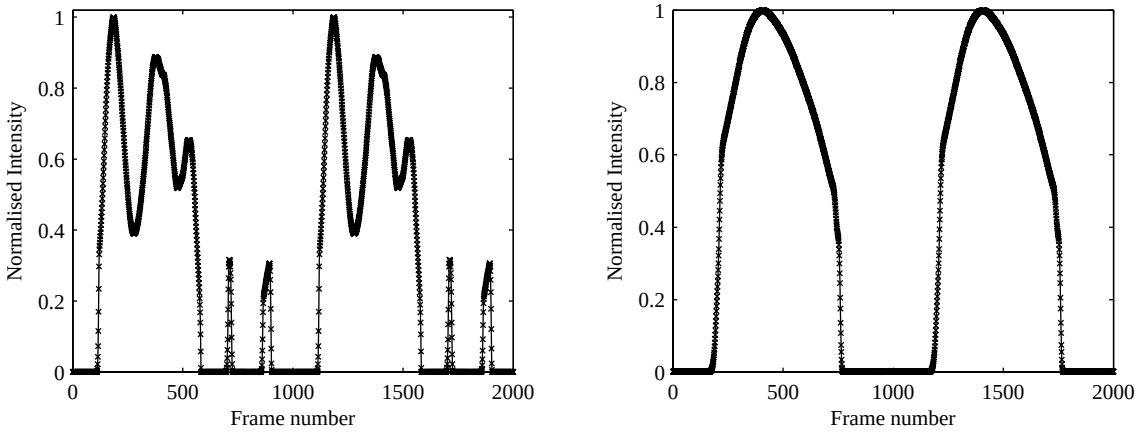

Figure 6.11: Measured image intensifier optical modulation temporal response for 10 MHz (left) and 65 MHz (right).

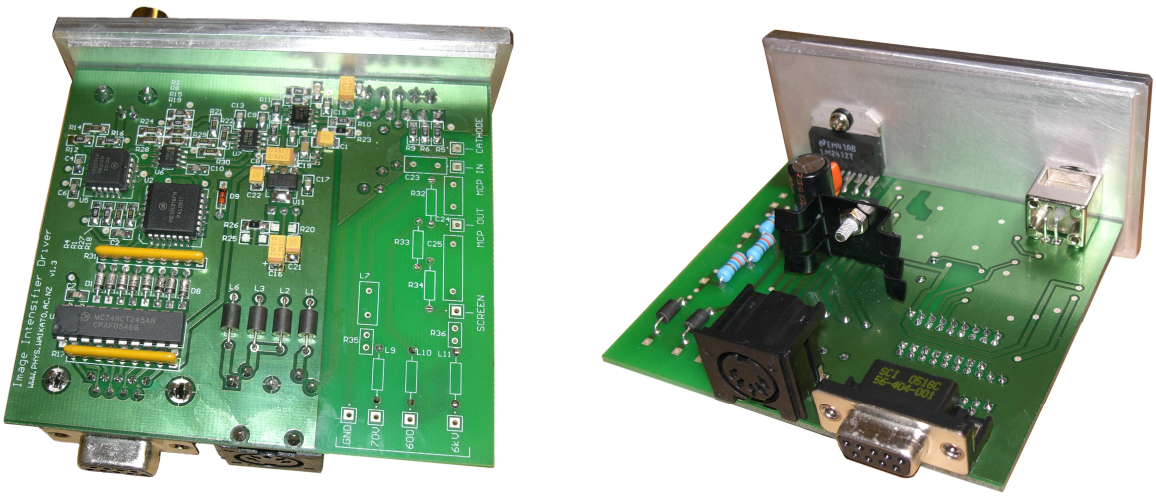

Figure 6.12: Revised version of the image intensifier driver PCB. The required image intensifier bias voltages, added separately in Figure 4.23, are incorporated into the RF amplifier PCB design.

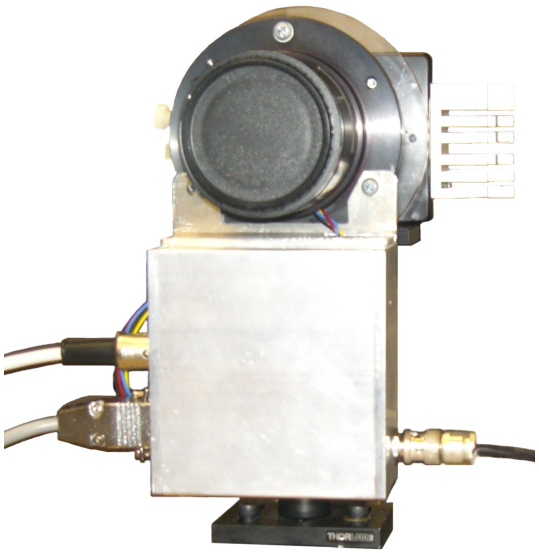

Figure 6.13: The revised image intensifier driver PCB with the connection to the photocathode made as short as possible. The aluminium housing acts as a heatsink for the driver, and also provides electrical shielding.

A second modification was also made to the MCP bias input voltage. The (relative) photocathode voltage from the driver PCB during capture in Figure 6.11 ranges from −40 V (on) to +10 V (off). To ensure that the voltage during the off state remains positive in the presence of electrical ringing, hence ensuring that no optical output occurs during this period, the bias voltage was decreased so that the relative photocathode voltage spans the range −38 V (on) to +12 V (off). This change lowers the image intensifier gain, but the reduction is minimal and is outweighed by the advantage that the spurious pulses seen in the 10 MHz capture in Figure 6.11 are removed. The response using the new PCB design is shown in Figure 6.14, where the former modulation waveform (dashed line) is also shown for reference.

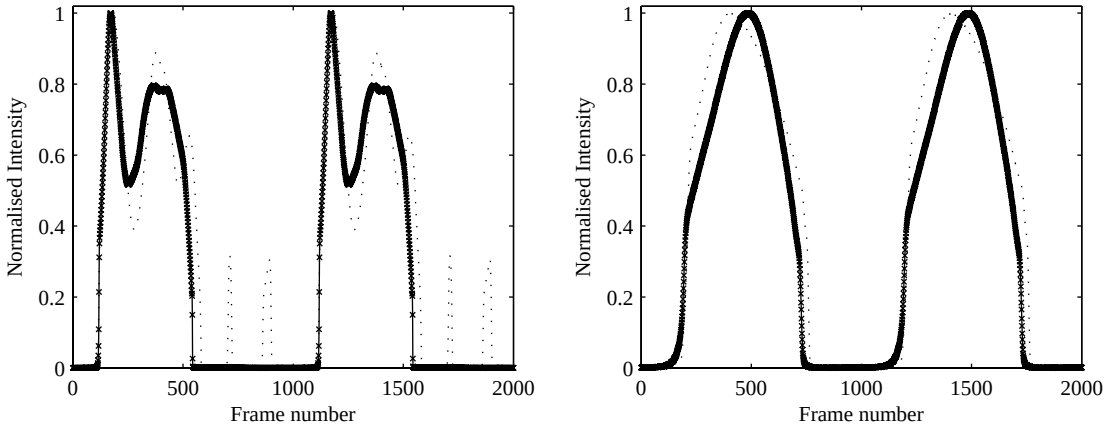

Figure 6.14: Measured image intensifier optical modulation temporal response for 10 MHz (left) and 65 MHz (right) using the modified driver PCB and bias voltages. The prior response from Figure 6.11 is shown for reference (dotted line).

The image intensifier optical modulation waveform using the modified PCB and bias voltages shows a significant reduction (approximately 20%) in the amplitude of the ringing on the rising edge, and the ringing on the falling edge no longer causes the output to turn on. The ringing still causes an undesirable modulation waveform containing significant harmonic content at the 10 MHz modulation frequency shown, which could potentially be further reduced through the use of impedance matching techniques between the amplifier output and the photocathode input. In practice it is not a significant problem for the range imaging application, as it is desirable to operate with as higher modulation frequency as possible to enhance the measurement precision (refer Section 3.1). Further characterisation of the temporal response at a modulation frequency of 85 MHz is shown in Figure 6.15, where the resultant profile approaches a much more ideal shape, providing an almost symmetrical profile containing no significant outliers. Due to RF power supply limitations, the response prior to revision of the image intensifier PCB was not measured at 85 MHz, hence is absent from the plot in Figure 6.15.

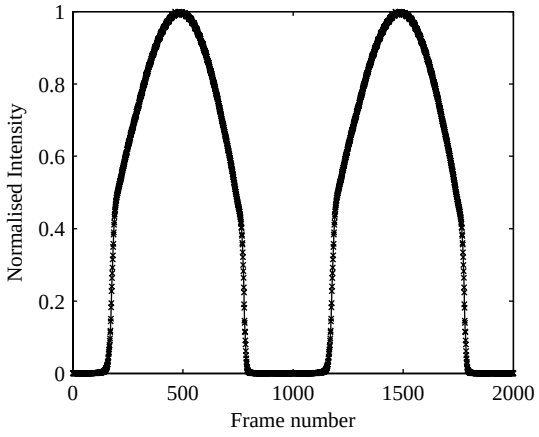

Figure 6.15: Measured image intensifier optical modulation temporal response at 85 MHz.

From Figures 6.14 and 6.15 it can be seen that the duty cycle of the optical modulation response increases with higher modulation frequencies. This is due to the asymmetrical voltage swing around zero volts, combined with the limited slew rate of the RF amplifier output. An attempt to compensate for this effect was made by reducing the duty cycle of the electrical waveform to approximately 40%, although further reduction is not possible due to the limited bandwidth of the CRT amplifier.

Within the image intensifier, the photocathode voltage is applied to a metal ring around the outer edge of the device and conducted through the photocathode material towards the

centre. At high modulation/gating frequencies a problem known as irising can occur (Detch and Noel, 1981; Thomas, 1994; Höß et al., 1998), where the modulation near the centre of the image intensifier is delayed relative to the outer edge due to the resistance of the photocathode material combined with capacitance to the MCP. This delay causes an error in the resultant range measurements, with objects near the centre of the image appearing at a larger distance (Dorrington et al., 2007a).

Rather than processing the captured data temporally, as in Figure 6.15, a single image frame can be selected and displayed spatially. Figure 6.16 shows one such frame from the rising (turn on) transition, where the image intensifier is modulated at 85 MHz. The data in Figure 6.16 have been spatially sub-sampled for clarity of the mesh. The spatial response clearly showing the irising effect, producing a bowl shape instead of a flat surface.

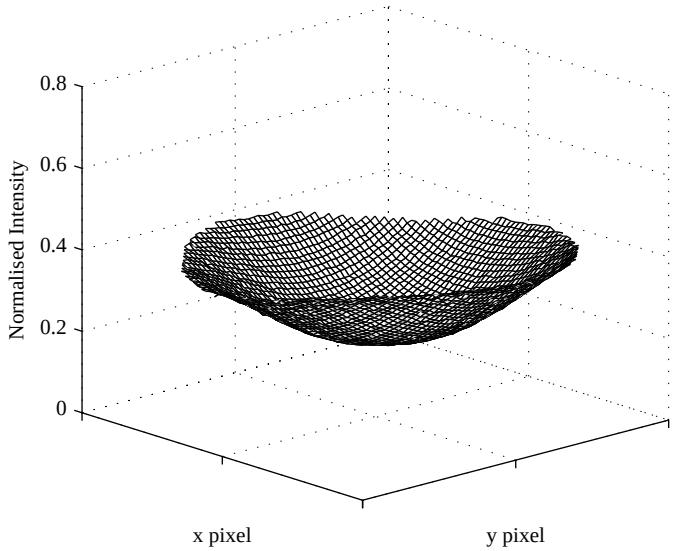

Figure 6.16: The instantaneous spatial response of the image intensifier during the rising edge, modulated at 85 MHz, shows significant irising.

Figure 6.15 illustrates the temporal response of a group of pixels in the central region of the image intensifier for a modulation frequency of 85 MHz, while Figure 6.16 illustrates the spatial response of the same data by taking a snapshot in time to plot a 3D surface. Figure 6.17 combines the data from each of these plots, illustrating the changing image intensifier gain over both time and space. A slice through the centre of Figure 6.16 is taken to produce a 2D line from the 3D bowl shaped intensity plot. Plotting the slices from a number of sequential frames, where each line represents a slice from a different phase delayed capture, illustrates the changing image intensifier gain over time. The left hand

side of Figure 6.17 demonstrates irising during the rising edge of the modulation waveform, where the plotted lines are ascending over time. Conversely, the right hand side of Figure 6.17 shows the falling edge with descending line plots. Only slices from every thirteenth captured frame have been shown in Figure 6.17 to improve the clarity of the plot, where each line plot is separated from its neighbours by a phase shift equivalent to a 153 ps time delay.

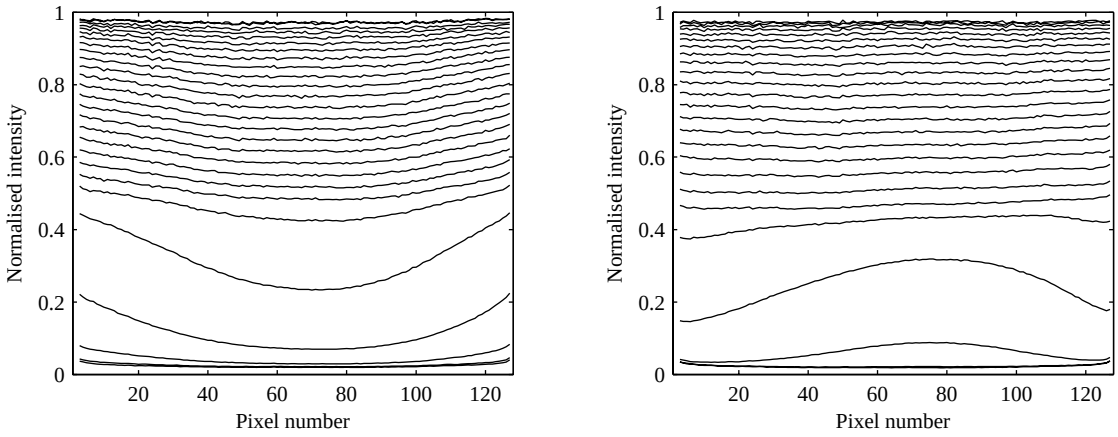

Figure 6.17: Irising during the rising (left) and falling (right) edges with an 85 MHz modulation frequency. Each plotted line is phase delayed, corresponds to a time delay of 153 ps.

Significant irising is visible during a ~500 ps window on both the rising and falling edges where a rapid change in intensity occurs (illustrated by a large spacing between the lines); however the irising is still present through the remainder of the modulation period despite being less pronounced. In Figure 6.17 the intensity at the edges of the image intensifier lead the intensity at the centre by approximately one line plot, hence the time required for the modulation voltage to traverse the photocathode from the edges to the centre can be estimated to be approximately 150 ps. From Equation (2.1), this delay corresponds to a range measurement deviation of approximately 23 mm, which is consistent with experimental observations (Dorrington et al., 2007a).

Irising is unavoidable when using an image intensifier modulated at high speed due to the physical design of the device, contributed to by the resistance of the photocathode material in conjunction with capacitance to the MCP. To reduce irising, the photocathode resistance can be reduced by applying a conductive coating or underlay during manufacture (Thomas et al., 1991), or a smaller diameter image intensifier can be used to reduce the photocathode capacitance. The image intensifier utilised already has a conductive coating, and although reducing the diameter (e.g. from 25 mm to 18 mm) is an

option, it comes at the expense of reducing the image quality while only reducing, and not completely eliminating, the irising effect. The irising effect should therefore be calibrated and corrected for within the range processing software to remove the influence on range measurements as suggested in Section 8.3.

6.4 Range Measurement Characterisation

To determine the measurement linearity and precision, a scene was constructed with eight flat white objects placed at distances between 1.5–5.5 m, with approximately 0.5 m spacing, as can be seen in Figure 6.18. An acquisition was performed with a modulation frequency of 15 MHz, a beat frequency of 1 Hz, and a camera frame rate of 29 Hz over a 10 second duration. The range, and scene intensity, is calculated from the resultant 290 image frame sequence, and the experiment repeated for modulation frequencies of 30,45,…,90 MHz.

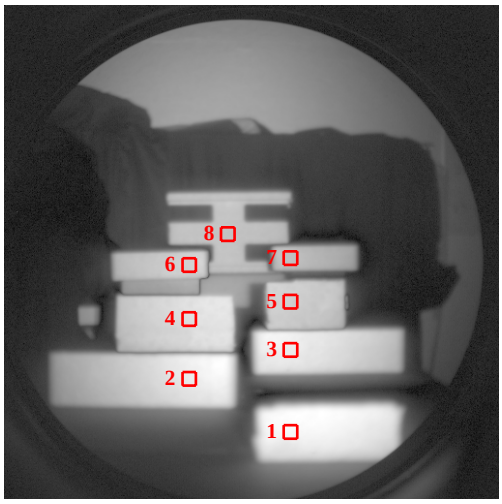

Figure 6.18: Gamma corrected intensity image from the range imaging system showing a field of targets at known depths, used to determine measurement linearity and precision.

The raw range data experienced systematic errors due to lens distortions and image intensifier irising (refer Section 6.3.4), and also requires calibration for the perspective projection coordinate system (Kahlmann and Ingensand, 2005; Dorrington et al., 2007a). The required calibration is outside the scope of the work performed here, and is conducted separately in the literature (Dorrington et al., 2007a). To make quantitative performance measurements from the raw range data without first performing this calibration, small areas within the field of view were analysed over which the systematic distortions can be

assumed to be negligible. Figure 6.18 shows the resultant intensity image; where a small region of 15×15 pixels was selected, and numbered, for each target.

The range data in the selected area of object number 4 are illustrated in Figure 6.19 from a capture at a modulation frequency of 90 MHz, and can be seen as a rough flat surface. From the assumption that the target is a perfectly flat object, a plane is fitted through the raw data using a least squares error fit, shown as a wire mesh, and is used as a local reference for the true range. A precision value is then determined by taking the standard deviation of the error between the measured range for each pixel and the fitted plane (Dorrington et al., 2006).

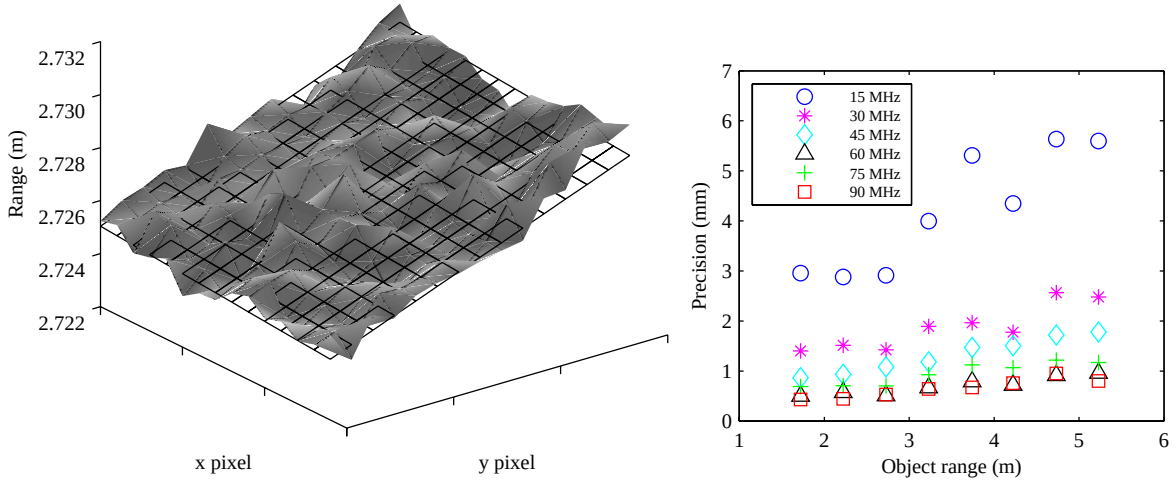

Figure 6.19: Left: Raw range data for target number 4, captured with a modulation frequency of 90 MHz. A plane is fit through the range data, illustrated by the mesh, which is used as a reference. Right: Measured one-sigma precision for each target for various modulation frequencies.

The right hand side of Figure 6.19 shows the measured precision versus the range of each object for the various modulation frequencies used. As predicted by Equation (3.1), the measurement errors tend to decrease for higher modulation frequencies. The 60 MHz measurement is an exception to this trend, outperforming the measurement at 75 MHz. Measurement precision is inversely proportional to not only the modulation frequency, but also the modulation index and square root of the SNR (refer Equation (3.1)), and, as demonstrated in Sections 6.2 and 6.3, the laser and image intensifier modulation waveform shapes vary with changes in modulation frequency. As a result, the beat signal modulation index is therefore also dependent on the modulation frequency. Measurement of the modulation index at each of the targets in Figure 6.18 for each modulation frequency is given in Table 6.1. Despite the 75 MHz measurement having a 25% higher modulation

frequency than the 60 MHz measurement, the modulation index value is reduced by 41%, explaining the anomaly in precision in Figure 6.19. This result reiterates the importance of optimising not only the modulation frequency, but also the modulation index of the high frequency waveforms.

Table 6.1: Mean modulation index value for the eight targets shown in Figure 6.18 for the various modulation frequencies used.

Frequency	Mean modulation index
15 MHz	0.36
30 MHz	0.44
45 MHz	0.42
60 MHz	0.54
75 MHz	0.32
90 MHz	0.40

From the precision values in Figure 6.19 it can be seen that the random error increases with range. This is expected as the amount of light reaching the receiver decreases by $1/r^2$, where r is the object range, hence resulting in a decrease of SNR. From the graph, the relationship between precision and range appears to be predominantly linear.

At the maximum frequency of 90 MHz, the one-sigma precision values are between 0.43 and 0.95 mm over the target range of 1.7 to 5.2 m. This result demonstrates the system's capability to perform measurements to the desired millimetre level precision, meeting the project requirements.

Despite the presence of uncalibrated lens distortion, image intensifier irising, etc., the measured raw distance to each target can be used to determine, to a first approximation, the linearity of the range measurements. Ideally a target mounted on an automatic interferometric track, such as that implemented by Kahlman and Ingensand (Kahlmann and Ingensand, 2006) could be used, where the target remains in the centre of the image so that no spatial distortions affect the measurement result, and the true change in the target distance can be measured with high accuracy. This would allow a comprehensive analysis of the measurement linearity, but such a system was not available. Instead, an independent range measurement to each of the eight targets in Figure 6.18 was made with a Leica DISTO™ laser distancemeter (Leica Geosystems, Switzerland), which is compared against the measured range values in Figure 6.20. The independent range measurements may use slightly different positions on each target to those selected in Figure 6.18, increasing the reference measurement uncertainty. The measured range values are from the 90 MHz

modulation frequency acquisition, with a beat frequency of 1 Hz, and a camera frame rate of 29 Hz over a 10 second duration. An offset of −1.676 m exists in Figure 6.20 between the measured and true range values due to propagation delays in the high frequency signal paths. This is a constant offset (for a given modulation frequency), hence can simply be subtracted from the range measurements during software processing.

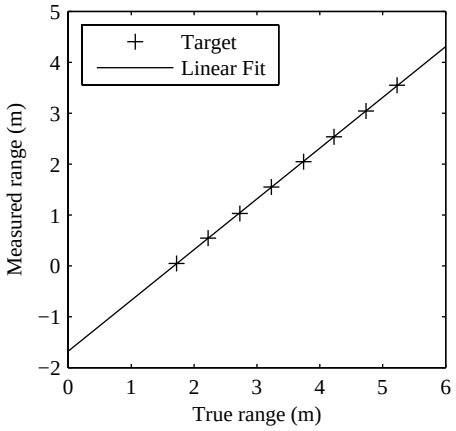

Figure 6.20: Range linearity measurement from the scene shown in Figure 6.18.

The range data exhibits excellent linearity over the depth of the scene as illustrated by the linear fit in Figure 6.20. The slope of the fit is 0.997, which is in close agreement with the expected value of 1 considering the error sources involved. The RMS value of the residuals is 7.5 mm.

6.5 Summary

To optimise measurement precision and performance, each of the major components within the range imaging system were independently characterised to identify and improve potential deficiencies in their operation. The DDS exhibited numerous unwanted spurious frequency components in the spectral output, which contribute jitter in the subsequent laser and image intensifier modulation signals. The addition of an external low pass filter significantly improved the spectral purity, increasing the SFDR from 34.5 to 51 dBc, reducing the resultant (theoretical) jitter from 140 to 60 ps peak to peak for the example 55 MHz modulation frequency studied. A similar improvement is realised for other modulation frequencies by selecting a low pass filter with a cutoff frequency slightly above the output frequency, removing contaminating spectral components from the output.

Measurement of the laser output modulation waveform showed a "relaxation oscillation" during the turn on transition, a phenomenon related to setting the bias current below the laser diode threshold level. The ideal bias current level is equal to the (temperature dependent) laser diode threshold current; setting the current too high reduces the output waveform modulation index value, while setting the current too low instigates the presence of the relaxation oscillation. Instead of tracking and compensating for changes in temperature, a simple alternative is to limit the transition speed of the electrical modulation current by placing a small capacitor across the laser diode connection. The speed of the laser diode is reduced using this method; however modulation frequencies up to 100 MHz can still be achieved as required by the range imaging system. Although the average output power from the laser diode still exhibits some temperature dependency, the laser turns off completely during each cycle to ensure that a high modulation index value is achieved to maximise the range measurement precision.

Electronic measurement of the image intensifier photocathode voltage is an inadequate method of characterising the image intensifier response, due to both the difficulties involved in performing such a measurement, and the inability to resolve variations across the 2D surface. An optical method of measuring the waveform was instead developed, utilising a circuit which exploits gain switching to generate picosecond laser pulses. The measured image intensifier response demonstrated the presence of significant electrical ringing, which was partially overcome by redesigning the driver PCB to reduce inductance. The MCP bias voltages were also adjusted to prevent ringing from turning the image intensifier on following the waveform falling edge. At high modulation frequencies, such as that shown in Figure 6.15 at 85 MHz, the temporal response is close to ideal with no undesirable peaks, while still exhibiting an excellent modulation index value. The optical characterisation technique also provides a method to measure spatial variations across the image intensifier, and measurement results clearly illustrate the presence of irising, where modulation is delayed near the centre relative to the outer edges. Irising is an inevitable consequence of image intensifier technology, and must be corrected for by the data processing software.

Quantitative performance measurements of the complete system were generated using a number of flat panel objects, spaced over a range of 1.7–5.2 m. The one-sigma precision

varied from 0.43 mm to 0.95 mm, with an apparent linear dependence on object distance (due to the associated decrease of intensity of the reflected light). This result demonstrates successful fulfilment of the project requirement to achieve millimetre range measurements. The uncalibrated lens system makes it difficult to accurately test the measurement linearity over a given scene, despite this limitation, the raw measurement data in Figure 6.20 fits well to the expected straight line.

7. Investigation of Digital Signal Generation

The heterodyne system requires a very small frequency difference between the two modulation signals, accomplished in Section 4.2 using a DDS with analogue stages within the design. The digital → analogue → digital process used by each DDS (and following comparator) is relatively complex, inefficient and more expensive than a completely digital path. In contrast, a phase stepped homodyne system can be designed with simpler digital timing signals which can be generated from a single logic device such as a CPLD or FPGA. This approach is much more efficient and cost effective. Ideally a hybrid of the phase stepped homodyne and heterodyne methods could be created, retaining the benefits of each technique. This would allow simple digital generation of the high frequency modulation signals,` while maintaining flexibility to easily change the number of samples collected during the beat signal period, avoiding errors due to harmonics and quantisation (refer Sections 2.1.8 and 3.2.2).

This chapter begins with an investigation of the effect of jitter on the ranging system's performance, where the high frequency signals experience a systematic cycle-by-cycle timing error. It is shown that no corruption of the range measurement occurs if the jitter mean value is zero (over the CCD integration period), and that if the modulation waveforms contain harmonics, jitter may actually be considered beneficial. It is proposed that the continuous (analogue) phase shifting in the heterodyne system is equivalent to

discrete (digital) phase stepping, with an additional systematic jitter component. Digital alternatives to the analogue (DDS) circuit are investigated, resulting in a unique and innovative modulation control scheme. The new scheme cancels the effect of harmonics from the modulation signals during data acquisition, producing range measurements with high accuracy (linearity) without the need to acquire a large number of samples. Finally, a modification to the illumination system is demonstrated, reducing the range measurement uncertainty. The combination of these two techniques enhances the range measurement linearity and precision, compared with the heterodyne system, while the acquisition time remains constant.

7.1 Effect of Jitter

The high frequency DDS outputs inherently contain spurious frequency components due to the synthesis process used as shown in Section 6.1. If these frequency components are not completely removed by the analogue filter stages, then the output signal will inevitably contain jitter when the signal is converted to a square wave through a comparator. In this section, an example of this jitter is shown in a simple time domain simulation, with a synthesised frequency of 74 MHz and a DDS clock frequency of 400 MHz. Figure 7.1 shows the staircase shaped analogue output from the DDS before it is filtered. In this instance the position of each "step" in the output waveform varies from cycle to cycle, an effect which occurs for all output frequencies which are not an integer factor of the clock frequency.

The staircase output is passed through an analogue "reconstruction" filter to create the desired sine wave. Ideally the reconstruction filter should consist of a narrow band pass filter centred at the desired output frequency, producing high spectral purity. In practice, a low pass filter must be used due to the requirement of generating a wide range of output frequencies. This, coupled with imperfections within the filter components and layout, DAC and amplifier distortions, and contamination from the reference oscillator, results in the spectral purity of the final output being less than ideal (refer Section 6.1). From a temporal point of view, the spurious components distort the waveform shape on a cycle by cycle basis. Figure 7.2 shows the staircase output from Figure 7.1 after it has passed through a single pole low pass filter. A basic filter is used in this example to clearly

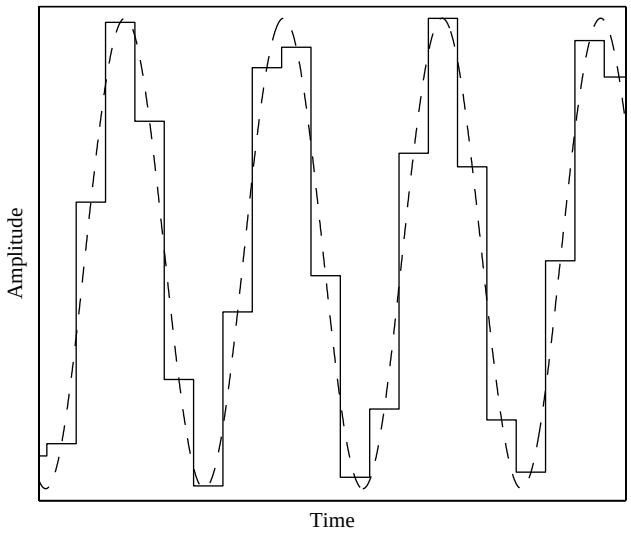

Figure 7.1: The output from the DDS before the low pass filter (solid line) compared with the desired sinusoidal output (dashed). Note in the DDS output that the "steps" occur in a different position in each iteration.

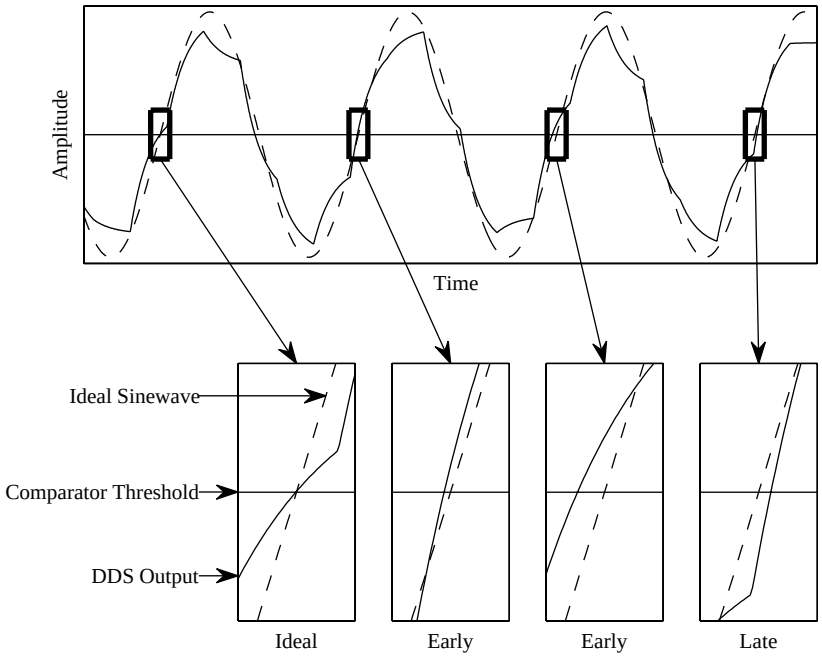

Figure 7.2: Top: The staircase output from Figure 7.1 has been low pass filtered (solid line), and is compared with the desired sinusoidal output. Bottom: Enlarged view of each rising transition. The time the DDS output crosses the comparator threshold to trigger the output varies cycle by cycle compared to the ideal sinusoidal output.

demonstrate the undesired effect on the output waveform, although this simplification exaggerates the resultant jitter. The simulation ignores the other sources of error mentioned above that also contribute spurious frequency components in the output. The ideal sine wave is again shown for reference, as well as the comparator threshold level which is used to generate the square output waveform. The rising edge of each cycle is

enlarged to show the point at which the ideal and actual waveforms cross the comparator threshold voltage. This transition point changes every cycle for the DDS waveform, creating jitter on the generated square wave. This is a well known effect when using a DDS to generate a clock source (Brandon, 2006), where the standard solution is to use a narrow band pass filter as described above.

To maintain wideband operation, a selection of discrete band pass filters could be switched in over various frequency ranges, or a field programmable analogue array (FPAA, the analogue equivalent of an FPGA) could be used to generate a dynamically programmable analogue filter. Coupled with an improved PCB design to reduce contamination from the digital reference clock, the resultant jitter could be significantly reduced, although it would still remain limited by the performance of the comparator and the spectral purity at the comparator inputs. The PCB design uses multiple AD9952 DDS ICs, each containing an independent phase locked loop (PLL) to multiply the 20 MHz reference clock to the required 400 MHz system clock rate. The PLL output frequency is constantly adjusted within a feedback loop to match the target frequency, which leads to further phase jitter at each DDS generated output. The magnitude of this source of jitter is dependent on the bandwidth of the feedback loop, and therefore can, to some extent, be controlled or minimised. To completely cancel this PPL jitter from the ranging system, a newer AD9959 DDS IC can be used to replace the AD9952. The AD9959 can generate up to four synthesised frequency outputs from a single IC, with the advantage that all outputs share a common PLL; hence any phase jitter is common to all outputs and is therefore effectively cancelled. Clock distribution from the reference oscillator to multiple DDS ICs, and the synchronisation clock between the DDS ICs is no longer required if an AD9959 is used; hence crosstalk between the digital input and analogue output can be significantly reduced, improving the spectral purity and further reducing the jitter. The prominent solution is therefore to redesign the DDS PCB to use the newer AD9959, and design a switched/dynamic narrow band pass filter to reduce the jitter magnitude.

Or is it?

Despite the presence of jitter, the constructed ranging system is capable of generating measurements with (sub) millimetre precision. The magnitude of jitter within the complete system (including the laser driver, etc.) was not directly measured, although was

significant enough that cycle by cycle variations in the laser pulse duty cycle were observed (due to fluctuations in the timing of the rising and falling edges). By studying Figure 7.2 it can be deduced that the resultant jitter is dependent on the system clock frequency, i.e. a higher clock frequency will produce a larger number of small 'steps' per output cycle, resulting in a better approximation of the desired sinusoid, therefore reducing jitter at the output. This is also evident in analysis of the produced frequency spectrum, as frequency components are generated at $nf_{clock} \pm f_{output}$, where n is an integer (0,1,2…), f_{clock} is the DDS system clock frequency, and f_{output} is the output frequency. By increasing f_{clock}, the undesirable frequency components fall at higher frequencies, where the low pass filter provides greater attenuation. For analysis purposes, it is assumed that the peak to peak magnitude of the jitter ($Jitter_{pp}$) is one tenth of the system clock period (a conservative estimate from experimental observations). Taking the 400 MHz clock used in the DDS design, using a single laser pulse and assuming a direct time of flight measurement, results in an uncertainty of

$$\begin{aligned} \frac{c}{2} \cdot Jitter_{pp} &= \frac{3.00 \times 10^{8}\ \text{ms}^{-1}}{2} \cdot \frac{1}{10 \cdot 400 \times 10^{6}\ \text{Hz}} \\ &= 32.5\ \text{mm}. \end{aligned} \tag{7.1}$$

It is quickly apparent that jitter severely constrains the measurement accuracy in a direct time of flight case. Even if the estimated jitter magnitude was decreased by an order of magnitude from the already conservative estimate, jitter would still prevent sub-millimetre range measurements from being performed.

In contrast, the heterodyne method (and phase stepped homodyne method) integrate a large number of laser pulses (~10^6) in a single image frame, and use multiple image frames, each with a different modulation envelope phase offset, to calculate the object range. Jitter from the DDS is largely systematic due to the "steps" in the synthesised waveform, leading to a frequency dependent, predictable pattern which is added to random fluctuations due to other sources of noise. When applied to the range imaging application, this frequency dependent output pattern can be divided into two cases; the first occurs where the average modulation envelope phase shift (over an image frame) deviates from the expected value, generating an error in the range measurement. The second case occurs when the jitter averages to zero.

7.1.1 Measurement Error Due to Systematic Jitter

A significant effect occurs when the DDS system clock frequency is an integer multiple of the generated output frequency: there is no cycle to cycle variation in the synthesised output waveform. As an example, an output of 66.$\dot{6}$ MHz from a 400 MHz system clock consists of six "steps", all occurring at the same phase angle during each cycle as shown in Figure 7.3. After being low-pass filtered, the signal may cross the comparator threshold voltage at a different time to the ideal sine wave; however it will remain constant on each repetition hence can be considered just as a fixed offset. As there is no variation in the cycle to cycle analogue waveform, no systematic jitter will occur in the output. At this point it is worth noting that a PLL can generate output frequencies which share an integer multiplicand/dividend relationship with the reference frequency; hence for this particular frequency the DDS does not contribute any useful function.

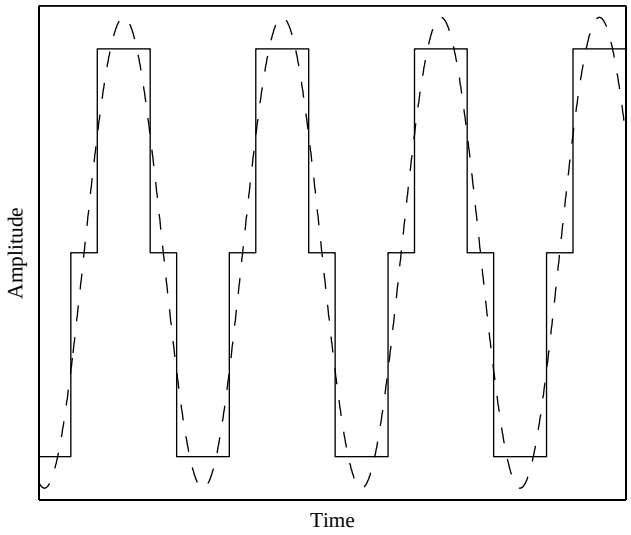

Figure 7.3: DDS DAC output at a frequency of 66.67 MHz. The output contains an integer number of "steps" which always occur at the same point in every cycle.

For the heterodyne method, a second frequency is required with a few hertz difference to that of the first. Assuming selection of the first frequency at 66.$\dot{6}$ MHz is a good choice as it contains no systematic jitter, a second frequency is generated with a small frequency difference, in this case at 66.86$\dot{}$ MHz, to illustrate the consequence of using a DDS. (The frequency difference of 200 kHz is used to clearly illustrate the effect; in practice, the difference is of the order of a few hertz.) The output waveform shape changes during each cycle as described earlier, however as the output frequency is *close* to an integer

relationship with the system clock frequency, only a very small change occurs in each cycle as shown by Figure 7.4.

The stair case waveform is low pass filtered (again just using a single pole filter for simplicity), and the output is enlarged around the comparator threshold voltage on each rising edge. In Figure 7.5 the DDS output consistently traverses the comparator threshold after the ideal signal, leading to retardation of the generated signal timing.

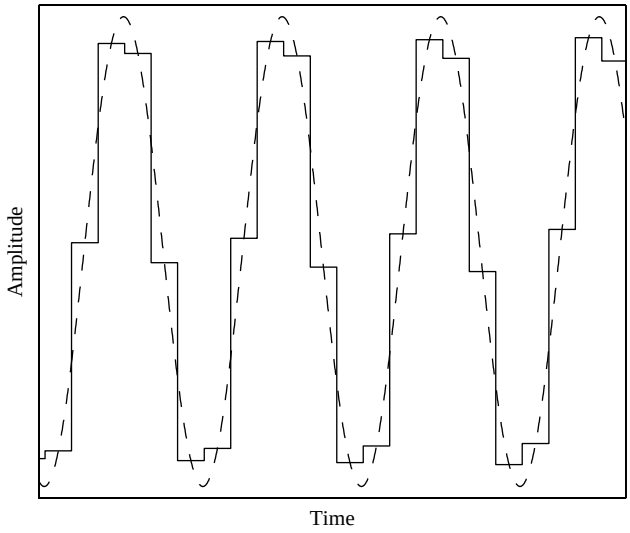

Figure 7.4: DDS DAC output at a frequency of 66.87 MHz. The "steps" occur at a slightly different location during each cycle.

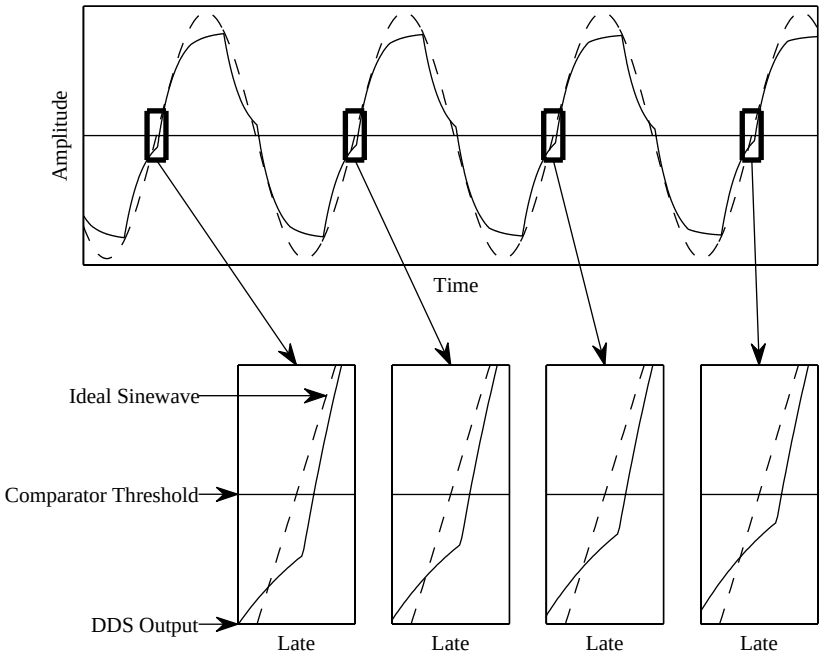

Figure 7.5: Jitter is not evenly distributed due to limited cycle to cycle variation in the waveform shape. In this example the DDS output consistently traverses the comparator threshold after the ideal sine wave, retarding the generated signal.

As the steps *slowly* move through the generated waveform, at some point the DDS waveform will traverse the comparator threshold earlier than the ideal sine wave, and the output waveform timing will advance. The slowly changing advancement/retardation is carried through the heterodyne mixing process and is present in the low frequency beat signal, with a final result equivalent to non-uniform phase shifted sampling of the beat signal. Figure 7.6 illustrates the error in the average sample positions, assuming a base frequency of $66.\dot{6}$ MHz, a frequency difference of 1 Hz, and a camera sampling rate of 10 Hz, where the size of each arrow is proportional to the magnitude of the error, and the direction indicates an advancement or retardation.

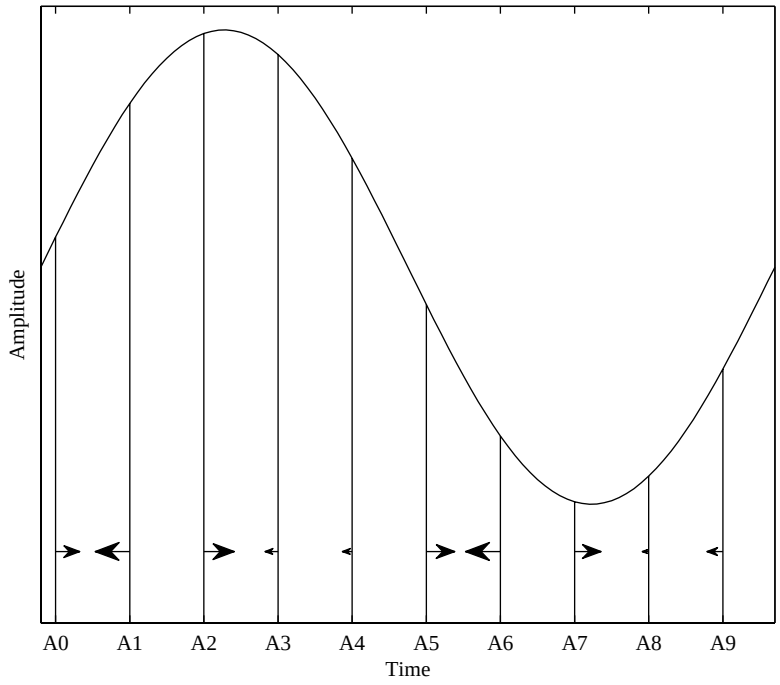

Figure 7.6: The low frequency beat signal is not sampled evenly due to slowly varying jitter. The plotted arrows indicate the offset from the ideal sample location, and the size of the arrow head indicates the magnitude of the offset.

Non-uniform sampling leads to an error in the range measurement as Equation (2.24) assumes the modulation envelope phase shift between samples is constant. Theoretically it is possible to compensate for the jitter by appending an additional offset term to Equation (2.24) for each sample. However, the term is dependent on the modulation frequency, PCB and filter design, and also on time (the jitter for each repeated sample will be different, albeit predictable, on each cycle of the beat frequency output). Compensating for this effect through post processing is therefore not a practical option.

7.1.2 Avoiding Error Due to Systematic Jitter

In the instance where the synthesised output frequency is *not* close to an integer relationship of the system clock frequency, the position of each step in the output waveform varies significantly from cycle to cycle as illustrated previously in Figure 7.1. The resultant jitter therefore varies rapidly (has a high frequency component), which the relatively long integration period integrates to a zero mean offset. The modulation envelope therefore has a constant average phase shift between each sample; hence Equation (2.24) is applicable, and the measured range can be calculated without the error discussed in Section 7.1.1.

Of course jitter does still affect the measurement. The resultant waveform shape, at the end of the integration period, is the contribution of the modulation waveform shape with a small phase offset applied during each cycle. If sinusoidal modulation is utilised, the waveform becomes slightly distorted, creating harmonics. A square modulation waveform is shown as an example in Figure 7.7. The simulation assumes that the jitter has a Gaussian profile, indeed, a similar outcome results from any distribution with zero mean. The resultant effective illumination waveform as seen at the end of the integration period shows that the sharp corners of the square waveform are "rounded off", corresponding to a

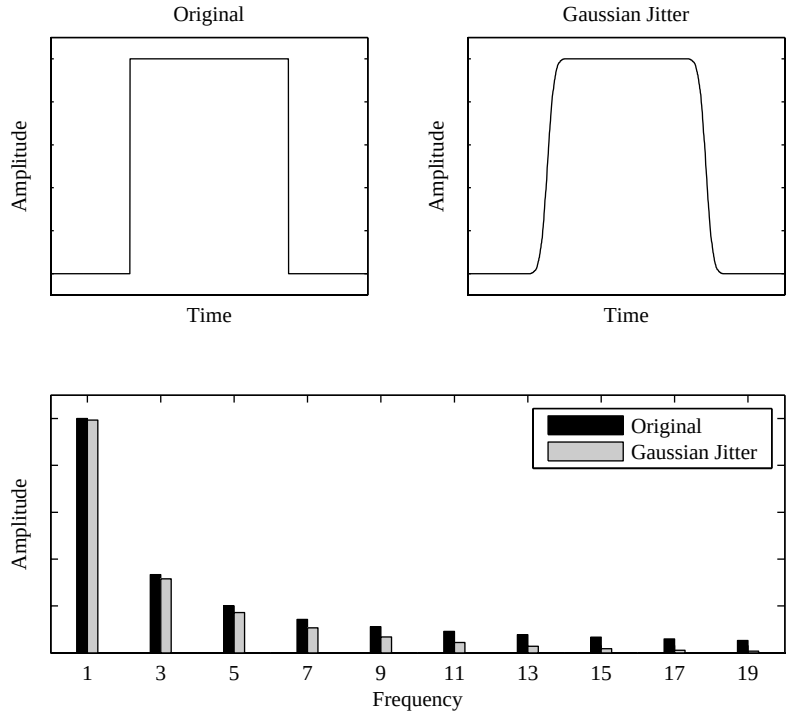

Figure 7.7: A square modulation illumination waveform (top left) is "rounded off" by jitter during the integration period (top right). The higher order harmonics are reduced in the resultant waveform (bottom).

reduction of higher order harmonics as shown in the lower half of Figure 7.7.

Harmonics can cause a systematic error in the range measurement (refer Section 3.2.2); hence for sinusoidal modulation addition of harmonics due to jitter can be detrimental to the measurement results, whereas for square wave modulation the reduction of harmonics can potentially be beneficial. The amplitude of the fundamental frequency component is also reduced, reducing the modulation index, hence measurement precision (refer Section 3.1.2). In the example shown in Figure 7.7, this change is less than 1% and can be considered as negligible.

For the designed hardware, the modulation waveforms have been characterised (refer Sections 6.2 and 6.3) and clearly contain harmonics, therefore, following from the above example, jitter can be advantageous to suppress the higher order harmonics. To ensure that the jitter for each sample has zero mean, consideration must be taken into account during selection of the synthesised modulation frequencies (alternatively the DDS system clock frequency can be altered if a particular modulation frequency is required).

7.2 Digital Frequency Divider

As jitter can actually be considered useful in the presence of harmonics, it is worth investigating a completely digital system (which can be envisioned as being equivalent to an analogue system with large systematic jitter content). A FPGA or CPLD can be used as a programmable counter that divides the input frequency by an integer value to produce the desired output frequency. Using an input frequency of 400 MHz (equal to the DDS system clock frequency used earlier) and integer division by eight that produces an output frequency of 50 MHz, it is possible to shift the phase of the waveform by 45° (360°/8) by programming the initial counter start value.

To capture 30 frames during the beat cycle, the modulation phase must advance by an average of 12° (360°/30) per frame, which is clearly beyond the capable phase resolution of the digital counter outputs. By breaking the long camera integration period into smaller parts, the modulation phase can be stepped at a specific moment during the integration period to produce the required *average* phase shift. For example, the first 12° phase shift could be produced by using modulation phases of 0° and 45° with a ratio of 0.73:0.27. Table 7.1 shows the ratio required for a selection of image frames.

The average modulation phase shift requirement can be met using this technique, but at the expense of an undesirable waveform shape. Using a square wave input, the resultant *effective* waveform shape for the first few frames are shown in Figure 7.8, along with the corresponding frequency spectrum.

Table 7.1: Ratio of integration times to generate the required phase angle per frame to capture 30 samples per beat cycle using a digital frequency divider.

Frame Number	Average Phase Shift	Integration Time at Each Input Phase Angle				
		0°	**45°**	**90°**	**…**	**335°**
A0	0°	1				
A1	12°	0.73	0.27			
A2	24°	0.47	0.53			
A3	36°	0.20	0.80			
A4	48°		0.93	0.07		
⋮						
A29	348°	0.73				0.27

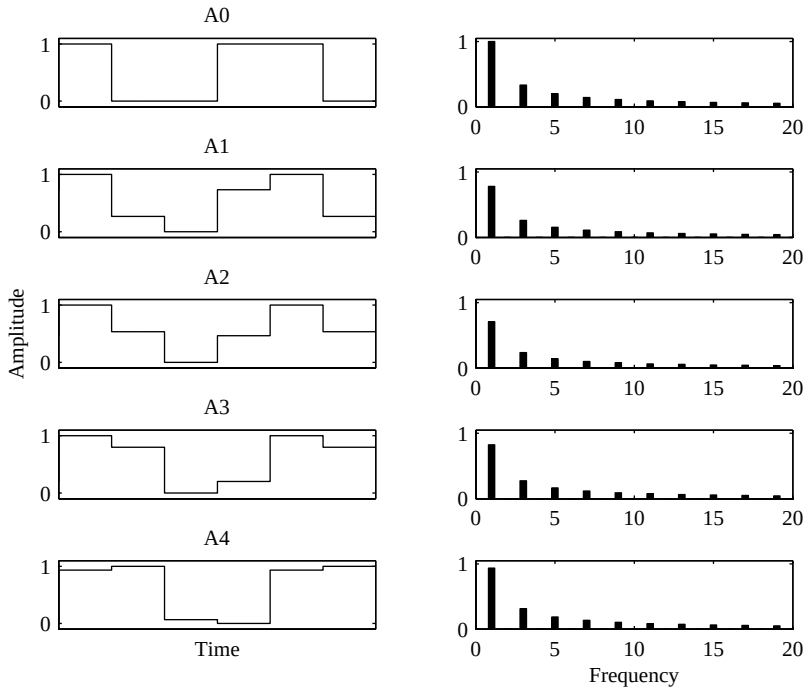

Figure 7.8: The first five resultant waveforms and their frequency spectrum using a digital frequency divider to generate the required average phase shift per frame, assuming a square wave input, 30 frames per beat cycle, a modulation frequency of 50 MHz and a system clock frequency of 400 MHz.

The considerable frame by frame variation in the effective waveform shape makes this method impractical. Of particular interest are the significant fluctuations in the magnitude of the fundamental frequency component, which can be seen in the frame by frame comparison of the frequency spectrum. This method can be considered analogous to using linear interpolation; the signal generator is capable of generating eight modulation phase steps in the above example, yet extraction of 30 samples is desired. The intermediate

samples do not create the true shape of the beat cycle waveform, and can only be considered as an estimate of the actual value. Given that the motivation for taking 30 samples is to remove the influence of harmonics and quantisation error, these contaminants still influence the range measurement result, hence this approach does not contribute any useful advantage.

7.3 Phase Stepped Heterodyne

From the previous section it is apparent that extracting 30 samples per beat cycle from hardware capable of only generating eight phase offsets does not produce the same result as an analogue system, hence to utilise a digital signal generator, finer phase step resolution is required. The phase step resolution of the hardware in the previous example is proportional to the system clock frequency of the FPGA or CPLD, assuming a programmable counter is used to divide the system clock frequency by an integer. Using this basic method to generate 30 phase offsets would require the system clock to operate at 30 times the output frequency (e.g. at 1.5 GHz to output 50 MHz), which far exceeds the maximum operating frequency of such devices.

An alternative method is to use a reconfigurable phase locked loop (PLL), such as that available within Altera Stratix II or III series FPGAs. The internal components of the PLL can operate at much higher frequencies than that of standard logic; hence finer phase step

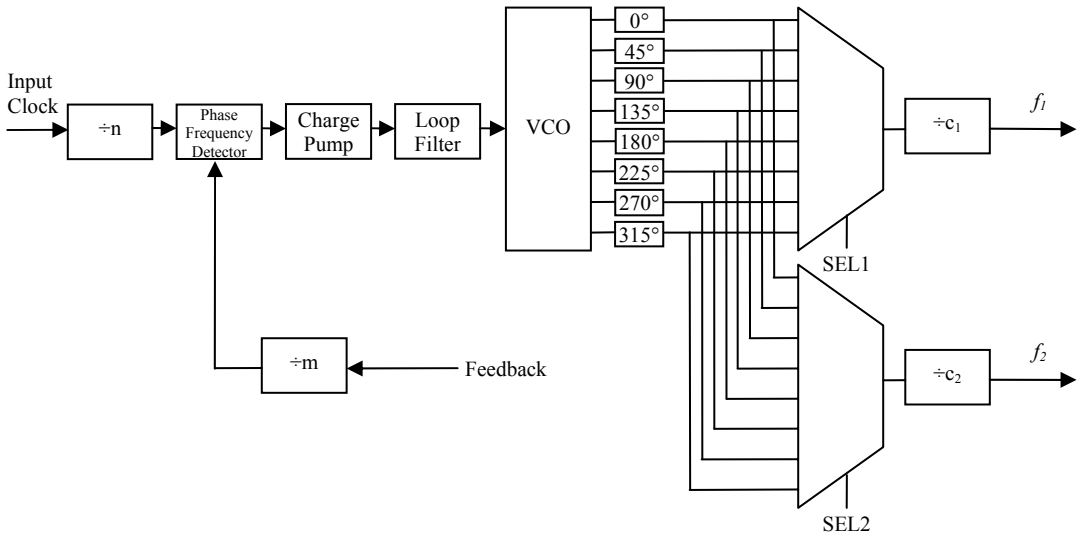

Figure 7.9: Schematic of Altera Stratix II Phase Locked Loop showing the components of interest for generating a phase stepped signal.

resolution can be produced. Additionally the output of the voltage controlled oscillator (VCO) for these particular phase locked loops has eight phase taps, each at a different phase offset, allowing even higher phase step resolution to be produced.

The basic components of interest within the PLL are shown in Figure 7.9. An integer divider in each of the input clock and feedback paths are used to set the VCO frequency, which is capable of operating at a maximum frequency of 1 GHz in the Stratix II devices. A multiplexer selects one of the eight VCO phase taps, which is then divided by a post scale counter to produce the required output frequency. The distinctive feature of this PLLis that the selected VCO phase tap can be reconfigured dynamically while the PLL is operating, changing the phase of an output signal with a resolution of

$$\text{Phase Step Resolution} = \frac{\text{Output frequency}}{\text{VCO frequency}} \times \frac{360°}{8}. \tag{7.2}$$

Using the same 50 MHz requirement as the example in Section 7.2, with a 1 GHz VCO frequency, the PLL is capable of generating phase steps with resolution of 2.25° (160 steps per 360°). This resolution is 20 times higher than that of the previous digital approach.

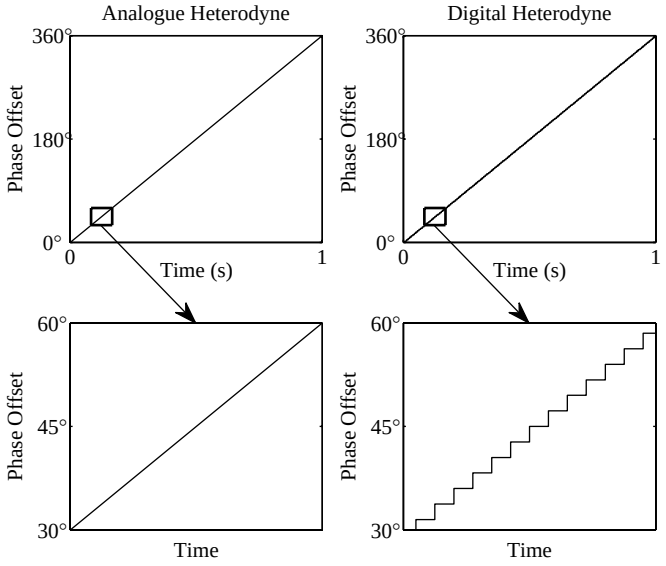

Figure 7.10: Comparison of the changing phase of the heterodyne beat signal generated with an analogue system (left), and the approximation generated with a phase stepped digital output from a PLL (right).

By generating two outputs from the PLL with the same frequency, and periodically incrementing the phase of one output, the continuous phase change (i.e. single hertz frequency difference) of the heterodyne configuration can be emulated. Figure 7.10 illustrates the difference between the true (analogue) heterodyne beat signal phase and the emulated (digital) approximation. The analogue system generates a beat signal with a

smooth, continuous phase change over a 1 second period (corresponding to a beat frequency of 1 Hz). In comparison, the digital approximation, or "digital heterodyne" system generates 160 small equally spaced phase steps over the same period.

Applying both the analogue and digital heterodyne techniques to a system using a square illumination modulation waveform results in an effective waveform as shown in Figure 7.11 (taking into account the integration period for a camera frame rate of 30 Hz). The simulation assumes a noise free environment without the presence of jitter. In the time domain representation, the analogue system effectively changes the square signal into a trapezoidal shape with a smooth slope on either side. The digital system response closely

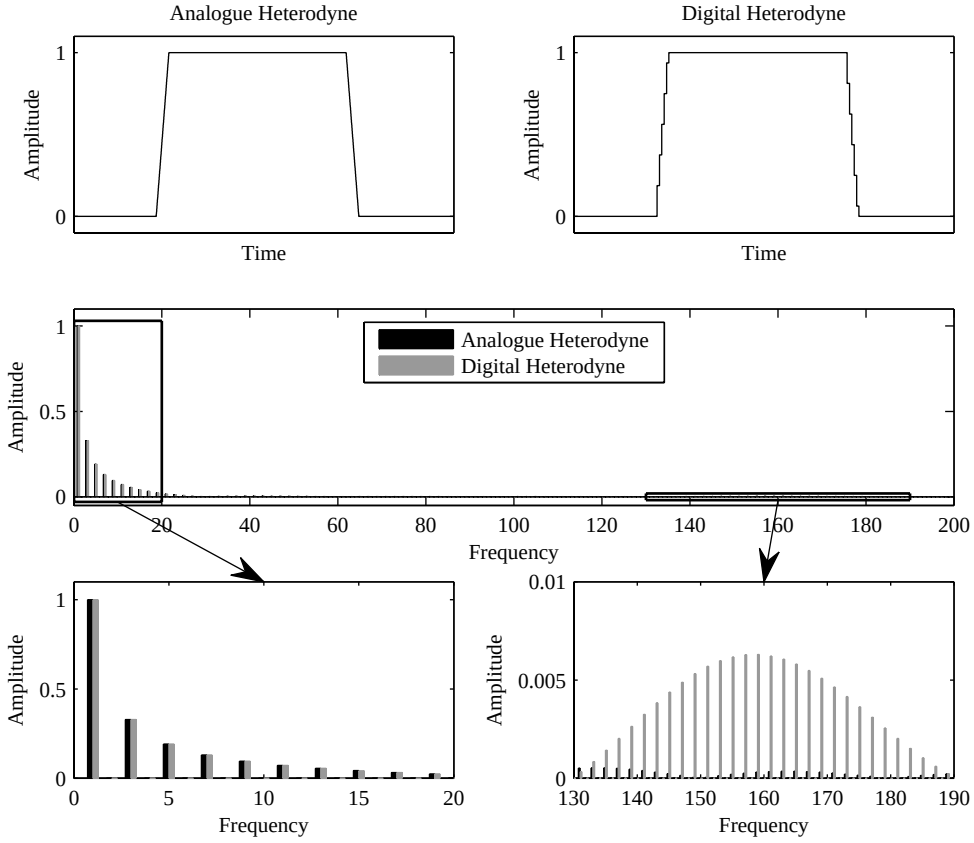

Figure 7.11: Temporal (upper) and frequency domain (mid and lower) comparison of the analogue and digitally generated heterodyne techniques applied to a square modulation waveform.

resembles that of the analogue system; however the sides of the trapezium are now jagged due to the abrupt phase steps during the integration period. Comparing the two waveforms in the frequency domain, it can be seen that the harmonic content at lower frequencies is identical between the two waveforms, however the digital phase stepped system introduces

some higher order harmonics. The additional harmonics are centred at the 160th order harmonic, which corresponds to the 160 phase steps (per 360°) used in the simulation.

Both the illumination and receiver modulation waveforms must contain harmonics of the same order for the beat frequency to also contain harmonics of that order (refer Section 3.2.2). The illumination modulation frequency in the example given is 50 MHz; hence the receiver modulation waveform would have to contain harmonics to 8 GHz for the introduced harmonics to be included in the resultant beat signal. This is unlikely to occur due to the limited bandwidth of the receiver, which coupled with the attenuation of higher order harmonics (refer Section 3.2.2), implies that the signal captured by the camera will be indistinguishable between the analogue and digital heterodyne systems.

This method of generating the modulation frequencies offers the advantage that the difference in modulation frequencies $f_2 - f_1$ is derived from a low frequency signal. The earlier system used a DDS with a high system clock frequency (~400 MHz) and a phase accumulator with a large number of bits to achieve the fine tuning required (single hertz resolution), hence deriving the beat signal from the high frequency source. A third low frequency signal was required, also derived from the high frequency clock, to act as a phase reference for the resultant beat signal. In contrast, the phase stepped PLL system generates the frequency difference $f_2 - f_1$ directly from a low frequency signal, which can be used directly as the reference, simplifying the design.

7.3.1 FPGA Implementation and Results

The digital heterodyne scheme incorporates the entire timing and control system within a single FPGA as illustrated in Figure 7.12. The FPGA communicates with a PC via an RS232 serial interface, allowing the PC to specify the modulation frequency, time interval between phase steps (i.e. the beat frequency), the camera integration time, etc. A control logic block uses this information to configure the PLL and a series of counters, and to initiate the output. The counters trigger the camera, prompt the PLL to advance to the next phase step, and generate a reference signal to synchronise the captured video sequence at a known phase offset by tagging one image frame to allow absolute range measurements to be made.

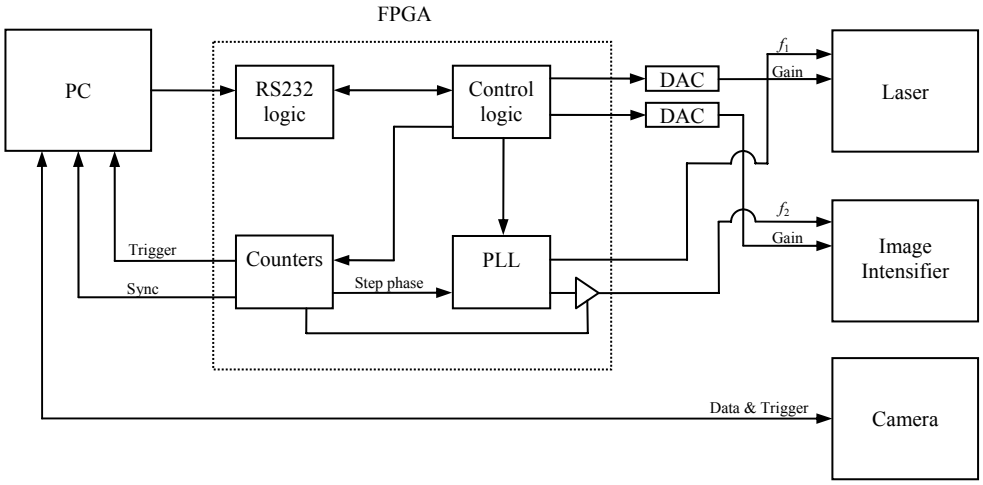

Figure 7.12: Implementation of digital heterodyne system using an FPGA to replace the DDS and associated timing and control components.

The image intensifier is enabled/disabled by gating the 2nd high frequency PLL output signal within the FPGA, removing the need for the additional logic on the image intensifier driver PCB to prevent image smear during the CCD frame transfer period. It is possible to also perform the "pulse skipping" function using logic within the FPGA, again removing the need for logic devices (the high speed counter) on the image intensifier PCB. Figure 7.13 illustrates the timing of the three main signals from the FPGA for a single captured image frame. The laser is modulated continuously in order to maintain a constant temperature, hence constant optical output power. The camera begins integration, and after a short time delay the image intensifier modulation signal is enabled. Near the end of the integration period, the image intensifier signal is disabled, and the frame trigger signal is lowered to initiate the CCD readout transfer.

Potentially, the camera trigger may experience a propagation delay between the FPGA signal transition and actual commencement of image exposure at the camera. To overcome any resultant error in the range measurement, the image intensifier operates during a window within the camera integration time, as illustrated in the timing diagram of Figure 7.13; hence any small variations in the camera trigger from the PC can not affect the captured image. This timing also ensures that the image intensifier has turned off completely (including the phosphor decay time) before the image is transferred from the CCD to prevent image smear. The 'sync' signal is not shown in the timing diagram, but operates in the same fashion as the earlier system, tagging one frame per beat cycle in the captured image sequence.

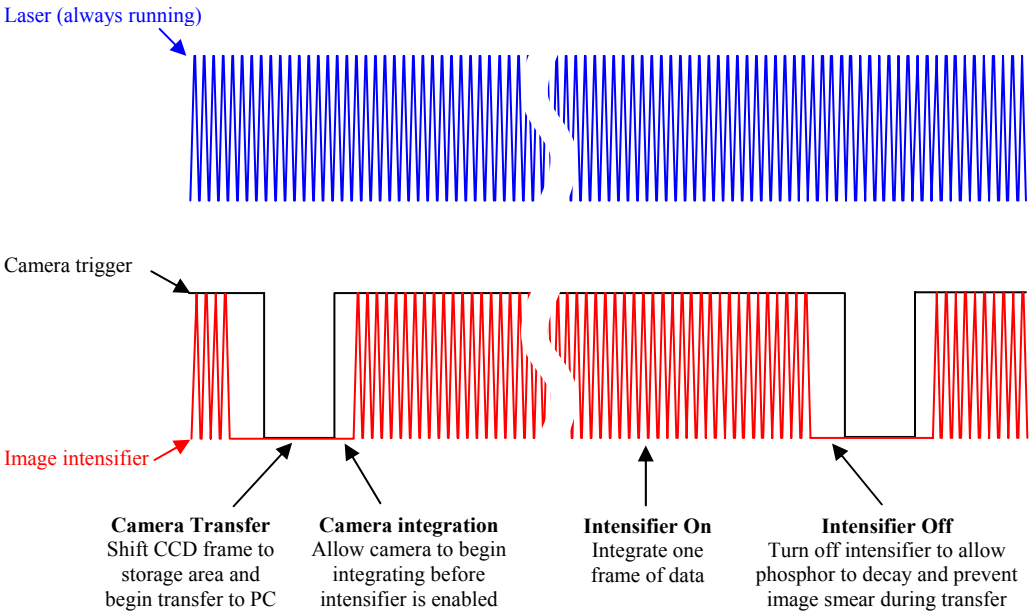

Figure 7.13: Timing diagram illustrating the camera trigger, laser, and image intensifier modulation signals over one captured image frame. The image intensifier signal is synchronised to the camera trigger to avoid image smear during the CCD transfer stage.

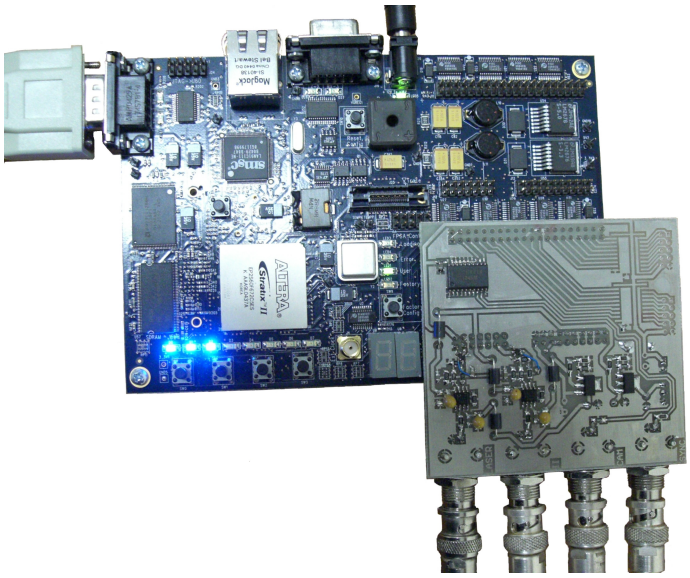

Figure 7.14: Altera Stratix II development kit with an add-on PCB to interface with the laser driver, image intensifier driver, and camera.

Code to implement the digital heterodyne modulation scheme was written in VHDL and downloaded to an Altera Stratix II development kit (Altera, CA, USA) shown in Figure 7.14. An interface PCB was designed to convert the LVCMOS signals from the FPGA to bipolar signals with 50 Ω impedance, as required by the existing laser and image intensifier driver circuits. The software allows the modulation frequency, beat frequency, and camera timing signals to be set by the PC using the serial interface. Data acquisition

can also be started and stopped from the PC, disabling the image intensifier when it is not in use to allow it to cool down.

In order to experimentally compare the digital heterodyne performance with the analogue system, a flat panel was imaged using both techniques. The scene, as well as the camera and illumination optics remained unchanged; the only difference between the acquisitions was the method used to generate the modulation and camera timing signals. The parameters used during the measurement were:

- 80 MHz modulation frequency.
- 1 Hz beat frequency.
- 29 fps camera frame rate.
- 10 second acquisition.
- Target approximately 1.5 m from camera.

The white panel is used to give a high intensity return signal, so that the two methods can be compared with little influence due to random noise. The centre of the image was divided into 25 regions of 30×30 pixels as shown in Figure 7.15, where lens distortion and vignetting are minimised. The precision was measured in each region using the technique shown in Section 6.4, in which a flat plane is fitted through the measured points and the standard deviation of the difference between each pixel and the fit gives the precision. Two captures were performed for each method, with the precision values shown in Figure 7.16 for each of the 25 regions.

The mean precision value for the analogue (DDS) system was measured as 0.44 mm, compared to 0.57 mm for the digital (FPGA) system. At first glance this appears to be a large difference, however to put the result into perspective, both methods are achieving a challenging feat of sub-millimetre precision at a distance larger than one metre by determining the beat signal phase with approximately 1.5 mrad precision. The difference between the two sets of data can be explained by the fact that the illumination and image intensifier modulation driver circuits derive both the frequency and duty cycle from the signal generator. Despite operating at the same frequency, the duty cycle is likely to be slightly different when the sinusoidal (analogue) signal from the DDS is replaced with the square wave (digital) signal from the FPGA. The duty cycle affects the modulation index

value and average illumination intensity (hence SNR), both of which are related to the measurement precision (refer Section 3.1). A limited attempt to match the duty cycle between the two systems was made; however the practicality and time required to measure, and precisely adjust, the duty cycle for multiple modulation frequencies being tested meant that it was only a rough estimate. From the result shown, the digital system is a suitable replacement for the analogue signal generator.

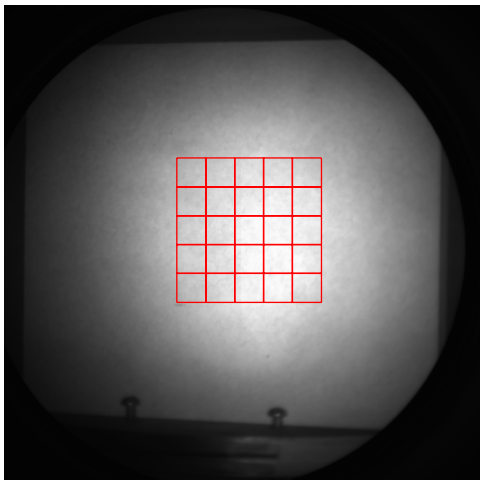

Figure 7.15: A white flat panel is used as a reference plane to compare precision using the analogue and digital heterodyne techniques. Precision is measured in small 30×30 pixels regions within the centre of the image.

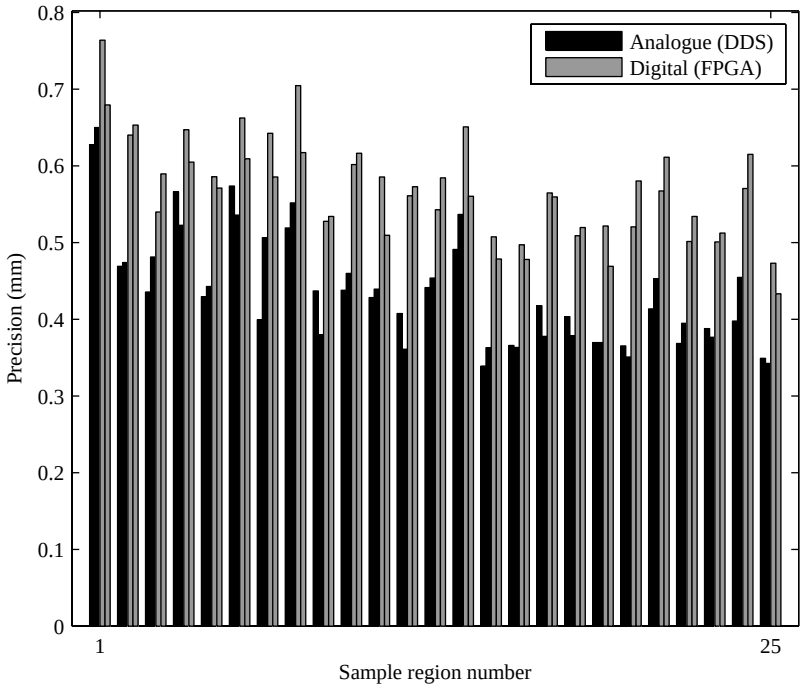

Figure 7.16: Precision of each 30×30 pixel region as shown in Figure 7.15 for the analogue and digital heterodyne techniques. The mean value for the analogue technique is 0.44 mm, compared to 0.57 mm for the digital method.

The use of an FPGA also brings a number of other advantages in addition in the compact and simplified design. Under certain conditions, electrical interference has been noticed in captured images using the DDS system. The interference manifests itself as visible horizontal bands across the image as shown in Figure 7.17, and is due to electrical cross talk between the image intensifier driver and the camera internal electronics (in particular the analogue to digital conversion stage).

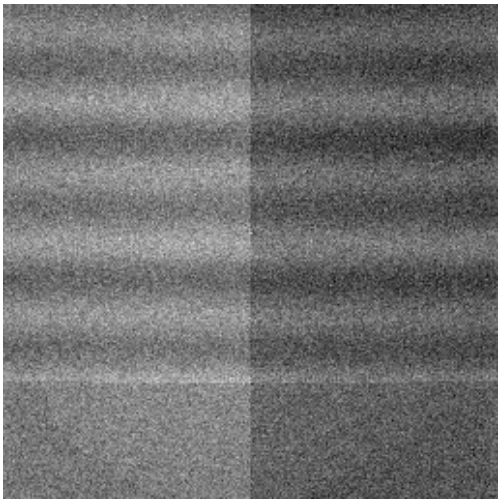

Figure 7.17: A dark image captured where the illumination and image intensifier waveforms are out of phase. Horizontal banding is visible due to electrical interference between the image intensifier driver and the camera internal electronics.

Figure 7.17 was taken with a modulation frequency of 40 MHz, and a camera frame rate of 100 fps. The illumination and image intensifier modulation waveforms were out of phase, leading to a dark image. The intensity difference between the left and right hand sides of the image are due to the CCD used, where the image is split down the centre and digitised by two ADCs to increase the camera frame rate, however in this case the two ADCs suffer from a slight gain mismatch. An interesting feature can be seen near the bottom of the image; the horizontal banding is non-existent. This is due to the fact that the image intensifier is disabled to prevent image smear during the CCD transfer period, which coincides with the time that the lower part of the image is being digitised and sent to the PC. The image intensifier driver is enabled to begin exposure of the next image frame, and it begins interfering with the electronics within the camera resulting in the banding seen near the top of the image as the remainder of the image is digitised.

Ideally the level of electromagnetic interference could be reduced by applying additional shielding and modifying the PCB layout, as described in Section 6.3.4. However, this

solution is not always a simple option. The influence on the resultant images is dominated heavily by the choice of operating conditions (in particular the modulation frequency and camera frame rate), so interference could be avoided by careful selection of these parameters. Alternatively the image intensifier could remain disabled for an extended period to allow the image to be completely transferred to the PC before starting exposure of the following image, although this approach reduces the acquisition speed due to the additional delay between image frames. Altering the image intensifier and camera timing would also not prevent interference with other electronic devices in the near vicinity. An alternate technique used in some electronic devices, such as PCs, is the use of a spread-spectrum clock. The FPGA PLL has the ability to generate such a clock signal, which modulates the frequency to spread the electromagnetic radiation over a wider frequency range, reducing interference to devices which are susceptible to a narrow band of frequencies (such as the camera used). It also makes it easier to meet EMI regulations. Both the illumination and image intensifier modulation signals are derived from the same PLL, hence when spread-spectrum clocking is enabled both signals experience the same variation in frequency. The variation can be analysed in the same manner as random jitter in Section 7.1, and has no significant ill effect due to the effective averaging over the integration period.

The logic design implemented uses very few of the available resources within the FPGA (approximately 1% of the available look up tables are used); hence ample resources are available for other tasks. The most prominent task is to process the data from the camera directly within the FPGA, returning the final range values, rather than the camera raw data, to the PC. The FPGA is suited to this type of task, with hardware elements capable of applying "multiply and accumulate" operations to rapidly perform the required calculations. By pre-processing the data within the FPGA before sending the results to the PC, valuable PC computation is made available for other tasks within the final application.

7.4 Intra-Frame Phase Shifting

In the previous section we saw that the FPGA is capable of emulating the analogue heterodyne configuration, achieving measurements with sub-millimetre precision by periodically incrementing the phase to approximate that of the DDS system. The FPGA

however, is capable of much greater flexibility and control of the phase, providing the opportunity to create more complicated phase shifting profiles.

In Section 3.2.2 it was shown that harmonics were aliased onto the fundamental frequency component during sampling, causing an error in the range measurement. In particular, when four samples per beat period are used, the third order harmonic is responsible for a significant contribution towards the systematic error (assuming square illumination and receiver modulation waveforms). The heterodyne configuration reduces the magnitude of the harmonics due to the sinc term in Equation (3.18), which is dependent on the phase shift $\Delta\theta$ of the fundamental component during the integration period. The DDS generates a continuous phase shift; hence the term $\Delta\theta$ is constrained by the number of samples collected per beat cycle. The continuous phase change of the heterodyne system is shown in the left hand side of Figure 7.18, where four samples are collected during the beat period; therefore $\Delta\theta$ is $\pi/2$ rad (90°). Using the FPGA, $\Delta\theta$ can be extended to $2\pi/3$ rad (120°) while capturing four samples by modifying the start and end phase offset of each integration period, and maintaining an average phase offset identical to the heterodyne case, as illustrated in the right hand side of Figure 7.18. The phase is clearly now discontinuous between samples, which cannot be achieved using the DDS system. The term $\Delta\theta$ was chosen as $2\pi/3$ (120°) to minimise the magnitude of the sinc function (Equation (3.18)) at the third order harmonic as illustrated in Figure 7.19.

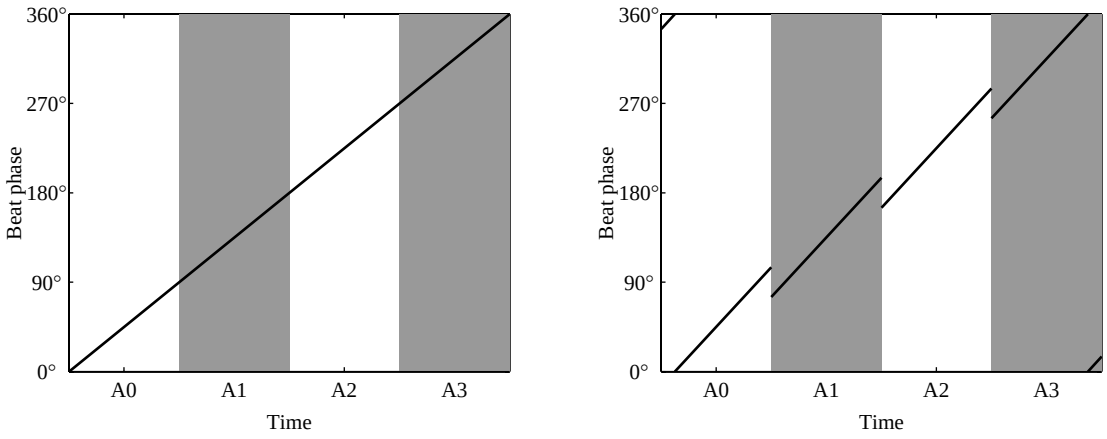

Figure 7.18: Left: For traditional heterodyne modulation the phase of the beat signal increases continuously. In this example, four images A0-A3 are collected during acquisition, each spanning a phase change of $\pi/2$ (90°). Right: Using the FPGA allows the phase change during each image period to be independently chosen, in this case to $2\pi/3$ (120°).

Simulating the response of the modified heterodyne system with square modulation waveforms, cancellation of the third order harmonic results in a significant reduction of the

cyclic error in the phase (range) measurement. The cyclic error is ±9.5 mrad, as shown in Figure 7.20, which is significantly less than ±46 mrad for the heterodyne case (refer Figure 3.8); hence this modified heterodyne technique offers an advantage over the standard heterodyne (and homodyne) methods. The increase in linearity comes at the expense of reducing the modulation index of the fundamental frequency component, and therefore reducing the measurement precision (refer Section 3.1.2); hence this trade-off must be taken into account when selecting modulation methods.

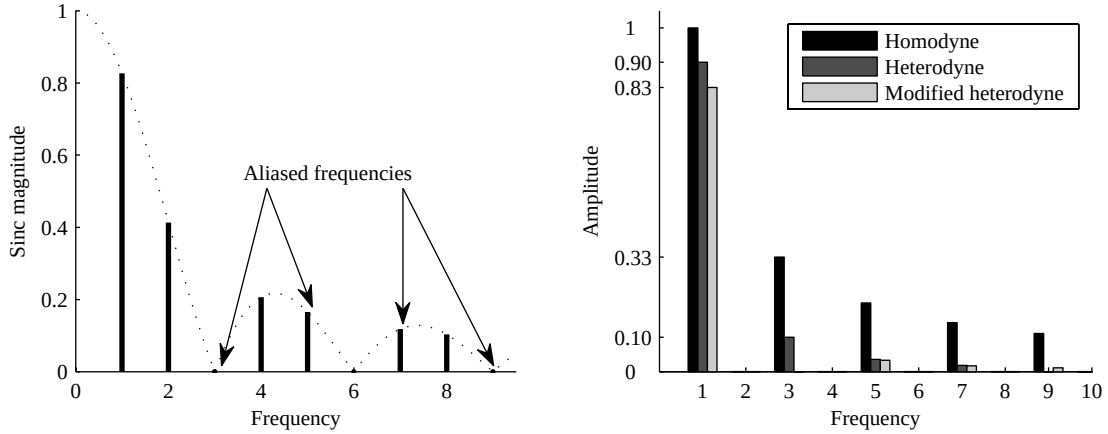

Figure 7.19: Left: By selecting $\Delta\theta$ equal to $2\pi/3$, the magnitude of the sinc function falls to zero at the first aliased frequency. Right: Comparison of the resultant magnitude of each harmonic for square wave illumination using the homodyne, heterodyne, and modified heterodyne configurations.

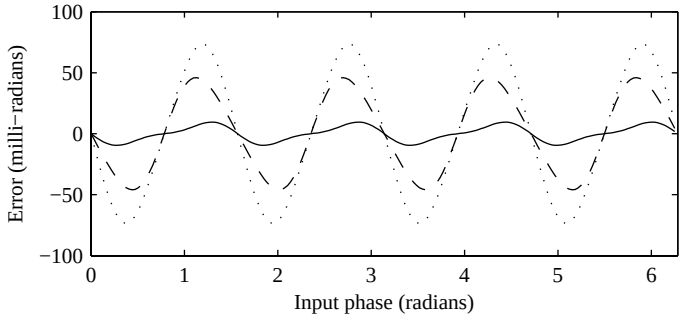

Figure 7.20: Simulation of the recovered phase error for square illumination and receiver gain modulation waveforms with four samples per beat cycle. The modified heterodyne technique error (solid line) is significantly less than that of the homodyne (dotted line) or heterodyne (dashed line) techniques.

7.5 Harmonic Rejection

A further enhancement can be made by recognising that the phase increments do not have to be periodic when using the FPGA; i.e. the intra-frame phase shifting profile does not have to be linear.

In the radio communications field, CMOS designs using digital switches for generating outputs at GHz frequencies have become common. The generation of a square wave is

used to simplify the RF electronic design, but as applications are affected by harmonics, the signal is subsequently filtered using discrete components. An interesting solution has been developed to eliminate the need for the costly, discrete component filter: three or more phase shifted square waves are added together, generating an approximation to a sinusoid (Weldon et al., 2001). A simple case is shown in Figure 7.21. Three square waves are generated with a common frequency, a π/4 rad (45°) phase offset between one another, and an amplitude ratio of $1:\sqrt{2}:1$. The waveforms are summed within the circuit, generating the resultant waveform shown in the bottom of Figure 7.21. The frequency spectrum of the new waveform is given in Figure 7.22, where it can be seen that the previously predominant third and fifth order harmonics from the square wave have now been cancelled.

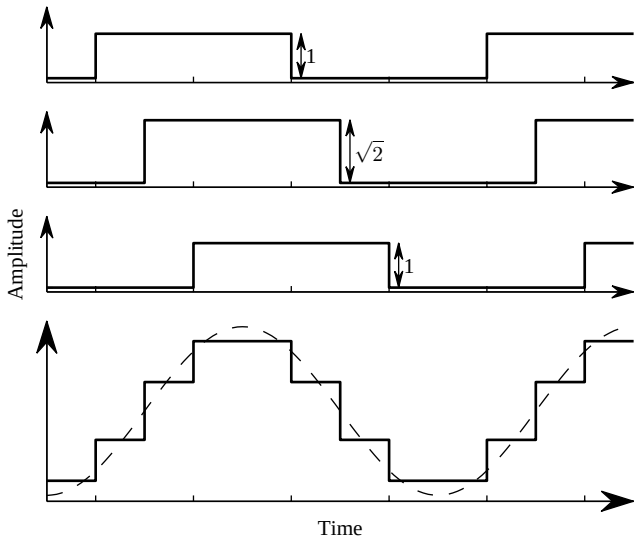

Figure 7.21: A sinusoid is approximated by summing three square waves offset by π/4 rad (45°) with amplitude ratio $1:\sqrt{2}:1$.

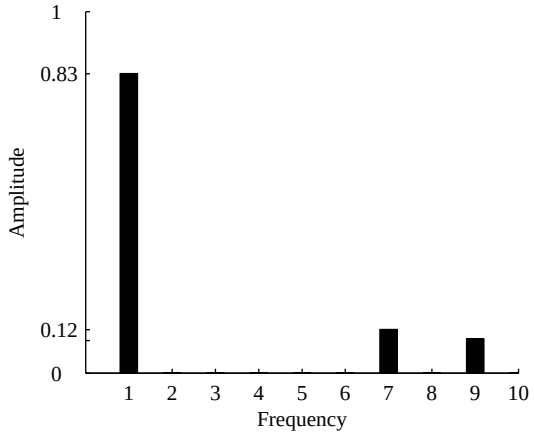

Figure 7.22: Frequency spectrum of the waveform generated in Figure 7.21.

To understand the cancellation, or "harmonic rejection", the three waveforms can be expanded using a Fourier series. Equation (7.3) gives the components up to the fifth order for each of the three waveforms.

$$\begin{aligned} f_1(t) &= \frac{4}{\sqrt{2}\pi}\left[\sin(\omega t)+\cos(\omega t)-\frac{1}{3}\left(\sin(3\omega t)-\cos(3\omega t)\right)\right. \\ &\qquad \left.-\frac{1}{5}\left(\sin(5\omega t)+\cos(5\omega t)\right)\right] \\ f_2(t) &= \frac{4}{\pi}\left[\sin(\omega t)+\frac{1}{3}\sin(3\omega t)+\frac{1}{5}\sin(5\omega t)\right] \\ f_3(t) &= \frac{4}{\sqrt{2}\pi}\left[\sin(\omega t)-\cos(\omega t)-\frac{1}{3}\left(\sin(3\omega t)+\cos(3\omega t)\right)\right. \\ &\qquad \left.-\frac{1}{5}\left(\sin(5\omega t)-\cos(5\omega t)\right)\right] \end{aligned} \tag{7.3}$$

Summing these components using the ratio $1:\sqrt{2}:1$ completely cancels the third and fifth harmonics. This approach can be used to cancel additional harmonics by increasing the number of square waves combined, where the phase shift resolution is increased, resulting in the waveform more closely approximating the desired sinusoid. For example, combining five waveforms, each offset by π/6 rad (30°), with an amplitude ratio of $1:\sqrt{3}:2:\sqrt{3}:1$, results in cancellation of odd harmonics up to the ninth order. An arbitrary number of odd order harmonics can be cancelled using this expansion technique; however in RF CMOS transmitter design this quickly becomes impractical due to the additional complexity and increase in circuit size.

The Altera FPGA is capable of generating up to six square wave signals from one PLL; hence this method could use five of the PLL outputs in this manner by configuring the required phase offsets at each output. An op-amp configured as a weighted adder could be used to combine the waveforms into the resultant analogue signal. The following electronic stages (e.g. amplifier and laser driver) would be required to be analogue (and linear) to ensure that the harmonics are not reintroduced into the system. The complexity of this approach makes it superfluous for use in the ranging system, given that an appropriately selected low pass filter could instead be used to convert the digital square wave to an analogue sinusoid by filtering out the higher order frequency components.

By leveraging the fact that each sample collected is not an instantaneous measurement, but instead is collected over the CCD integration period, the phase-shifted square waveforms can be summed temporally (Payne, 2007; Payne et al., 2008b). For the simplest case, the integration period is divided into three elements with the ratio $1:\sqrt{2}:1$ and the phase of the square illumination waveform is stepped by $\pi/4$ rad at the end of each element during the integration period. The length of each element dictates the amplitude contribution towards the captured image, resulting in an *effective* waveform over the CCD integration period as shown in Figure 7.23.

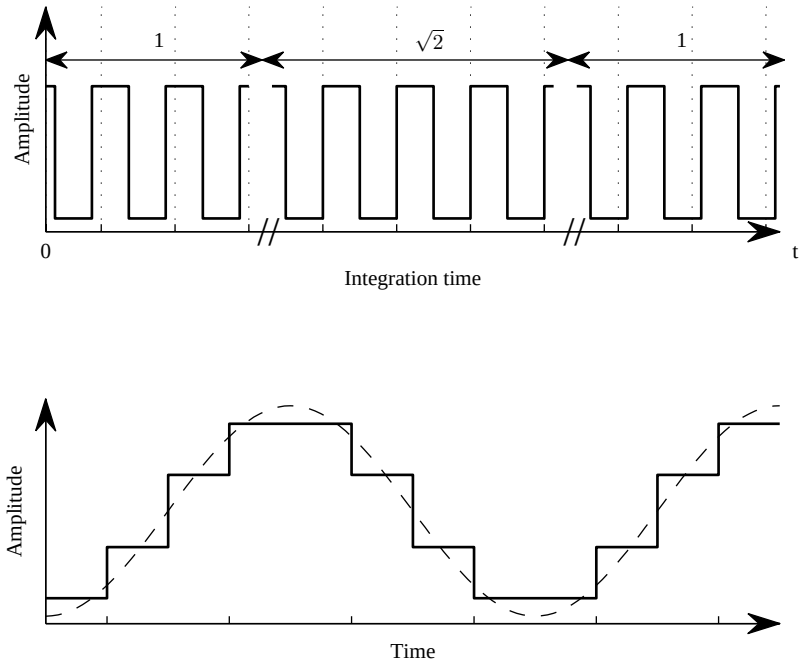

Figure 7.23: Top: The CCD integration time is divided into three elements, and the illumination waveform phase is changed by $\pi/4$ rad in each element. The integration time of each element dictates the amplitude contribution towards the final captured image. Bottom: The resultant *effective* waveform over the CCD integration period.

The advantage of this method is that square (digital) illumination and receiver modulation waveforms can be used, and the interfering harmonics are cancelled from the beat signal during acquisition. There is no longer a requirement to use analogue sinusoidal signals to avoid harmonics in the beat signal. In RF mixer design, only lower order odd harmonics are typically cancelled due to the requirement of generating and instantaneously summing together, a number of waveforms; instead one waveform is used in the range imaging system, which is phase shifted over time. Expanding the order of the cancelled harmonics is therefore relatively easy; the CCD integration period is divided into fine elements, and a small phase step between each element is used.

The cancellation technique only removes the odd order harmonics; hence even order harmonics can still be aliased onto the fundamental frequency and contaminate the range measurement. Equation (3.17) defines the relationship between the aliased harmonic components and the ratio of the sampling frequency f_{samp} to the signal frequency f_{sig}. By ensuring that f_{samp}/f_{sig} is an even integer, i.e. that an even number of samples are collected over the beat signal period, only odd order harmonics are candidates for aliasing, and range measurement errors due to the presence of even order harmonics are avoided. The phase measurement is therefore independent of the shape of the high frequency laser and image intensifier modulation waveforms, while only requiring acquisition of a minimum of four image frames.

7.5.1 Harmonic Rejection Implementation and Results

The FPGA hardware was configured in the same manner as Section 7.3.1, and a flat white target was placed in the scene. Different software designs were downloaded to the FPGA to allow comparison between the linearity of the harmonic rejection technique, and the traditional homodyne and digital heterodyne techniques. The FPGA was configured with an output frequency of 66.6̇ MHz, and an internal PLL VCO frequency of 1 GHz, producing a phase step resolution of $\pi/60$ rad. Two harmonic rejection profiles were written, the first with coarse phase steps ($\pi/4$) capable of cancelling the 3rd and 5th harmonics, and the second with fine phase steps ($\pi/60$) capable of cancelling odd harmonics up to the 117th order. To calculate the integration time for each element, a sine wave is sampled at the phase step resolution, and the amplitude difference between terms, as illustrated in Figure 7.24, is used to define the integration ratio. Only the resolution between phase steps is defined; hence an arbitrary phase offset can be applied when sampling the sinusoid. In Figure 7.24, this offset was chosen to generate a minimal number of symmetrical terms, although a different offset could equally be chosen while still rejecting the harmonics from the resultant waveform.

The corresponding phase profile over the acquisition period is shown in Figure 7.25. In comparison with those in Figure 7.18, the average phase during each captured image remains constant, but the phase profile is no longer constrained to being linear. The laser and image intensifier modulation waveforms were characterised at 66.6̇ MHz using the techniques outlined in Chapter 6, giving responses shown in Figure 7.26. The significance

of the relatively low duty cycle of the illumination waveform is outlined in Section 7.6. A beat signal was simulated from these data, and the output phase measured for various input

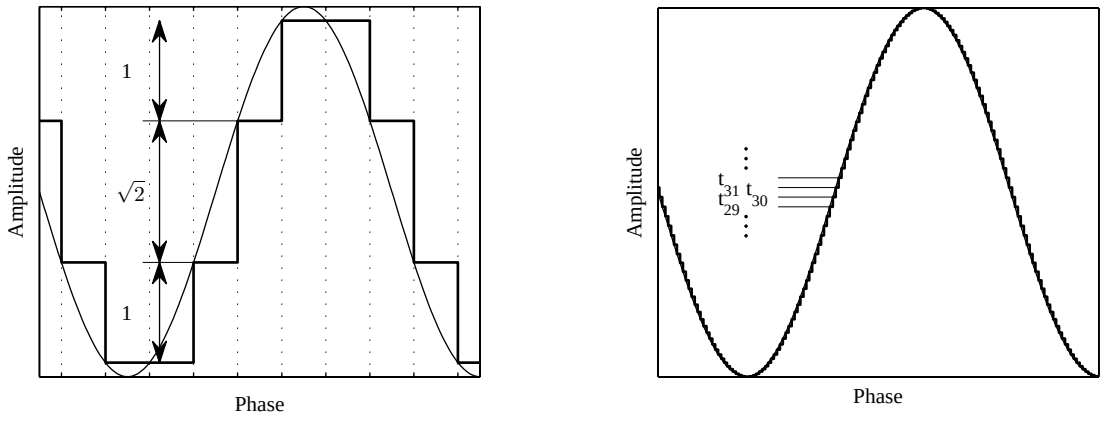

Figure 7.24: Left: A sine wave is sampled at a resolution of π/4, giving the familiar ratio $1:\sqrt{2}:1$. Right: A sine wave is sampled at a resolution of π/60, giving terms t_0–t_{58}.

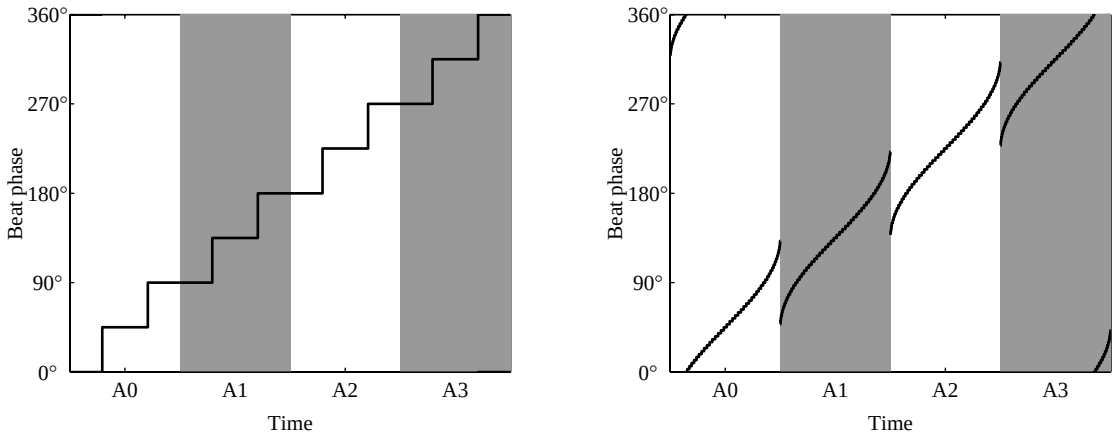

Figure 7.25: Phase profile over the acquisition period for the harmonic rejection method for π/4 phase step resolution (left) and π/60 phase step resolution (right). The acquisition period is divided into four image frames A0-A3.

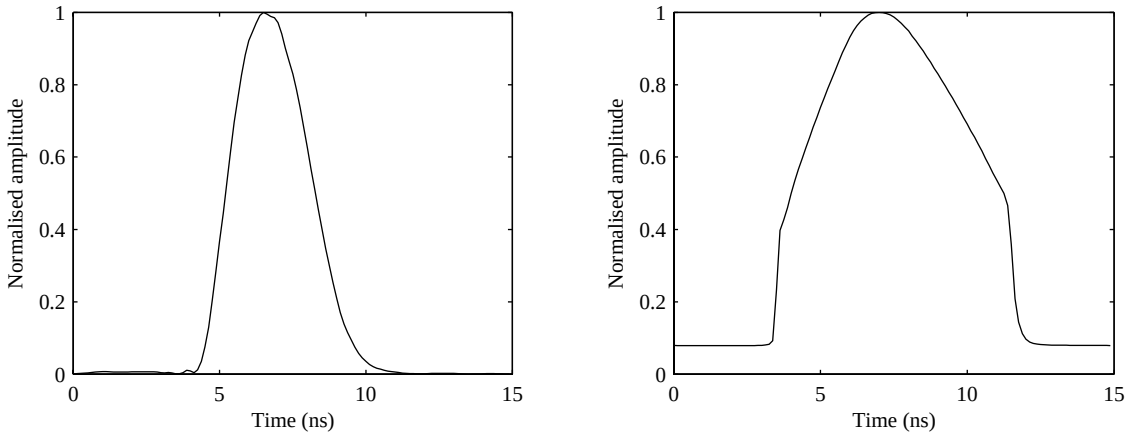

Figure 7.26: Measured response of laser (left) and image intensifier gain (right) at a modulation frequency of 66.7 MHz. From these responses, the linearity of range measurements using various techniques can be simulated.

phase angles. The phase error for each of the homodyne, heterodyne and harmonic rejection techniques were calculated, capturing four samples per beat period, and are shown in the left hand side of Figure 7.27.

Experimentally measuring the system linearity requires the distance to the flat target to be varied over the ambiguous range (2.25 m for 66.6̇ MHz). An automated tracked trolley has been used by others for this purpose (Kahlmann et al., 2006), where an interferometer gives a relative accuracy to the order of microns. Such a device was unavailable, and the combination of poor accuracy and required time to manually move the target over the required range made this an unappealing option. To overcome this limitation, the flat panel target remained stationary and range was synthetically varied by adding a time delay to the illumination signal. This approach has the advantage that the target remains in focus and in exactly the same position in the camera field of view (hence lens distortions will not affect the results). The received illumination intensity remains constant over the synthetic range as the true target distance is not changing, whereas in a real system it decreases with a $1/r^2$ relationship. The real system will therefore suffer a reduction in SNR with increasing range. A low SNR will create random noise in the measurement results, but will not affect the measurement linearity; hence the synthetic range technique is suitable for studying the effect of harmonics, maintaining constant high SNR and measurement precision over the measurement range.

Each measurement was performed by capturing four image frames (for one range measurement), followed by a single phase step to add a 125 ps time delay to the illumination signal before capturing the following four image frames. Each technique (homodyne, digital heterodyne and harmonic rejection with phase steps of $\pi/4$ and $\pi/60$) was captured five times, and an area of 5×5 pixels was analysed for each capture, with the corresponding error plotted in the right hand side of Figure 7.27. A solid line has been plotted through the mean error value.

A good agreement exists between the expected results from the simulation and the experimental results. A slight linear slope can be seen in the experimental results, and is most probably related to a temperature dependency due to heating of the image intensifier (and driver) while the system is operating. The homodyne method shows the largest error as expected, which is primarily due to the 3rd and 5th order harmonics. The heterodyne

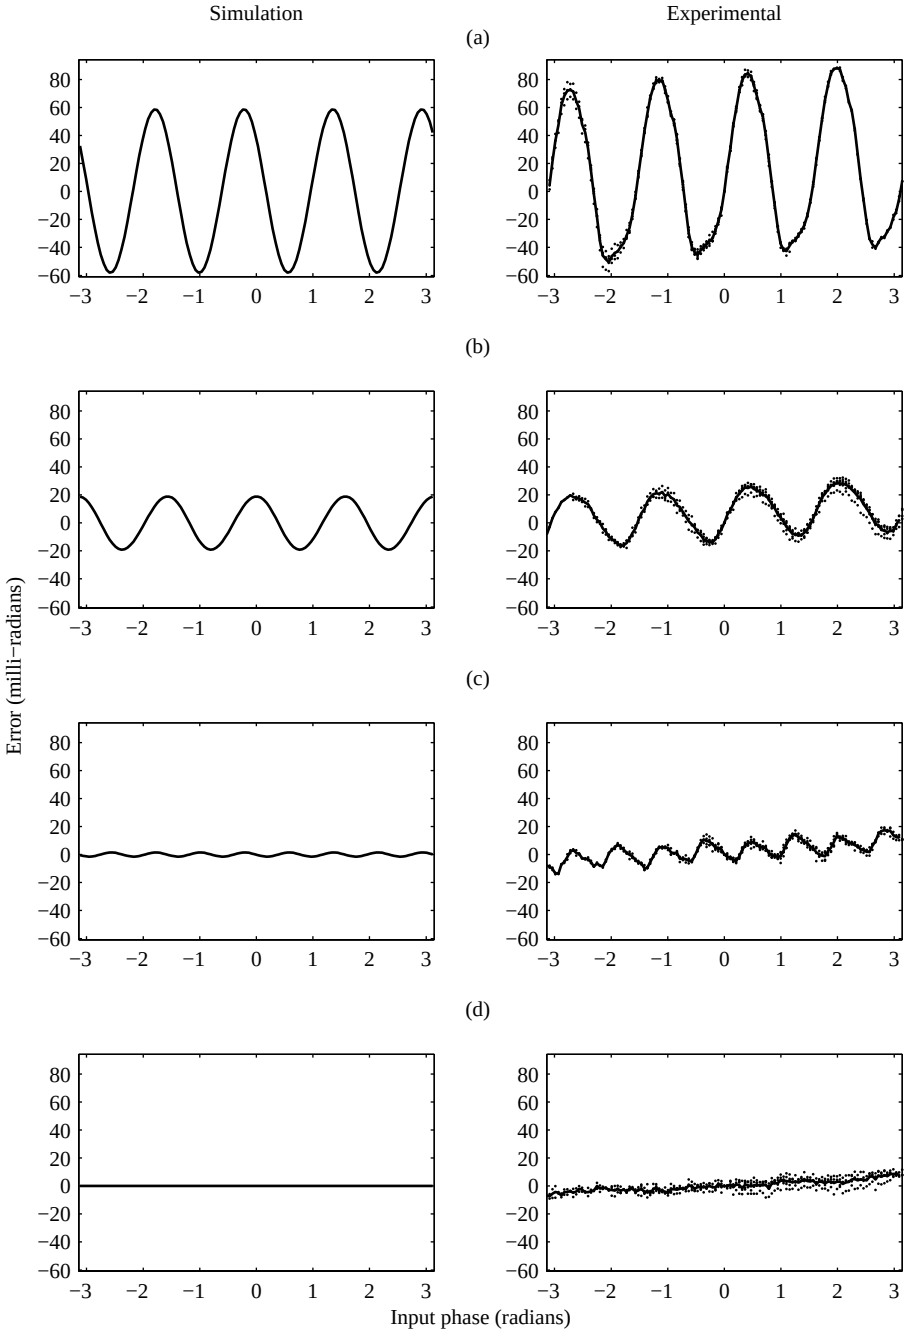

Figure 7.27: Simulation and experimental phase error of various techniques; (a) homodyne, (b) heterodyne, (c) harmonic rejection with π/4 rad phase steps, and (d) harmonic rejection with π/60 rad phase steps.

method partially attenuates these problematic harmonics, resulting in a reduction in systematic error. The first harmonic rejection profile (π/4 rad phase steps), completely cancels the 3rd and 5th order harmonics, hence the systematic error is reduced significantly. The 7th and 9th order harmonics are now the most significant contributors towards the error, and lead to the eight oscillations over the 2π input phase range (Payne et al., 2008b). The second harmonic rejection profile (π/60 rad phase steps) cancels all of the odd harmonics

present in the system, resulting in no discernable oscillation in the phase measurement. The results show that the harmonic rejection technique produces excellent range measurement linearity by cancelling the harmonics during the sampling process.

7.6 Increasing Measurement Precision

The measurement precision is dependent on the modulation index of the beat signal fundamental component (refer Section 3.1). As with the heterodyne system, the changing phase of the harmonic rejection method causes a reduction in the modulation index. The additional modulation index factor for the π/4 rad phase step method is $(2\sqrt{2})/(2+\sqrt{2}) = 0.83$, and with finer phase step resolution, such as π/60, the factor approaches π/4 = 0.79.

The range precision was measured over a number of regions using the homodyne technique, which does not suffer from a reduction of the modulation index, using the data captured in Section 7.5.1 and is plotted against the measured signal amplitude in Figure 7.28. The SNR is proportional to the signal amplitude[**]; hence from Equation (3.1) the precision is proportional to the inverse root of the amplitude shown by the solid line in Figure 7.28. The precision was analysed in each corresponding region for the harmonic rejection captured sequence (using π/60 rad phase steps), and is also plotted against the *homodyne* amplitude for each region. This allows the precision to be compared between the two techniques for each 5×5 region in the flat panel target.

The random error for the harmonic rejection method is approximately 25% worse than that of the homodyne method, and is in close agreement with a 27% increase predicted by taking the inverse of the change in modulation index (π/4). The trade-off between measurement precision and accuracy (linearity) therefore exists in the selection of the homodyne or harmonic rejection techniques.

[**] In this instance we assume that sensor noise is considerably larger than quantum noise. If a perfect sensor were used, then the SNR would be proportional to the square root of the illumination intensity.

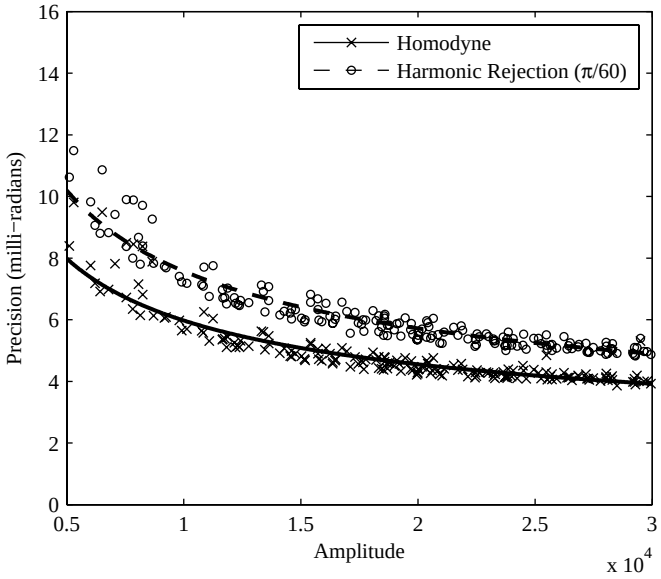

Figure 7.28: Precision comparison between the homodyne and harmonic rejection techniques.

The harmonic rejection technique has the advantage that the measurement linearity is now independent of the illumination and receiver modulation waveform shape. Aliasing from even harmonics is avoided by taking an even number of samples, and odd order harmonics are cancelled during sampling. This affords flexibility in the selection of the generated waveform shape, which can be manipulated to enhance the modulation index.

Assuming that the illumination source is a rectangular wave, the modulation index value of the fundamental frequency component is dependent on the waveform duty cycle as given by Equation (7.4). By decreasing the duty cycle, the modulation index value can be increased as illustrated in Figure 7.29. A reduction from 50% to 25% duty cycle improves the illumination waveform modulation index, hence beat signal modulation index value, by a factor of $\sqrt{2}$. A narrow pulse (duty cycle < 1%) provides the maximum possible modulation index value, twice that of a sinusoidal waveform.

$$m_{rectangle} = 2sinc(\pi \cdot duty\,cycle) \tag{7.4}$$

Reducing the duty cycle in this manner causes a significant increase in the harmonic content of the waveform, which exaggerates the linearity problem when using the homodyne technique, whereas the harmonic rejection technique alleviates this problem due to reduced harmonic sensitivity.

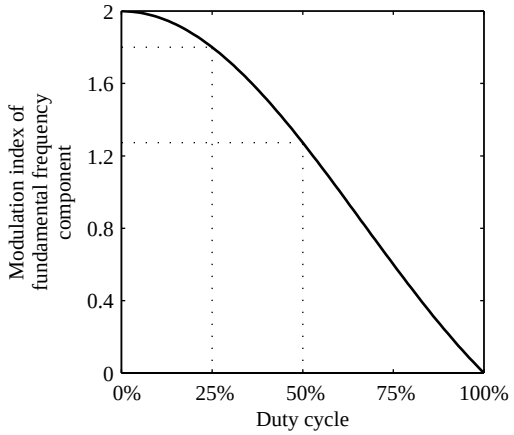

Figure 7.29: The modulation index value is dependent on the duty cycle of a rectangular wave. Reducing the duty cycle from 50% to 25% gives a $\sqrt{2}$ increase in the modulation index.

The laser driver was configured to operate with approximately 38%, 20% and 2.5% duty cycles at $66.\dot{6}$ MHz, while the image intensifier response remained constant as shown in Figure 7.26. The resultant captured beat signal is shown in Figure 7.30, with modulation index values of 0.68, 0.79, and 0.88 respectively.

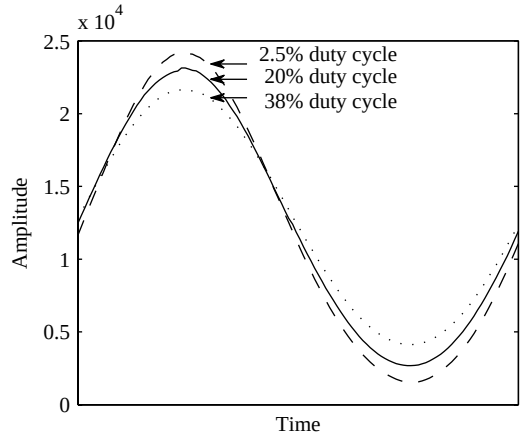

Figure 7.30: The beat signal modulation index increases when the illumination source duty cycle is reduced.

Exclusively reducing the duty cycle does not give an overall gain in measurement precision however, as the illumination intensity, hence system SNR, is also proportional to the duty cycle. The peak power must therefore be increased to maintain the same average illumination intensity, hence maintain the SNR. The maximum pulse power rating of the ML120G21 laser diodes is 160 mW, twice the CW rating of 80 mW. When operated at 50% duty cycle, the laser diodes are operated at both of these limits. To run at 25% duty cycle, the pulse power requirement is doubled to 320 mW to maintain the same average power, exceeding the device specifications. Twice as many laser diodes would therefore be required, each operating within the pulse power limit. As laser diode technology

improves this problem is likely to become less of a problem, for example HL6535MG (Hitachi Ltd., Toyko, Japan) laser diodes, are rated with a pulse power of 240 mW, and CW rating of 90 mW, allowing a duty cycle down to 37% before the average power must be reduced. In contrast LEDs are typically constrained by the average power, hence can be operated with a constant average power while the duty cycle is reduced. The major limitation with LEDs however, is their limited bandwidth; therefore a significant reduction in duty cycle may not always be possible at high modulation frequencies.

Despite these practical limitations, a modest halving of the duty cycle from 50% to 25%, requiring double the peak power and double the bandwidth, provides an increase in the modulation index by $\sqrt{2}$. Coupled with the decrease due to use of the harmonic rejection technique (a factor of $\pi/4$), the combined outcome is an increase in the modulation index of 11%. This is a significant accomplishment, as this combination provides both improved precision *and* accuracy (linearity) over the traditional homodyne system.

The duty cycle reduction can be applied to either, or both, the illumination or receiver modulation waveforms. Caution must be applied to ensure that the beat signal modulation index does not exceed one, or the beat signal will become distorted and again contain harmonics. Table 7.2 gives a selection of modulation waveform shapes and the corresponding beat signal modulation index value.

Table 7.2: Selection of waveforms and the corresponding beat signal modulation index and harmonic content.

Waveform 1	Waveform 2	Technique	Beat signal modulation index	Harmonics
Sine	Sine	Homodyne	0.50	None
Sine	Square	Homodyne	0.64	None
Square	Square	Homodyne	0.81	Present
Square	Square	Harmonic rejection	0.64	None
Square	25% duty cycle	Harmonic rejection	0.90	None
Square	1% duty cycle	Harmonic rejection	1.00	None

7.7 Summary

This chapter investigated the possibility of creating a “digital heterodyne” system, removing the requirement of high frequency analogue signals generated by the DDS and corresponding analogue filters and amplifiers. The signals were instead generated using an FPGA, providing a compact and efficient digital output. A simulation showed that the analogue and digital systems function identically if the digital phase step resolution is

sufficiently high, and experimental results confirmed this by producing comparable precision in measurements from each system.

The additional flexibility and control afforded by use of the FPGA over the DDS system allowed other phase modulation techniques to be investigated. Extending the linear intra-frame phase shift allowed cancellation of the third harmonic, significantly improving the range measurement linearity over a homodyne or heterodyne system. By recognising that the FPGA is not constrained to producing a linear intra-frame phase shift over time, a novel phase modulation technique was developed, based on the theory from CMOS RF harmonic rejection mixers. The harmonic rejection technique is capable of cancelling all odd harmonics in the beat signal, and by ensuring an even number of samples are collected during acquisition, no harmonics are aliased to the signal of interest. The range measurement linearity then becomes independent of the shape of the high frequency modulation waveforms. The modulation scheme was implemented by a simple software change within the FPGA, and the experimental results showed good agreement with the simulation.

As the dependency on the modulation waveform shape is removed using the harmonic rejection technique, alteration of the illumination waveform was investigated to improve the beat signal modulation index, hence range precision. Reducing the signal duty cycle, while simultaneously maintaining a constant average power, increases the modulation index value of the fundamental frequency component, improving the range measurement precision. This approach is not appropriate for use with the homodyne or heterodyne systems, as the harmonic content also increases with a change in duty cycle causing a reduction in measurement accuracy (linearity).

Combining the harmonic rejection technique, and the change in duty cycle, allows a digital system to be constructed which is capable of outperforming the traditional homodyne or heterodyne system, producing both higher measurement precision and linearity. Arbitrary illumination and receiver modulation waveforms can be used, alleviating problems such as the image intensifier non-linear transfer function, while also allowing simple switched (digital) circuits to be used. Use of the FPGA potentially brings other advantages such as spread spectrum clocking of the modulation signals, and integration of the control signals, including gating of the image intensifier, requiring less external logic. The potential also

exists to process the data from the CCD within the FPGA itself, affording a significant reduction in PC computation which can be used for other tasks.

8. Evaluation and Conclusion

8.1 Range Measurement Examples

To demonstrate the performance of the range imaging system under different operating conditions, a number of scenes were acquired in either a high precision or a high speed mode of operation, and are illustrated Figures 8.1 to 8.8. The first example, a bowl of fruit, shown in the photograph of Figure 8.1, is imaged at an approximate distance of 1.5 m with the system configured for a high precision measurement, using an 80 MHz modulation frequency, a beat frequency of 1 Hz, a camera frame rate of 29 fps, and a total acquisition period of 10 s. The camera is configured to operate in a 2×2 pixel binning mode, providing a resolution of 512×512 pixels. The resultant intensity image is shown in the right hand side of Figure 8.1, and the captured range image is illustrated in Figure 8.2 using a colour scale to represent the distance from the camera. The range image clearly illustrates the three dimensional shape of each piece of fruit, and individual grapes can be identified from within the bunch. The pear and apple appear to be beside one another in the photograph (and intensity image); however, the range image clearly shows that the pear was placed approximately 2 cm further from the camera than the apple. It is worth noting that the captured intensity and range figures effectively illustrate the raw captured data, where image processing has been restricted to a gamma correction and a simple truncation of the range data beyond 1.65 m for clarity to the reader. Further image processing techniques, such as smoothing and edge detection, could be applied to the acquired data to reduce measurement noise as deemed necessary by the intended application.

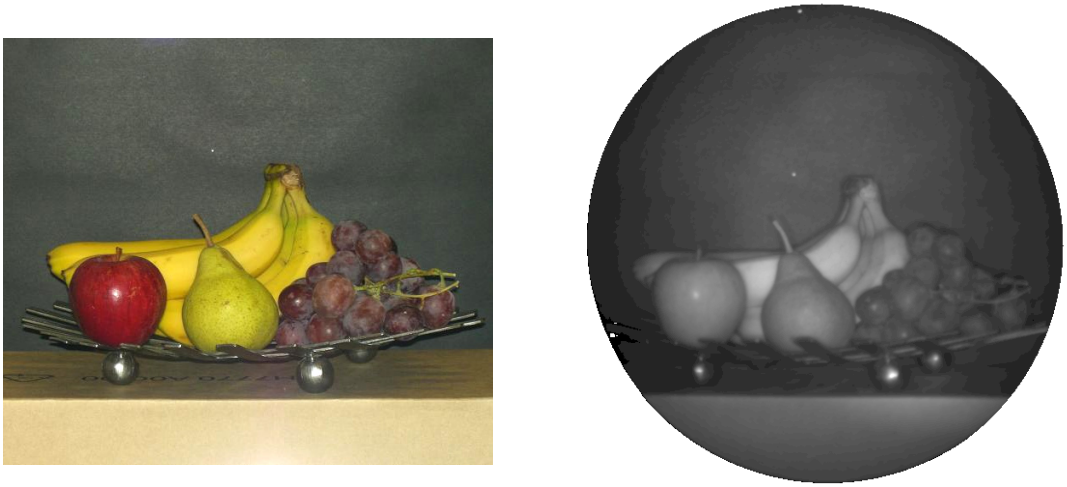

Figure 8.1: Photograph (left) and acquired intensity image (right) of a bowl of fruit.

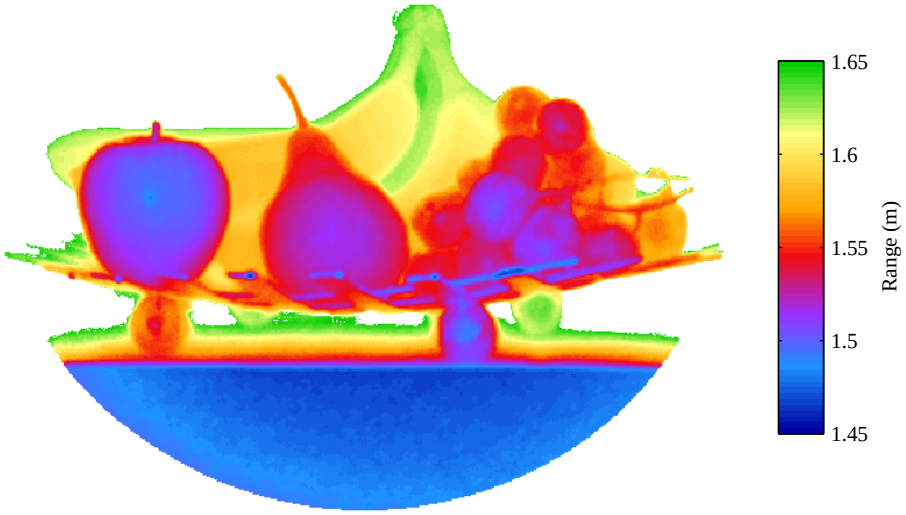

Figure 8.2: (Colour) Range image of a bowl of fruit. Each pixel within the image is artificially coloured to depict range.

The second example is 'Stumpy' – a garden gnome. The modulation frequency used was 75 MHz, and other configuration parameters were identical to that used to capture the fruit bowl. Again, a photograph and captured intensity image are provided in Figure 8.3, and a colour coded range image is illustrated in Figure 8.4, where distances beyond 1.95 m have been truncated for clarity. A 3D reconstruction of Stumpy is provided in Figure 8.5, where range data have been employed to generate a 3D surface, which is then overlaid with intensity data. Once again fine detail, such as facial features, and deviations within Stumpy's beard, can clearly be seen within the range images.

The final example is a moving person within an indoor scene, where rapid acquisition is vital to minimise measurement errors due to target motion. The modulation frequency

Figure 8.3: Photograph (left) and acquired intensity image (right) of 'Stumpy' the garden gnome.

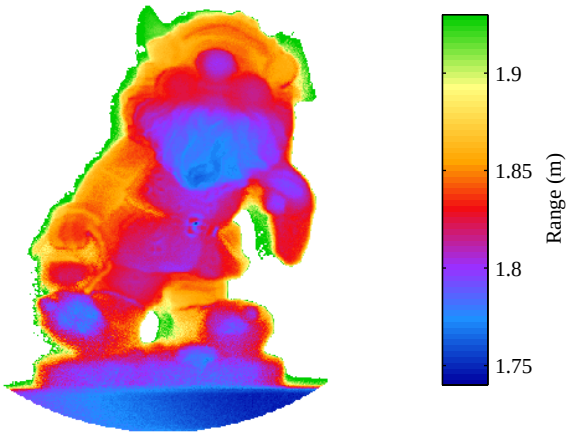

Figure 8.4: (Colour) Range image of Stumpy. Each pixel within the image is artificially coloured to depict range.

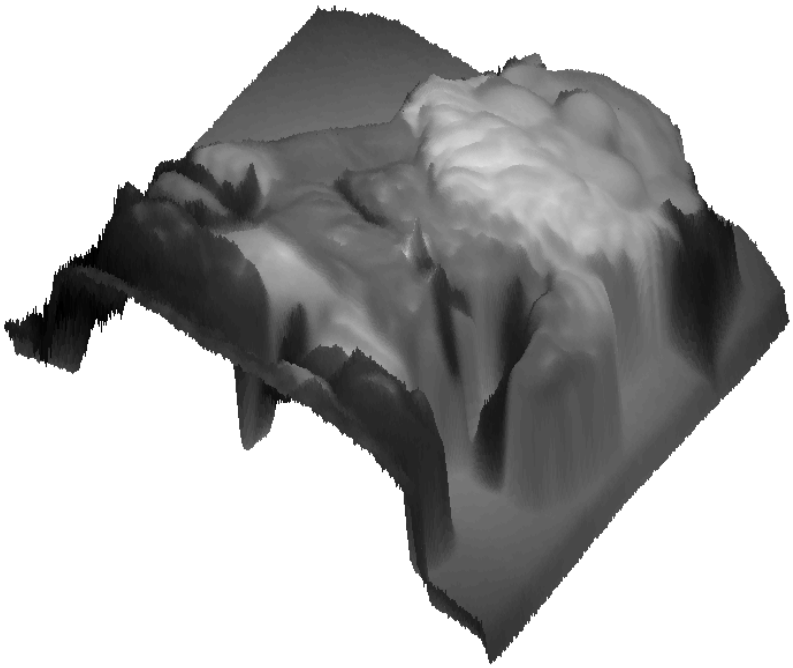

Figure 8.5: A 3D reconstruction of Stumpy overlaid with intensity information.

used was 80 MHz, the beat frequency 25 Hz, and the camera frame rate 100 fps. The increased beat frequency and camera frame rate allow range and intensity data to calculated from series of four captured frames, producing an independent measurement every 40 ms (i.e. updating at 25 fps) at a resolution of 512×512 pixels. A ten second capture was performed, providing 250 independent intensity and range images of the moving person. A selection of eight of the captured images is shown in Figures 8.7 and 8.8, illustrating various different poses.

A number of observations can be made from the captured motion sequence. Despite the short acquisition period, the range imaging system provides high quality depth and intensity information. The demonstrated sequence illustrates the inherent potential for applications where a human pose must be estimated with high precision, for example, as an input device for a video game system. Detail, such as objects and shelving in the background, can easily be identified in the range measurements. Range precision was measured for the lower half of the door at the regions shown in Figure 8.7 (g), as the door provides an unobstructed, flat object which can be fitted to a reference plane in the same manner as that used in Figures 6.19 and 7.15. The mean one-sigma precision is 18.2 cm over the selected region, or approximately 3% of the range value. As precision is dependent on the intensity of the reflected illumination, regions such as the torso, which provide almost double the signal amplitude, will result in significantly smaller errors.

The 80 MHz modulation frequency provides an unambiguous range of 1.875 m (refer Section 3.2.1), which in this instance is less than the depth of the scene, resulting in identical measured phase values for objects at different distances. For example, the 3.9 m range to the outstretched hands in Figure 8.8 (c) and the 5.8 m range to the shelving to the left of the person, both result in a phase value approximately equal to π (shown as a dark red colour), hence the system is unable to distinguish between these two ranges. The 80 MHz modulation frequency was selected for this example to provide enhanced measurement precision compared to that achieved using a lower frequency that is capable of imaging the scene unambiguously. This result highlights the need to combine phase measurements from multiple modulation frequencies to achieve unambiguous range measurements with high precision, refer Section 3.2.1. The current hardware can only

perform such measurements on static scenes, where the two modulation frequencies are used sequentially.

The example scenes in Figures 8.1 to 8.8 include a selection of different surface textures and colours at various angles of incidence, resulting in a wide variation of reflected intensity. As with all optical measurement systems, the intensity of the received light affects the measurement SNR, hence precision. The examples show that the constructed system is, in most instances, capable of producing quality range measurements despite the variations within the scene. The MCP gain must be carefully selected to avoid sensor saturation, particularly in the presence of specular reflection such as that seen on the front of the apple in Figure 8.2 and Stumpy's belt in Figure 8.3, while also providing high gain for the remainder of the scene. In these examples the high localised reflectivity causes a slight discrepancy in the resultant range measurement. One possible solution to this issue is to capture the scene twice using low and high gain settings as suggested in Section 5.2.4, allowing bright and dull regions respectively to be meaured without concern. In Figure 8.8, the person's black hair produces an extremely weak signal, where the amplitude is approximately 3% of that in brighter regions (such as the torso), resulting in increased noise, particularly at the top of the head. One solution to overcome this poor SNR is to extend the camera integration period, however, the requirement of a rapid update rate in this application precludes the use of this approach. Alternative methods to improve SNR include increasing the illumination power, reducing ambient lighting at the wavelength used by the system (the three example scenes shown were acquired in the presence of standard fluorescent room lighting), or increasing pixel binning within the

Figure 8.6: Photograph of a moving person within an indoor scene.

Figure 8.7: Selected intensity images of a moving person captured at a video update rate (25 fps). Corresponding range data are shown in Figure 8.8. Range precision is measured for the flat, bare region of the door in (g) for 25×25 pixel segments illustrated by the grid pattern.

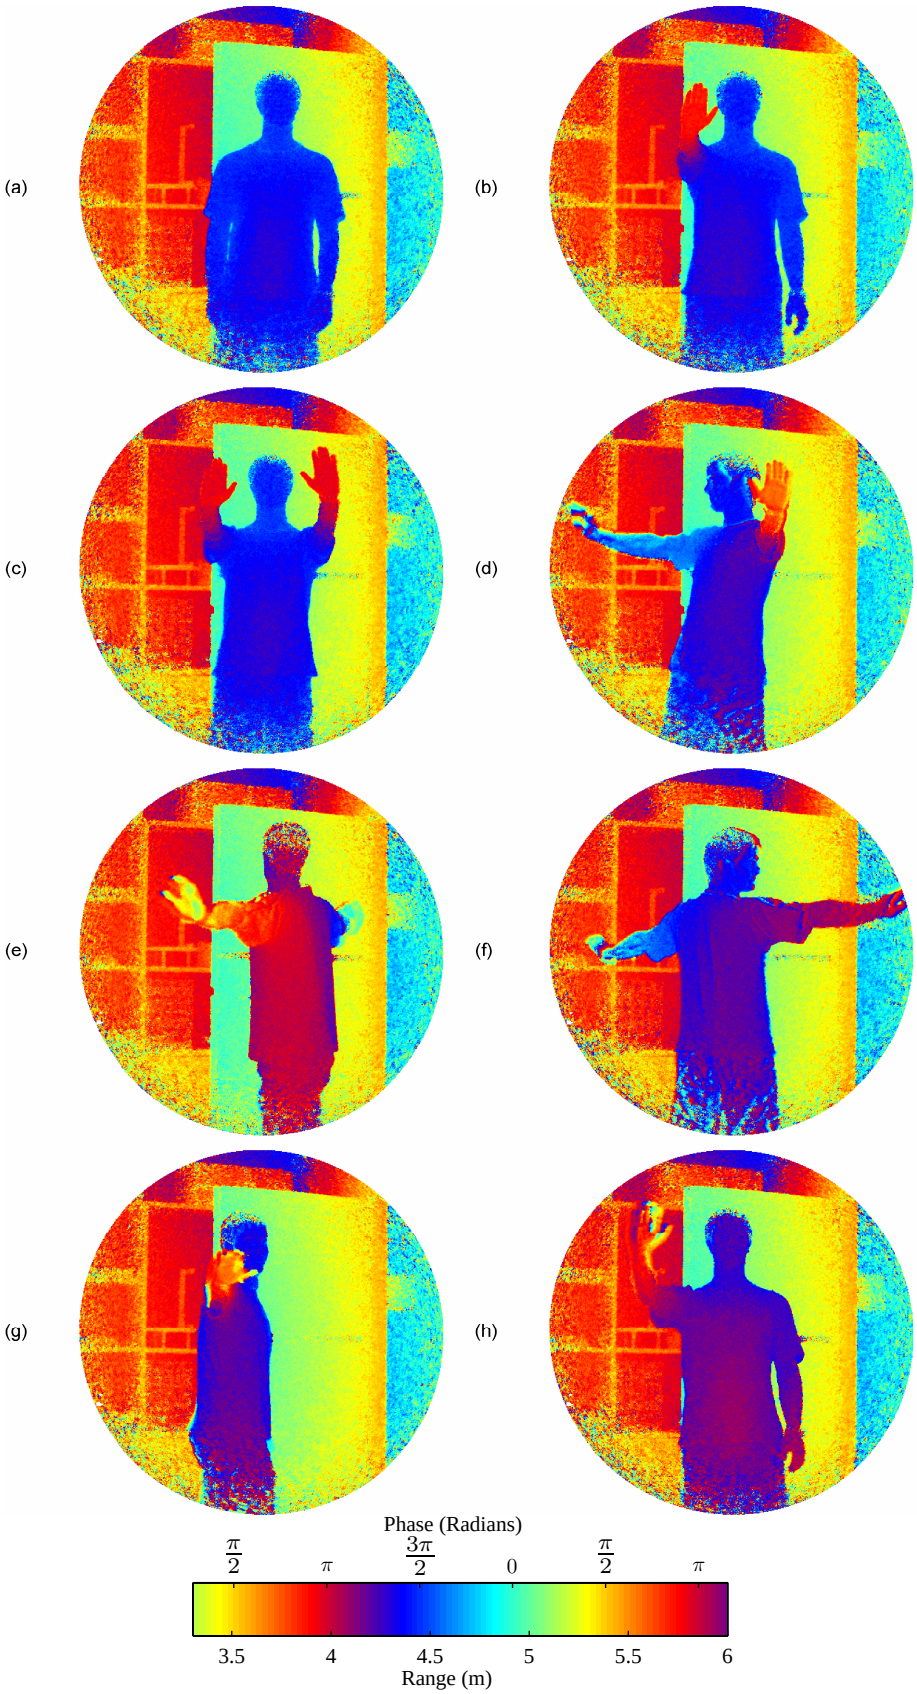

Figure 8.8: (Colour) Range images of a moving person, captured simultaneously with the intensity images in Figure 8.7 at a video update rate (25 fps). Each pixel within the image is artificially coloured to depict range. Measured precision is 18.2 cm for the lower half of the door in the regions illustrated in Figure 8.7 (g).

CCD (at the expense of reducing spatial resolution). The system employs a 645 nm long pass filter, which combined with the photocathode response rejects most ambient lighting, however, fluorescent lighting produces some spectral output beyond this 645 nm cutoff (Kawakita et al., 2004c). The outer edges of the range measurements in Figure 8.8 also tend to exhibit increased noise, primarily due to the use of a non-uniform illumination profile that is suited to a narrower field of view such as those demonstrated in Figures 8.1 and 8.3. This problem can therefore be reduced by optimising the illumination optics to match the wider scene.

Few motion artefacts are visible in the captured range images, with the exception of the rapidly waving hand in Figure 8.8 (h). Motion artefacts such as this, could be identified and minimised using software processing techniques. Alternatively, by increasing pixel binning within the CCD camera (at the expense of spatial image resolution), the camera frame rate can be increased, reducing the acquisition period and the corresponding motion blur from the captured range image.

The examples shown successfully demonstrate the operation of the range imaging system in high precision and high speed modes of operation, producing range and intensity data with high spatial resolution.

8.2 Summary

A full-field, time of flight image ranging system, or "3D camera" has been developed, capable of adapting to the requirements of various existing, and future, applications.

8.2.1 Literature Review

This thesis firstly reviews the state-of-the-art range measurement technology available at the outset of the research, and different full-field range measurement techniques are discussed and evaluated. Interferometric techniques are suited for measuring shorter distances (μm–cm) than that required for the applications of interest here (typically 1−100 m). Other techniques using cues such as shading and focusing provide limited performance with respect to range estimates. Geometric techniques, both active (structured lighting) and passive (photogrammetry), suffer from occlusion problems, an inverse range accuracy dependency, and require a relatively large baseline to measure

distances over large ranges. A time of flight system is preferable for the applications of interest here, where the propagation delay of light travelling to and from a target is measured to determine range.

To minimise the required acquisition time, a full field system is preferred; capable of determining range simultaneously for every pixel within a scene. Typically, full field direct time of flight measurement systems require complicated electronic design, consisting of a large array of high speed optical receivers and high speed digital counters, where the range resolution is limited by the operating speed of the design. Various indirect time of flight methods have been devised that reduce the electronic speed requirements by encoding flight time into an alternate variable such as intensity, phase, or frequency. Demonstrated systems using these techniques are reviewed in some depth, where limitations are identified, and the range measurement performance is evaluated. The heterodyne system, which encodes range into the phase of a low frequency beat signal, is favoured as it is provides improved immunity to the affects of ambient lighting, sensor quantisation and dynamic range limitations compared to intensity encoding systems. It can also operate with relatively few captured images to provide a short acquisition period. A heterodyne range imaging system is reviewed that achieves centimetre range resolution (Carnegie et al., 2005).

8.2.2 System Analysis and Enhancements Achieved

The objective of the work presented here is to enhance this existing heterodyne range imaging system into a flexible, high performance range measurement device for a variety of applications capable of acquiring high spatial resolution range measurements, with millimetre range resolution, or alternatively operating at rapid video update rates.

Parameters dictating the measurement precision and accuracy of the heterodyne range imaging system are examined and analysed. Measurement precision was found to be dependent on the modulation frequency used, the modulation index value, and the SNR; hence it is important to maximise these parameters where possible. The object distance is encoded into the phase of a sampled signal, hence range ambiguity problems can occur due to phase being cyclic, and aliased harmonics can contaminate the phase measurement producing a nonlinear measurement. Techniques to overcome these limitations were

examined, such as using multiple modulation frequencies, and acquiring a large number of samples over the beat signal period.

To enhance the performance of the existing range imaging system, the hardware components were redesigned to permit higher modulation frequencies to be used. A discrete digital synthesiser (DDS) was constructed with three synchronised, frequency-locked outputs, capable of generating frequencies in excess of 100 MHz with 0.09 Hz tuning. The DDS controls the illumination and shutter modulation, while a third output provides a synchronised reference enabling the acquisition of absolute range measurements. An image intensifier is used as the optical shutter element, where the photocathode voltage controls the optical gain of the device. The operating mechanism of the image intensifier is explained, and the photocathode modulation voltage investigated in terms of gain linearity, spatial resolution and power dissipation. A driver circuit was constructed to output 50 V_{pp} at frequencies up to 100 MHz, where the modulation voltage and modulation frequency capabilities were increased compared to the original system, providing enhanced spatial resolution and measurement SNR.

A laser illumination circuit was designed and constructed to replace the existing LED light source, which was not capable of being modulated at the required high frequencies. The laser circuit produces a combined average optical output power of 320 mW, at frequencies beyond 100 MHz, from four laser diodes. The light is coupled through optic fibres placed equidistantly around the camera lens providing, on average, coaxial illumination to avoid shadowing. An FPGA was introduced into the system design, used to allow synchronous triggering of the camera with minimal jitter, removal of image smear by disabling the image intensifier during the CCD image transfer period, and control over the gain of the lasers and image intensifier.

Each of the electronic systems were characterised to verify and optimise their operation. The DDS output was measured using a spectrum analyser, and found to contain a number of unwanted spurious frequency components. Spurious frequency components can create jitter in the timing of the laser and image intensifier modulation waveforms; therefore an external low pass filter was added to the DDS to attenuate unwanted frequencies. The modulated laser output was measured with a photodiode, and was found to contain a large amount of ringing on the rising edge, known as a relaxation oscillation. This ringing is

undesirable as it causes the peak output power to exceed the manufacturer rating for the laser diode, and also adds harmonics to the modulation waveform. The oscillation was critically damped by adding a capacitor to the driver circuit, decreasing the transition speed of the electrical drive current, while maintaining the ability to modulate the laser at frequencies up to 100 MHz.

An optical method of characterising the image intensifier modulation response was developed, replacing the illumination source with a pulsed laser. A circuit was designed to generate pulses with a FWHM of 266 ps using the previously problematic laser relaxation oscillation to illuminate the image intensifier input window. The laser pulses act as a temporal probe, allowing the image intensifier optical gain to be measured both temporally and spatially. Electrical ringing, particularly at low frequencies (<10 MHz) caused a far from desirable modulation waveform shape, leading to the redesign of the driver PCB. The redesigned PCB reduced stray inductance (by incorporated the high voltage components previously on a separate PCB, thereby reducing interconnect length), resulting in a significantly improved modulation response. Irising, where the modulation voltage is delayed near the centre of the image intensifier, was observed and measured to be approximately 150 ps, resulting in a range measurement error of ~20 mm. This effect is unavoidable due to the physical design of the image intensifier, but can be corrected in the processing software by applying a calibration generated using data from the image intensifier characterisation.

Range measurement performance was determined by imaging flat objects at various distances from the camera, at a variety of modulation frequencies. The measurement precision was determined by fitting a plane to a small region of pixels within each object, and measuring the standard deviation of the residuals. At a 90 MHz modulation frequency, the precision values are between 0.43 and 0.95 mm over the range of 1.7 to 5.2 m, meeting the high measurement precision objective. The linearity of the range measurements matches that expected, however, due to the presence of uncalibrated radial lens distortions and other sources of error, a precise linearity measurement is complicated and outside the scope of this thesis.

Jitter was found in the high frequency modulation signals, but further analysis determined that it does not have a negative influence on range measurements in some instances, and

can actually be beneficial. By envisioning a digital modulation system as being equivalent to an analogue system with large systematic jitter, a FPGA successfully utilised digital phase stepping to emulate the continuous phase shift of the analogue heterodyne system. Further advanced phase stepping techniques were investigated, resulting in a novel (and patent pending) method of rejecting harmonics from the captured signal. This technique allows arbitrary laser and image intensifier modulation waveforms to be utilised without compromising range measurement linearity. As the range can be accurately measured from an acquisition with a minimum of four image frames using this technique, the system is capable of achieving range measurements at rapid video update rates.

Further analysis also showed that it is possible to increase measurement precision by reducing the laser duty cycle, while maintaining the average output power. This increases the modulation index of the modulated illumination waveform; but also increases the harmonic content. However, as the digital phase stepping technique is insensitive to modulation waveform harmonics, the increased modulation index results in improved range measurement precision without compromising measurement linearity. The resultant system can therefore produce range measurements with higher accuracy and precision than the equivalent analogue heterodyne system, while the average illumination power, hence eye safety restrictions, remain unchanged.

Measurement examples included in this chapter show the performance of the improved range imaging system in real-life scenes, where objects have an assortment of colours, textures, and are at various angles of incidence. The examples demonstrate high precision operation, using a 10 s acquisition period to produce 512×512 intensity and range images with sub-millimetre depth resolution. Alternatively, high speed operation is demonstrated, again producing 512×512 intensity and range images, this time at a video update rate of 25 fps.

8.3 Future work

The work presented here exceeded the requirements for the original objectives, however a number of further enhancements could be implemented to extend the capabilities of the range imaging system, both in terms of measurement performance, and to increase the prospect for commercialisation.

To achieve high measurement accuracy, systematic errors must be calibrated and compensation performed by the processing software. A calibration for image intensifier irising can be generated from the characterisation data. However, a number of other sources of error exist, including radial lens distortions for the zoom lens imaging the scene onto the image intensifier, and the relay lens that couples the image intensifier output to the CCD. A principal point and principal distance must also be accurately determined to convert the acquired range data into a Cartesian coordinate (Dorrington et al., 2007a). The influencing parameters are liable to change, for example by altering the focal length of the zoom lens, requiring the calibration procedure to be repeated. A simple, rapid calibration routine is therefore required, extending traditional 2D calibration methods to correct distortions in 3D space.

The current hardware implementation only has the provision to modulate the laser and image intensifier with a single modulation frequency; hence to achieve high precision measurements while avoiding the occurrence of range ambiguity, two measurements must be sequentially acquired at dissimilar modulation frequencies, and combined during software processing. The current hardware therefore prohibits high speed applications from using this method, resulting in either the use of a lower modulation frequency, hence poor range precision, or the presence of range ambiguity in the measurement results, such as that illustrated in Figure 8.8. To overcome this limitation, two modulation frequencies can be used simultaneously by selecting different beat frequencies (refer Section 3.2.1). The drawback to this approach, however, is that a minimum of five image frames must be captured to resolve the two phase values (or six frames if harmonic rejection is implemented), increasing the measurement acquisition time.

The FPGA responsible for generating the modulation waveforms in Chapter 7 is largely underutilised, with surplus resources available which have the capacity to process the range data in hardware (refer Section 7.7), rather than in software on the PC. Obtaining raw data from the CCD camera is valuable during development of the ranging system as it can be used for characterisation and testing purposes, however, only the computed range and intensity data are required for final applications. Reducing the data processing load on the PC is highly desirable as the range camera is likely to be one of many sensors for a given application. The FPGA would ideally interface to the CCD to retrieve image data

using a CameraLink interface, process the data using hardware logic modules (multipliers, adders, etc.), while storing the interim results in RAM, and output range and intensity data to a PC using a USB, Ethernet network, or other suitable interface.

The system presented uses a lens to couple the output from the image intensifier to the CCD. Lens coupling provides high flexibility as the camera can be removed for fault diagnosis, or to be used for other purposes; however the coupling efficiency is relatively low. Alternatively, optical fibre coupling can be used to transfer the output image from the image intensifier directly to the CCD, providing an order of magnitude increase in coupling efficiency. As a result of the improved coupling, the image intensifier gain (MCP voltage) can be reduced, extending the MCP lifetime and improving the linearity of the image intensifier (Tomkins and Lyons, 1999; Princeton Instruments, 2007). Fibre coupling removes image distortions caused by the coupling lens, as well as stray light present in a lens coupled system, where light can contaminate the range measurement due to crosstalk between pixels, particularly when light from a bright region intersects a dull region in the image. The physical size of the system can also be reduced using fibre coupling, due to the minimal separation required between the image intensifier and the CCD.

A relatively large diameter image intensifier (25 mm) was selected during the design of the presented system to provide range images with high spatial resolution. However, the maximum modulation frequency of the driver designed in Section 4.3.2 is primarily limited by the capacitive loading of the image intensifier photocathode, where the capacitance is proportional to the area of the device. By decreasing the image intensifier diameter from 25 mm to 18 mm, capacitive loading is approximately halved, thereby allowing an increased modulation frequency to be used to improve the range measurement precision. Although image resolution will decrease with a smaller diameter image intensifier, this reduction can be compensated for by using an image intensifier with finer MCP channels that provide higher spatial resolution per unit area. For example, a 25 mm diameter image intensifier manufactured by Photek uses 10 μm MCP pores and can achieve 25-36 lp/mm, whereas smaller devices using 6 um MCP pores can achieve 50 lp/mm (Photek, 2003)

The visible illumination source is distracting and often undesirable when operating in the vicinity of humans, especially for security applications, or multimedia applications where a standard colour camera is to be incorporated with the range camera. The visible light

source wavelength was selected over an infrared wavelength because of the response of the image intensifier (refer Section 5.1), which has poor quantum efficiency in the invisible infrared region. Third-generation (Gen III) image intensifiers use a different photocathode material, allowing operation at visible to near infrared range wavelengths, while providing enhanced quantum efficiency. Typically Gen III image intensifiers have required the addition of an ion barrier within the design to prevent damage to the photocathode (Tomkins and Lyons, 1999). To overcome the effect of the ion barrier, a photocathode voltage approximately four times higher than that of a Gen II image intensifier is required, making high speed modulation difficult. Newer *filmless* Gen III image intensifiers no longer include the ion barrier, hence are suited for high speed modulation due to the use of lower voltages similar to our Gen II image intensifier. The requirement to use visible illumination can therefore be avoided by using a filmless Gen III image intensifier, in conjunction with an illumination source in the 800–900 nm near infrared range. To provide high measurement precision, the illumination source should be designed to operate with a short duty cycle, while maintaining high average optical power as discussed in Section 7.6.

A number of ranging systems included in Table 2.1 perform gain modulation directly within the image sensor, eliminating the need for an image intensifier (Schwarte, 2004; Büttgen et al., 2006; Hsu et al., 2006). The image intensifier is relatively expensive, noisy, fragile, bulky, and requires the use of high voltage power supplies; hence it is desirable to remove it from the system design. The modulated sensors currently available have reduced spatial resolution (120×160 and 176×144 pixels), and operate at lower modulation frequencies (20–44 MHz) than the image intensifier based range imaging system developed in this work (which typically operates at 512×512 pixels with modulation frequencies up to 90 MHz). The image intensifier based system is therefore the preferred solution for applications requiring high spatial and depth measurement resolution. However, rapid advancement of these image sensors suggests that the modulation frequency and spatial resolution may be comparable to that of image intensifiers in the near future. Modulated image intensifiers discard half of the incoming light (i.e. the duration when the shutter is closed); whereas "multi-tap" modulated image sensors can potentially capture *all* of the incoming light by temporally segmenting the light into multiple storage regions (Schwarte, 2004; Büttgen et al., 2006; Hsu et al., 2006), providing an inherent

doubling of efficiency. Such an image sensor should therefore be used in future implementations of the range imaging system to overcome the drawbacks imposed by the use of image intensifier technology, and provide increased optical efficiency.

The work performed here is directly applicable to a system using of a modulated image sensor, as many of the issues addressed are inherent to the indirect heterodyne time of flight measurement method used. The modulated image sensor temporal response can be characterised using a picosecond pulsed laser in the manner described in Section 6.3.2 to measure the modulation phase, gain, and duty cycle of each pixel within the sensor. The harmonic rejection phase stepping technique in Section 7.5 can be applied to avoid phase contamination due to the presence of harmonics, and the modulation index, hence measurement precision, can be improved by altering the illumination source as described in Section 7.6.

Multiple signal returns at a pixel site can compromise the measured range. There are numerous means that may lead to this occurrence including:

- Pixels falling on the edge location of an object. The pixel intensity is partially comprised of the object signal return, and partially from the background signal return.
- The superposition of multiple illumination sources (laser diodes). Point source illumination is assumed, although for very short measurement distances this assumption is not valid due to the use of multiple sources.
- Multiple reflections occurring within the scene.
- Alignment between the CCD pixels and the MCP channels, where multiple MCP pores contribute to the same CCD pixel.

In this instance, the measured signal intensity is the result of a weighted average of the multiple returns, with the weighting proportional to the relative contribution of each signal. The measured range may then be incorrectly determined from the captured intensity values. For example, if two objects are separated by a phase value of less than π radians (or $cT/4$ meters, where c is the speed of light and T is the modulation period), then the resultant range measurement is also the weighted average of the two distances, resulting in a relatively graceful smearing of the range values between the edges of two objects in a scene. However, when the phase separation exceeds π radians, the resultant distance

measurement does not fall between the two objects, and instead lies in front of, or behind, the pair of objects. This effect can result in large errors at object edges. The phase and amplitude of each signal return can be computed through further post processing of the captured data using either knowledge of the harmonic content of the beat waveform, or by using data captured at multiple modulation frequencies.

8.4 Conclusion

A high performance full-field time-of-flight range imaging system, or 3D camera, has been realised. The system is demonstrated to produce intensity and range measurements with sub-millimetre depth resolution and high 512×512 pixel spatial resolution. Alternatively, high speed operation is possible due to the flexible design of the system, and the system is demonstrated producing the same high 512×512 pixel resolution at 25 fps. This achievement required the design and construction of a number of electronic circuits. A frequency-locked DDS circuit was built to generate frequencies beyond 100 MHz, controlling a constructed high voltage (50 V_{pp}) image intensifier driver, and a laser diode driver circuit. The camera frame capture was synchronised to the modulation signals using a FPGA, allowing the system to generate absolute range measurements.

The DDS, laser, and image intensifier response are characterised, revealing deficiencies that were mitigated where possible through hardware modifications. A novel phase stepping technique was developed to cancel a systematic error caused by harmonics in the high frequency modulation signals. The digital technique is demonstrated using an FPGA, significantly reducing the electronic design complexity and enhancing flexibility. Range measurement precision and accuracy can be enhanced by combining the phase stepping technique with a modification to the laser illumination duty cycle. The objectives of the project have been exceeded, resulting in the design of a flexible system suitable for a wide range of existing and future applications.

Appendix A. PCB Schematics

A.1 Direct Digital Synthesiser – Power

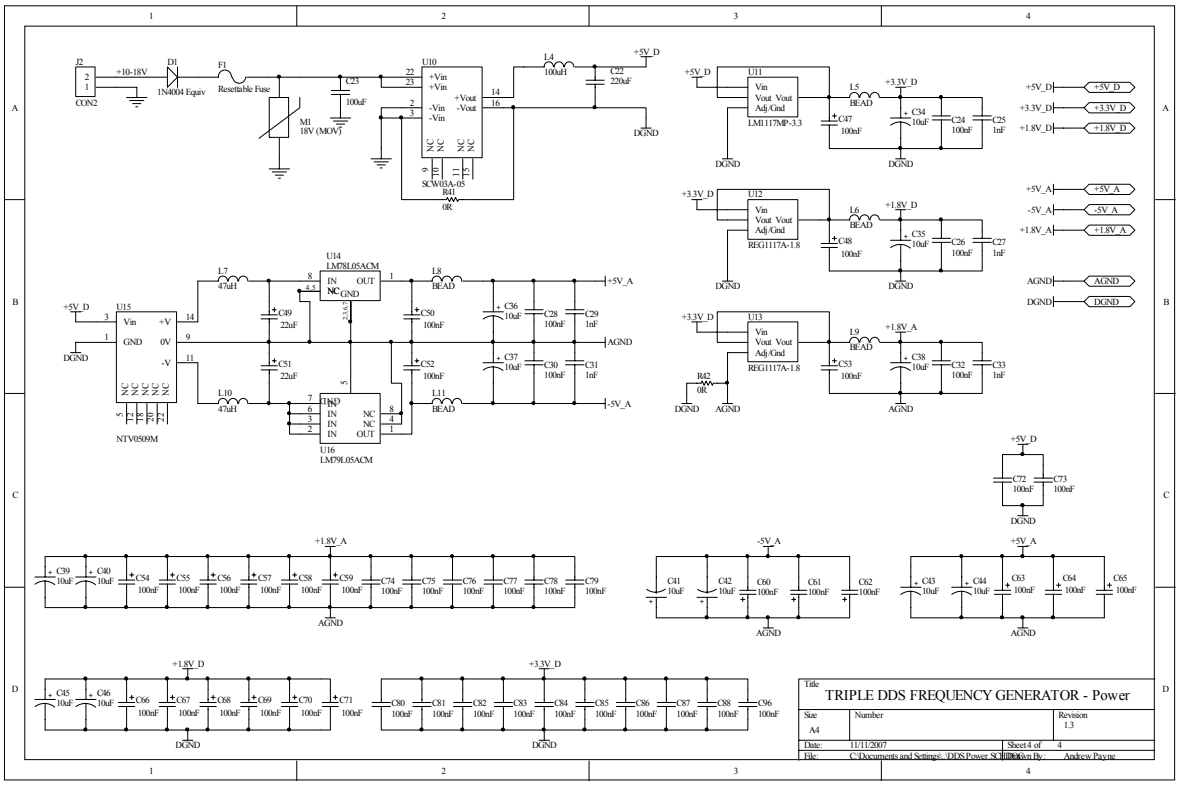

A.2 Direct Digital Synthesiser - Digital

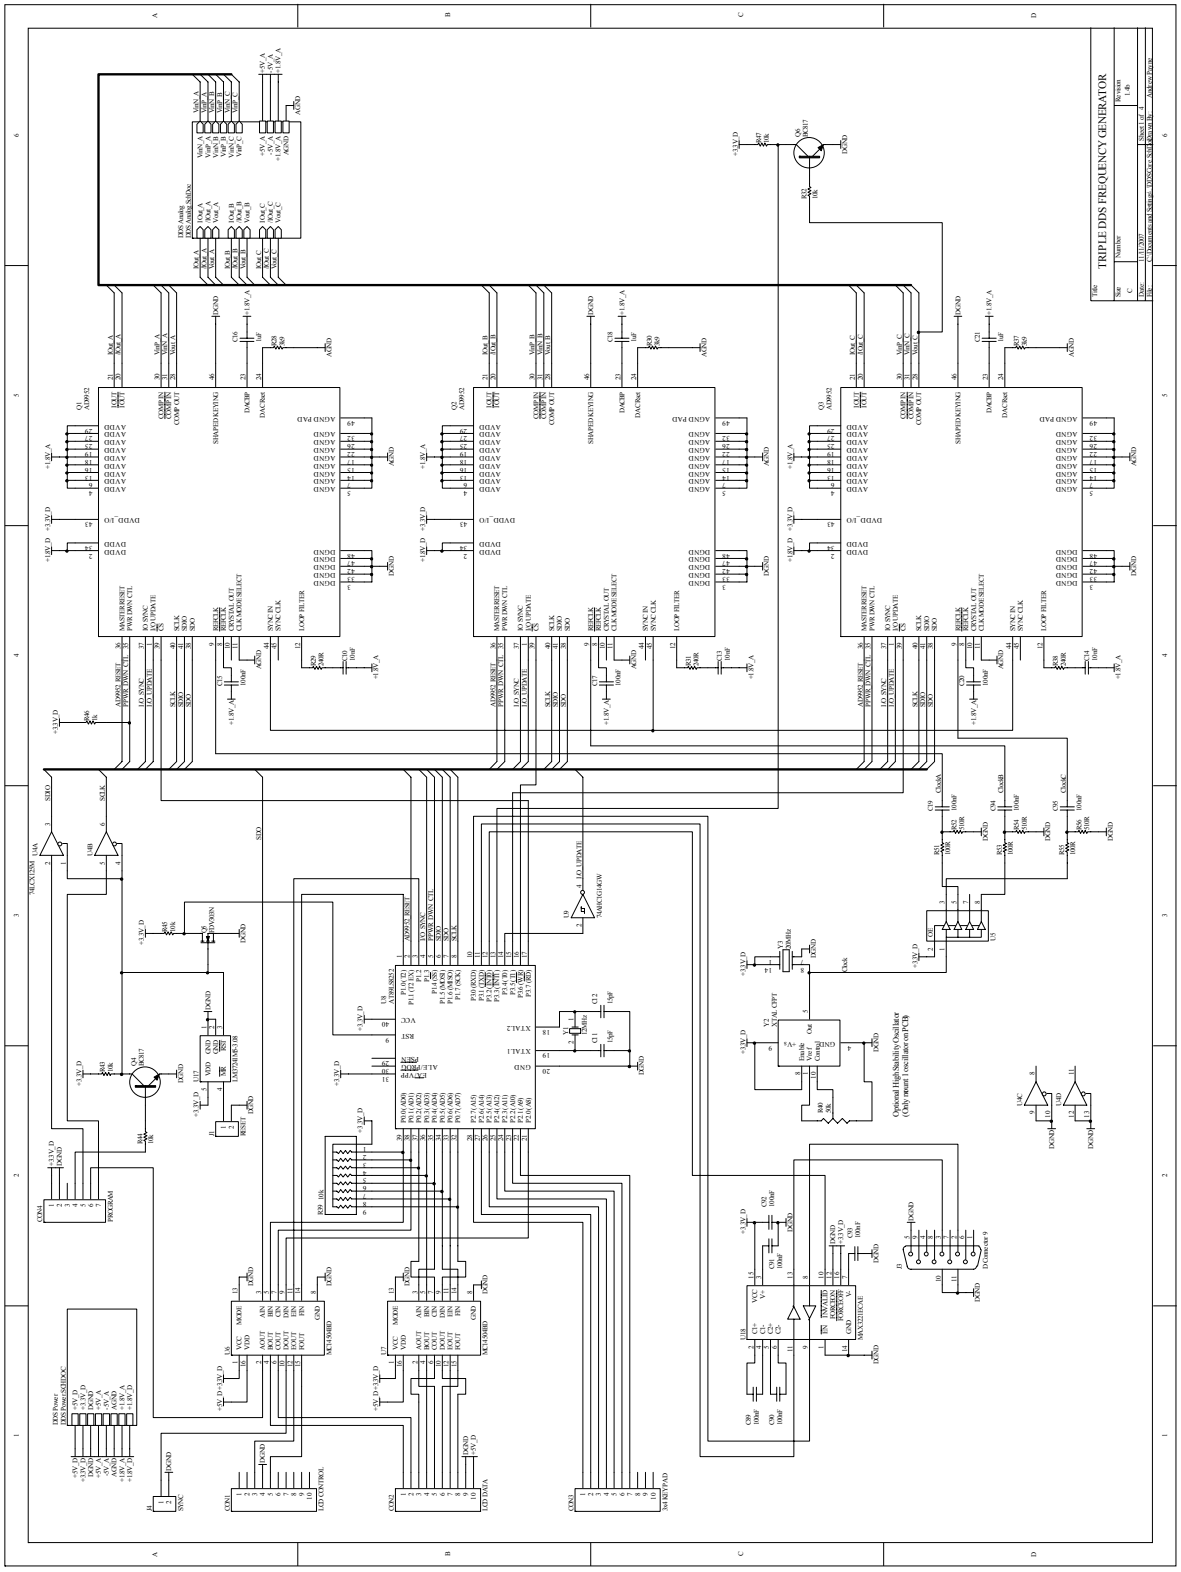

A.3 Direct Digital Synthesiser - Analogue

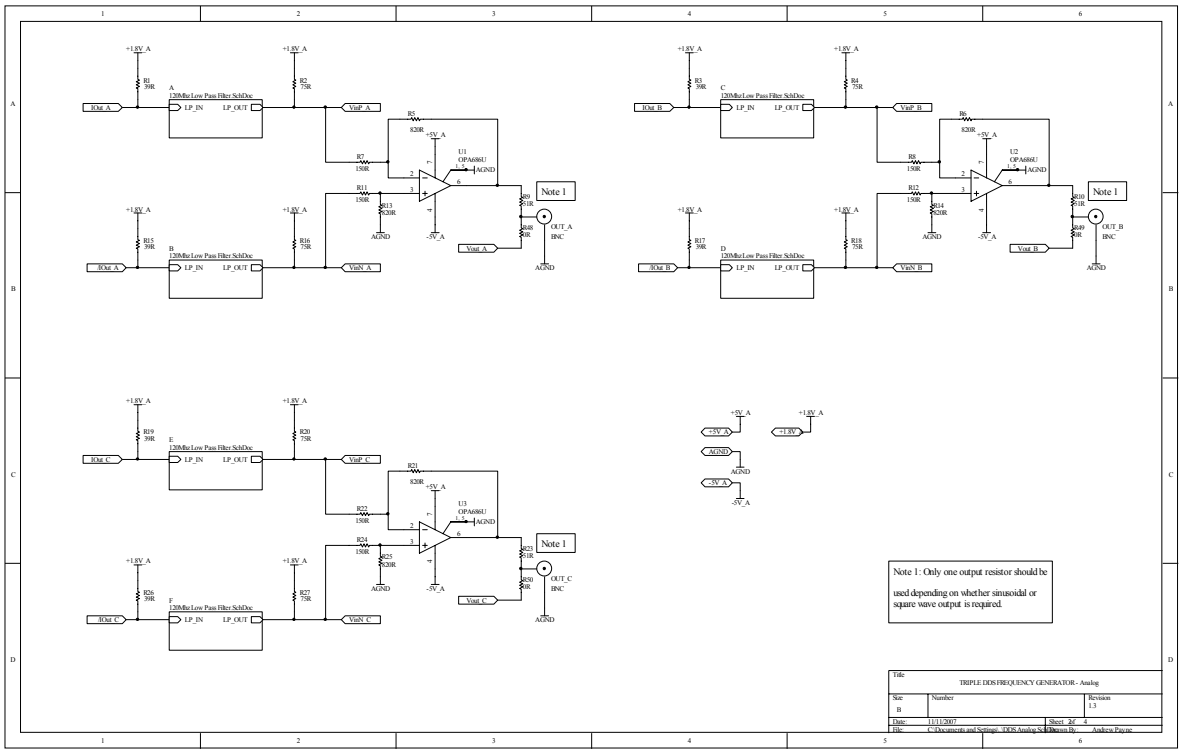

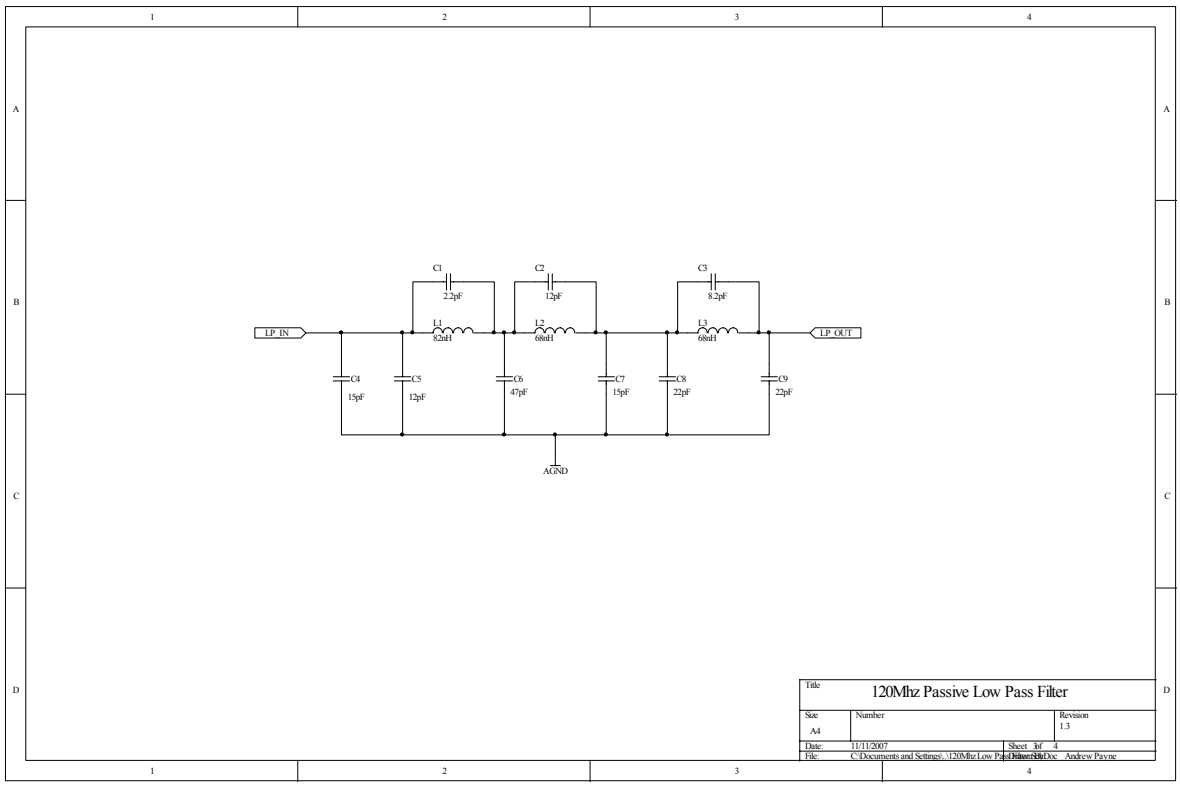

A.4 Direct Digital Synthesiser – Cable

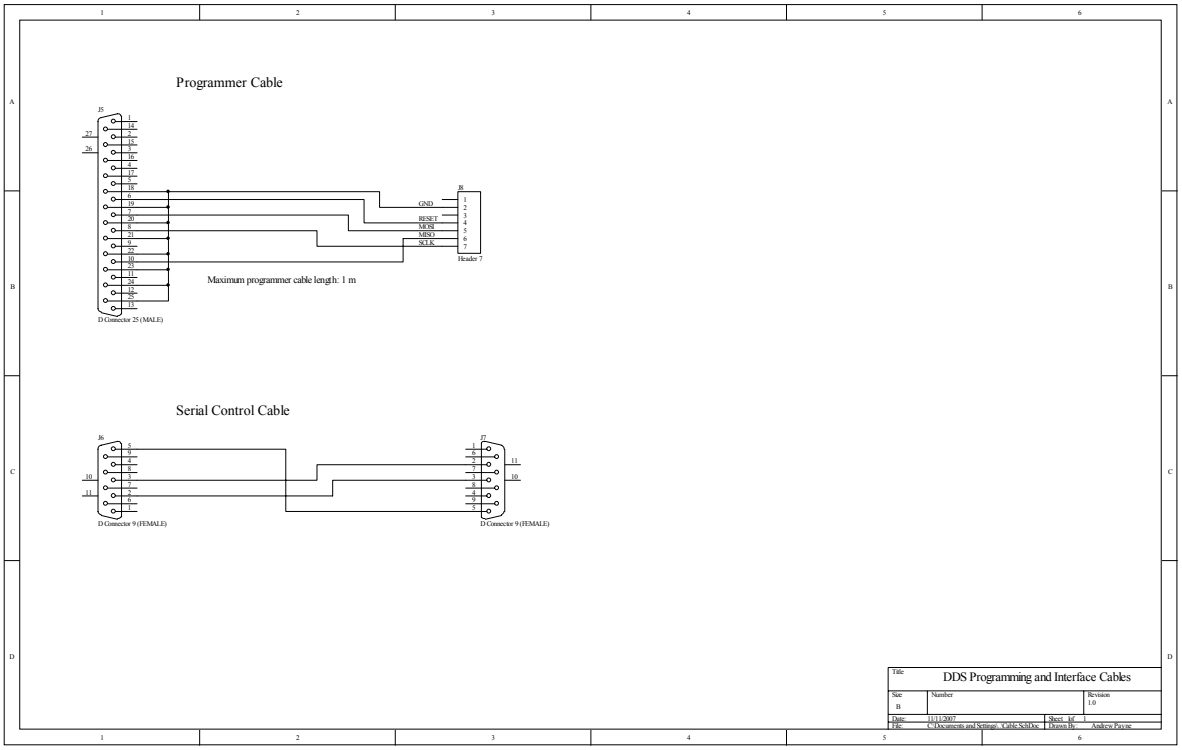

A.5 Image Intensifier Driver

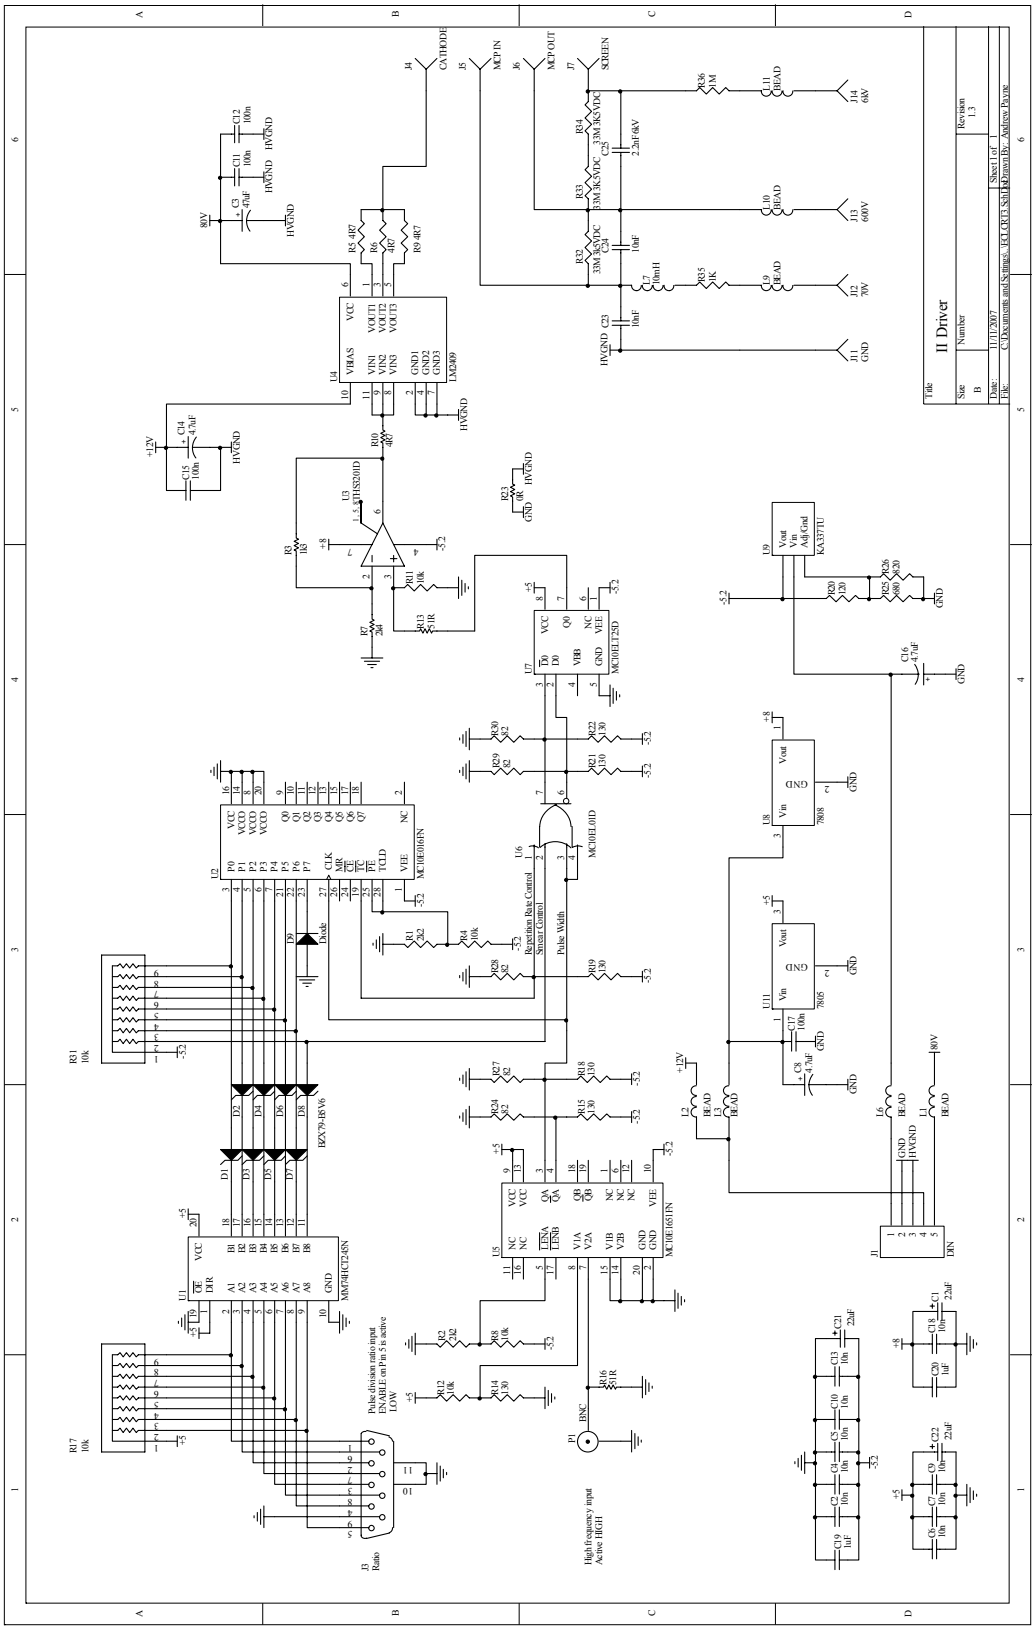

A.6 Laser Driver – Modulated Output

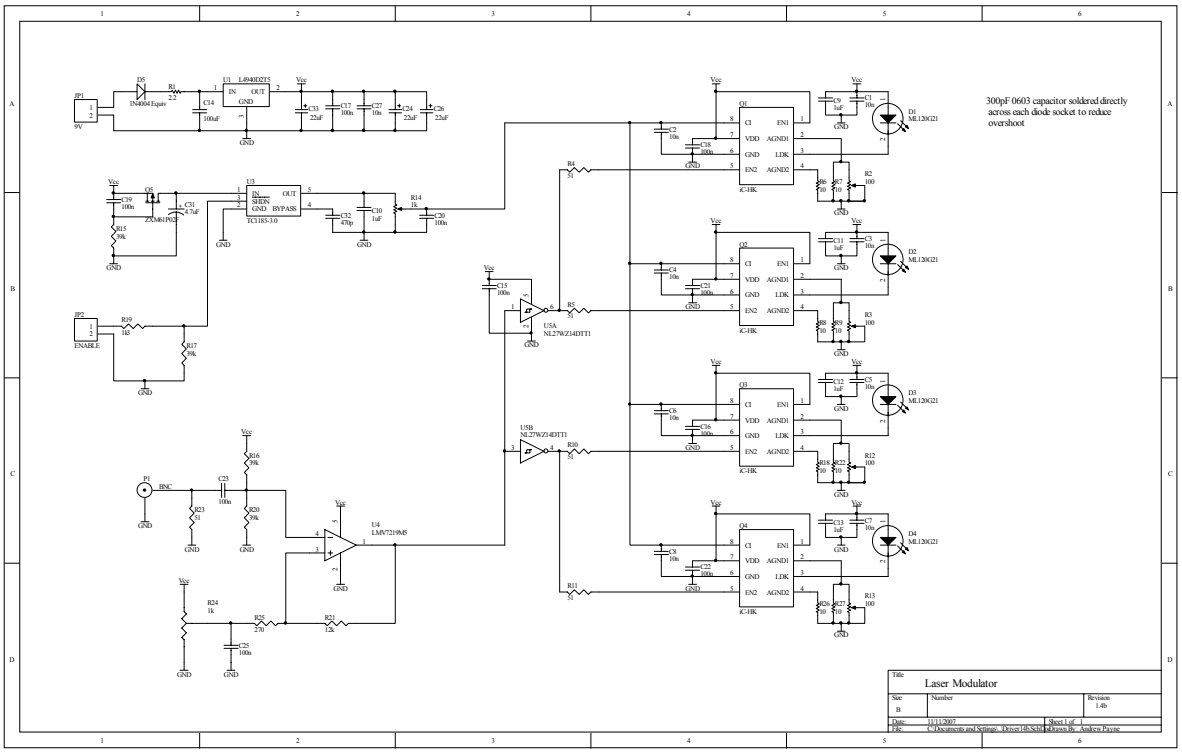

A.7 Laser Driver – Pulsed Output

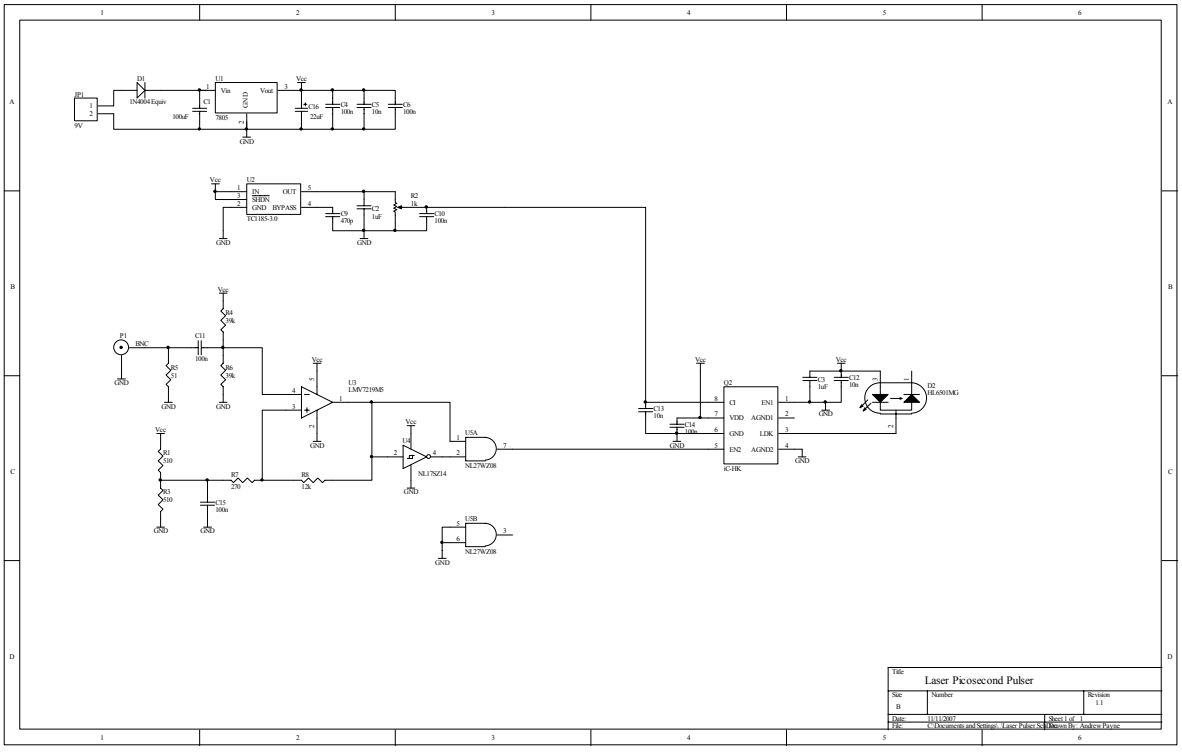

Appendix B. Eye Safety Calculations

Eye safety calculations require the determination of a maximum permissible exposure (MPE), which are determined from system variables such as the illumination wavelength, geometry of the optics, pulse width and repetition rate. This appendix follows example A.4-1 and subclause 13.3 from the New Zealand safety standard AS/NZS 2211.1:2004 (Standards New Zealand, 2004) to calculate the MPE for a repetitively pulsed system.

B.1 Average Power Limitation

The MPE is determined from the most restrictive of three criteria; the exposure from a single pulse within a pulse train, the average exposure for a pulse train of duration T, and the average exposure from a pulse train multiplied by a correction factor. For simplicity, the calculations assume worst case conditions utilising a point source (C_6 = 1), accidental viewing (t = 0.25 s), and a repetition rate of 1 MHz at the wavelength of λ=658 nm. The following equations specify the MPE for each condition:

a) The exposure from any single pulse within a pulse train shall not exceed the MPE for a single pulse

$$\begin{aligned} H_{single} &= 5\times10^{-3} C_6 \ \mathrm{Jm^{-2}} \\ &= 5\times10^{-3} \ \mathrm{Jm^{-2}} \end{aligned} \qquad \textbf{(B.1)}$$

b) The average exposure for a pulse train of exposure duration T shall not exceed the MPE given in Table 5.3 for a single pulse of exposure duration T. For the 0.25 s exposure duration, the radiant exposure is limited to

$$\begin{aligned} H_T &= 18\,t^{0.75}\,C_6 \text{ Jm}^{-2} \\ &= 18(0.25)^{0.75} \text{ Jm}^{-2} \\ &= 6.36 \text{ Jm}^{-2} \end{aligned} \tag{B.2}$$

given that there are $N = 2.5\times10^5$ pulses in 0.25 s, the single pulse radiant exposure is

$$\begin{aligned} H_{\text{single-ave}} &= \frac{H_T}{N} \\ &= \frac{6.36}{2.5\times10^5} \\ &= 2.55\times10^{-5} \text{ Jm}^{-2} \end{aligned} \tag{B.3}$$

c) The average exposure from pulses within a pulse train shall not exceed the MPE for a single pulse multiplied by a correction factor $N^{-1/4}$. Since the laser is operating at high repetition rates, multiple pulses which occur within the time period T_i = 18×10^{-6} s are counted as a single pulse, hence the effective pulse repetition frequency is

$$\begin{aligned} F_E &= \frac{1}{T_i} \\ &= 55.56 \text{ kHz} \end{aligned} \tag{B.4}$$

The MPE for a pulse of duration T_i is the same as a single pulse, 5×10^{-3} Jm^{-2}. The effective number of pulses in 0.25 s is

$$\begin{aligned} N_E &= T\,F_E \\ &= 0.25\cdot55.56\times10^3 \\ &= 1.39\times10^4 \end{aligned} \tag{B.5}$$

The radiant exposure is therefore

$$\begin{aligned} H_{\text{train}} &= H_{\text{single-eff}}\cdot N_E^{-\frac{1}{4}} \\ &= 5\times10^3\cdot(1.39\times10^4)^{-\frac{1}{4}} \\ &= 4.6\times10^{-4} \text{ Jm}^{-2} \end{aligned} \tag{B.6}$$

Dividing the result of Equation (B.6) by the number of pulses during period T_i allows the three MPEs to be compared.

$$\begin{aligned} H_{\text{single-train}} &= \frac{H_{\text{train}}}{T_i \cdot f} \\ &= \frac{4.6\times10^{-4}}{18\times10^{-6}\cdot 1\times10^{6}} \\ &= 2.55\times10^{-5}\ \text{Jm}^{-2} \end{aligned} \qquad \textbf{(B.7)}$$

Conditions b) and c) are the most restrictive; hence eye safety for this system is constrained by the average optical power and not the peak pulsed power. The calculations given assume the minimum repetition rate expected from the system of 1 MHz; for higher rates the eye safety constraint clearly remains restricted by the average power level.

B.2 MATLAB Code to Calculate NOHD

```
% Calculates the Nominal Ocular Hazard Distance (NOHD) for the laser source used in the
% range imaging system where four 658nm laser diodes are configured in a square pattern
% around the camera lens
%
% This is the minimum distance considered to be eye safe according to standard "AS/NZS
% 2211.1:2004 Safety of laser products"
%
% The limit for each possible combination of laser diodes is calculated, and the most
% restrictive should be used.  To simplify this calculation a high repetition rate is used
% (MHz), therefore the average power is the limiting constraint (rather than the power from
% individual pulses).
%
% Andrew Payne 13/07/2006

clear all;

phi = 10;            % divergence angle (degrees) between the two points where the
                     % irradiance has fallen to 37% (1/e) of its peak value
P = 0.050;           % average power level (watts) per source
a = 0.003;           % individual source diameter (metres) at exit point (beam waist)
seperation = 0.08;   % distance between sources (metres)
wavelength = 658;    % wavelength (nanometres)

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
if (wavelength < 600) || (wavelength > 700)
    fprintf('\nThis version only calculates Retinal Thermal Hazard for 600-700nm\n')
    return
end

phi = phi * pi/180;                  % convert to radians
alpha_min = 1.5e-3;
alpha_max = 100e-3;

% Accidental Viewing (t < 0.25 seconds)
% One diode
t = 0.25;

alpha = alpha_min;   % assume alpha = alpha_min for initial calculation (ie point source)
                     % as a starting point for the iterative process

C6 = 1;
MPE = 18*t^0.75*C6;                  % J/m^2
MPE = MPE / t;                       % W/m^2
```

```
NOHD = (1/phi)*sqrt(4*P/(pi*MPE)) - a/phi;

if 2*atan(a/(2*NOHD)) > alpha_min
    for i=1:100
        NOHD_old = NOHD;              % store old result to compare with next iteration
        alpha = 2*atan(a/(2*NOHD));
        if alpha < alpha_min
            C6 = 1;
        elseif alpha > alpha_max
            C6 = alpha_max/alpha_min;
        else
            C6 = alpha/alpha_min;
        end
        MPE = 18*t^0.75*C6;           % J/m^2
        MPE = MPE / t;                % W/m^2
        NOHD = (1/phi)*sqrt(4*P/(pi*MPE)) - a/phi;
        if (abs(NOHD_old - NOHD) < 0.001)
            break;
        end
    end
    if i == 100
        fprintf('\nERROR OCCURED\n');
        return;
    end
end

NOHD_accidental1 = NOHD;

% Accidental Viewing (t < 0.25 seconds)
% Two diodes

NOHD = 0;

for i=1:100
    NOHD_old = NOHD;                  % store old result to compare with next iteration
    alpha1 = 2*atan(a/(2*NOHD));              % angle in horizonal dimension
    alpha2 = 2*atan(a+seperation/(2*NOHD));  % angle in vertical dimension

    if alpha1 < 1.5e-3                % alpha values are limited as shown in Example A.2-4
        alpha1 = 1.5e-3;
    elseif alpha1 > 0.1
        alpha1 = 0.1;
    end

    if alpha2 < 1.5e-3
        alpha2 = 1.5e-3;
    elseif alpha2 > 0.1
        alpha2 = 0.1;
    end

    alpha = (alpha1 + alpha2)/2;

    C6 = alpha/alpha_min;

    MPE = 18*t^0.75*C6;               % J/m^2
    MPE = MPE / t;                    % W/m^2
    NOHD = (1/phi)*sqrt(4*P*2/(pi*MPE)) - a/phi;     % power is doubled for two diodes
    if (abs(NOHD_old - NOHD) < 0.001)
        break;
    end
end
if i == 100
    fprintf('\nERROR OCCURED\n');
    return;
end

NOHD_accidental2 = NOHD;

% Accidental Viewing (t < 0.25 seconds)
% Four diodes

NOHD = 0;

for i=1:100
```

```
    NOHD_old = NOHD;                % store old result to compare with next iteration

    alpha2 = 2*atan(a+seperation/(2*NOHD));

    if alpha < 1.5e-3
        alpha = 1.5e-3;
    elseif alpha > 0.1
        alpha = 0.1;
    end

    C6 = alpha/alpha_min;

    MPE = 18*t^0.75*C6;             % J/m^2
    MPE = MPE / t;                  % W/m^2
    NOHD = (1/phi)*sqrt(4*P*4/(pi*MPE)) - a/phi;  % power is 4 times for four diodes
    if (abs(NOHD_old - NOHD) < 0.001)
        break;
    end
end
if i == 100
    fprintf('\nERROR OCCURED\n');
    return;
end

NOHD_accidental4 = NOHD;

% Intentional Viewing (t > 10 seconds)
% One diode
alpha = alpha_min;   % assume alpha = alpha_min for initial calculation (ie point source)
                     % as a starting point for the iterative process
MPE = 10;            % W/m^2
NOHD = (1/phi)*sqrt(4*P/(pi*MPE)) - a/phi;

if 2*atan(a/(2*NOHD)) > alpha_min
    for i=1:100
        NOHD_old = NOHD;            % store old result to compare with next iteration
        alpha = 2*atan(a/(2*NOHD));
        if alpha > alpha_max
            C6 = alpha_max/alpha_min;
        else
            C6 = alpha/alpha_min;
        end
        T2 = 10 * 10^((alpha*1000-1.5)/98.5);
        if(T2 < 10)
            T2 = 10;
        elseif(T2 > 100)
            T2 = 100;
        end
        MPE = 18*C6*T2^-0.25;       % W/m^2
        NOHD = (1/phi)*sqrt(4*P/(pi*MPE)) - a/phi;
        if (abs(NOHD_old - NOHD) < 0.001)
            break;
        end
    end
    if i == 100
        fprintf('\nERROR OCCURED\n');
        return;
    end
end
T2_1 = T2;
NOHD_intentional1 = NOHD;

% Two diodes
% Intentional Viewing (t > 10 seconds)
% Extended source, so don't bother testing for point source constraint.

NOHD = 0;

for i=1:100
    NOHD_old = NOHD;                % store old result to compare with next iteration
    alpha1 = 2*atan(a/(2*NOHD));            % angle in horizonal dimension
    alpha2 = 2*atan(a+seperation/(2*NOHD));  % angle in vertical dimension

    if alpha1 < 1.5e-3              % alpha values are limited as shown in Example A.2-4
```

```
        alpha1 = 1.5e-3;
    elseif alpha1 > 0.1
        alpha1 = 0.1;
    end

    if alpha2 < 1.5e-3
        alpha2 = 1.5e-3;
    elseif alpha2 > 0.1
        alpha2 = 0.1;
    end

    alpha = (alpha1 + alpha2)/2;

    C6 = alpha/alpha_min;

    T2 = 10 * 10^((alpha*1000-1.5)/98.5);
    if(T2 < 10)
        T2 = 10;
    elseif(T2 > 100)
        T2 = 100;
    end
    MPE = 18*C6*T2^-0.25;             % W/m^2
    NOHD = (1/phi)*sqrt(4*P*2/(pi*MPE)) - a/phi;     % power is doubled for two diodes
    if (abs(NOHD_old - NOHD) < 0.001)
        break;
    end
end
if i == 100
    fprintf('\nERROR OCCURED\n');
    return;
end
T2_2 = T2;
NOHD_intentional2 = NOHD;

% Four diodes
% Intentional Viewing (t > 10 seconds)
% Extended source, so don't bother testing for point source constraint.

NOHD = 0;

for i=1:100
    NOHD_old = NOHD;                  % store old result to compare with next iteration

    alpha = 2*atan(a+seperation/(2*NOHD));

    if alpha < 1.5e-3                 % alpha values are limited as shown in Example A.2-4
        alpha = 1.5e-3;
    elseif alpha > 0.1
        alpha = 0.1;
    end

    C6 = alpha/alpha_min;

    T2 = 10 * 10^((alpha*1000-1.5)/98.5);
    if(T2 < 10)
        T2 = 10;
    elseif(T2 > 100)
        T2 = 100;
    end
    MPE = 18*C6*T2^-0.25;             % W/m^2
    NOHD = (1/phi)*sqrt(4*P*4/(pi*MPE)) - a/phi;   %power is 4 times for four diodes
    if (abs(NOHD_old - NOHD) < 0.001)
        break;
    end
end
if i == 100
    fprintf('\nERROR OCCURED\n');
    return;
end
T2_4 = T2;
NOHD_intentional4 = NOHD;

fprintf('\n\nCalculated from AS/NZS 2211.1:2004, Safety of laser products\n');
fprintf('Wavelength: \t\t\t\t\t\t%d nm\n',wavelength);
```

```
fprintf('Individual Source Average Power: \t%.3f W\n',P);
fprintf('Divergence Angle: \t\t\t\t\t%.1f degrees\n',phi*180/pi);
fprintf('Source Seperation: \t\t\t\t\t%.3f m\n\n',seperation);

fprintf('Nominal Ocular Hazard Distance for accidental viewing 1 diode (<0.25 seconds)...
is %.3fm\n',NOHD_accidental1);
fprintf('Nominal Ocular Hazard Distance for accidental viewing 2 diodes (<0.25 seconds)...
is %.3fm\n',NOHD_accidental2);
fprintf('Nominal Ocular Hazard Distance for accidental viewing 4 diodes (<0.25 seconds)...
is %.3fm\n',NOHD_accidental4);
fprintf('Nominal Ocular Hazard Distance for intentional viewing 1 diode (>%.1f seconds)...
is %.3fm\n',T2_1,NOHD_intentional1);
fprintf('Nominal Ocular Hazard Distance for intentional viewing 2 diodes (>%.1f seconds)...
is %.3fm\n',T2_2,NOHD_intentional2);
fprintf('Nominal Ocular Hazard Distance for intentional viewing 4 diodes (>%.1f seconds)...
is %.3fm\n',T2_4,NOHD_intentional4);
fprintf('\nResults are only valid if distance is greater than 0.1 m, but are shown ...
for completeness\n');                %see section 3.7
```

References

3DV Systems, (2003). "ZCAM", Product specifications available from http://www.3dvsystems.com/products/zcam.html, accessed November 2005.

Advanced Scientific Concepts, (2006). "Portable 3D Camera", Product specifications available from http://www.advancedscientificconcepts.com/camera1.html, accessed May 2007.

Agilent Technologies, (2004). "T-1 3/4 (5 mm) Precision Optical Performance AlInGaP LED Lamps", Technical datasheet available from http://www.farnell.com/ datasheets/70661.pdf, accessed July 2004.

Ailisto, H., Heikkinen, V., Mitikka, R., Myllyla, R., Kostamovaara, J., Mantyniemi, A. and Koskinen, M., (2002). "Scannerless imaging pulsed-laser range finding," *Journal of Optics A - Pure and Applied Optics*, vol. 4 (6), pp. S337-S346.

Albota, M. A., Heinrichs, R. M., Kocher, D. G., Fouche, D. G., Player, B. E., O'Brien, M. E., Aull, B. F., Zayhowski, J. J., Mooney, J., Willard, B. C. and Carlson, R. R., (2002). "Three-dimensional imaging laser radar with a photon-counting avalanche photodiode array and microchip laser," *Applied Optics*, vol. 41 (36), pp. 7671-7678.

Analog Devices, (2003). "AD9952 400 MSPS 14-Bit, 1.8 V CMOS Direct Digital Synthesizer rev. 0", Technical Datasheet available from http://www.analog.com/ UploadedFiles/Data_Sheets/AD9952.pdf, accessed June 2004.

Analog Devices, (2004). "AD9852 CMOS 300 MSPS Complete-DDS rev. C", Technical Datasheet available from http://www.analog.com/UploadedFiles/Data_Sheets/ AD9852.pdf, accessed May 2004.

Andersson, P., (2006). "Long-range three-dimensional imaging using range-gated laser radar images," *Optical Engineering*, vol. 45 (3), pp. 034301.

Andressen, C., (2005). "LADAR: The Emergence of an EO Technology to Defend Against Future Threats," *Technology Today, Highlighting Raytheon's Technology* (2), pp. 11

Atkinson, K., (2001). *Close Range Photogrammetry and Machine Vision*, Whittles Publishing, Scotland, UK.

Aull, B., Loomis, A., Young, D., Heinrichs, R., Felton, B., Daniels, P. and Landers, D., (2002). "Geiger-mode avalanche photodiodes for three-dimensional imaging," *Lincoln Laboratory Journal*, vol. 13 (2), pp. 335-349.

Becker, W. and Bergman, A., (1993). "Lifetime Imaging Techniques for Optical Microscopy", available from http://www.becker-hickl.de/pdf/tcvgbh1.pdf, accessed May 2004.

Behrend, M. R., Sun, Y., Vernier, P. T., Kuthi, A. and Gundersen, M. A., (2004). "Four-channel pulse generator for real-time biological investigations," *Proc. 26th IEEE Int. Power Modulator Symp.*, pp. 210-215.

Bernius, M. and Chutjian, A., (1990). "Improved high-voltage, high-frequency square-wave generator," *Review of Scientific Instruments*, vol. 61, pp. 925.

Bosch, L. A., (1995). "Dynamic uses of image intensifiers," *Proc. SPIE - Photoelectronic Detectors, Cameras, and Systems*, vol. 2551, pp. 159-172.

Bostelman, R., Hong, T., Madhavan, R. and Weiss, B., (2005). "3D range imaging for urban search and rescue robotics research," *IEEE International Workshop on Safety, Security and Rescue Robotics*.

Brandon, D., (2003). "AN-605, Synchronizing Multiple AD9852 DDS-Based Synthesizers", Technical Application Note available from http://www.analog.com/UploadedFiles/Application_Notes/371092853519044414816844703 5AN605_0.pdf, accessed June 2004.

Brandon, D., (2004). "DDS Design", *EDN*, pp. 71-84 available from http://www.edn.com/ednmag/article/CA415103.html, accessed July 2004.

Brandon, D., (2006). "AN-823 Direct Digital Synthesizers in Clocking Applications Time Jitter in Direct Digital Synthesizer-Based Clocking Systems", Technical Application Note available from http://www.analog.com/UploadedFiles/Application_Notes/44165304775709740692131461831AN823_0.pdf, accessed April 2006.

Brandon, D. and Gentile, K., (2006). "AN-837 DDS-Based Clock Jitter Performance vs. DAC Reconstruction Filter Performance", Technical Application Note available from http://www.analog.com/UploadedFiles/Application_Notes/351016224AN_837.pdf, accessed November 2007.

Breuer, P., Eckes, C. and Muller, S., (2007). "Hand Gesture Recognition with a Novel IR Time-of-Flight Range Camera–A Pilot Study" In: *Computer Vision/Computer Graphics Collaboration Techniques*. Eds, A. Gagalowicz and W. Philips, Springer-Verlag, Berlin / Heidelberg, pp. 247-260.

Briers, J. D., (1996). "Speckle techniques" In: *Optical Measurement Methods in Biomechanics*. Eds, J. F. Orr and J. C. Shelton, Chapman and Hill, London, pp. 76-98.

Busck, J., (2005). "Underwater 3-D optical imaging with a gated viewing laser radar," *Optical Engineering*, vol. 44 (11), pp. 116001.

Busck, J. and Heiselberg, H., (2004). "Gated viewing and high-accuracy three-dimensional laser radar," *Applied Optics*, vol. 43 (24), pp. 4705-4710.

Büttgen, B., Oggier, T., Lehmann, M., Kaufmann, R., Neukom, S., Richter, M., Schweizer, M., Beyeler, D., Cook, R., Gimkiewicz, C., Urban, C., Metzler, P., Seitz, P. and Lustenberger, F., (2006). "High-speed and high-sensitive demodulation pixel for 3D imaging," *Proc. SPIE - Three-Dimensional Image Capture and Applications VII*.

Carnegie, D. A., Cree, M. J. and Dorrington, A. A., (2005). "A high-resolution full-field range imaging system," *Review of Scientific Instruments*, vol. 76 (8), pp. 083702.

Carnegie, D. A., Cree, M. J. and Dorrington, A. A., (2006). "Range and Velocity Sensing System," Patent Application No. PCT/NZ2005/000241, filed 13 September 2005.

Chang, K., Bowyer, K. and Flynn, P., (2003). "Face recognition using 2D and 3D facial data," *Proc. ACM Workshop on Multimodal User Authentication*, pp. 25–32.

Christie, S., Hill, S. L., Bury, B., Gray, J. O. and Booth, K. M., (1995). "Design and Development of a Multi-Detecting 2-Dimensional Ranging Sensor," *Measurement Science & Technology*, vol. 6 (9), pp. 1301-1308.

Creath, K., (1987). "Step height measurement using two-wavelength phase-shifting interferometry," *Applied Optics*, vol. 26 (14), pp. 2810-2816.

Dawson, R. W., (1980). "Led Bandwidth Improvement by Bipolar Pulsing," *IEEE Journal of Quantum Electronics*, vol. 16 (7), pp. 697-699.

Detch, J. L. J. and Noel, B. W., (1981). "Radial-pulse propagation and impedance characteristics of optically shuttered channel intensifier tubes," *Proc. SPIE - Los Alamos Conf. on Optics '81*, vol. 288, pp. 434-446.

Dinish, U. S., Chao, Z. X., Seah, L. K., Singh, A. and Murukeshan, V. M., (2005). "Formulation and implementation of a phase-resolved fluorescence technique for latent-fingerprint imaging: theoretical and experimental analysis," *Applied Optics*, vol. 44 (3), pp. 297-304.

Dorrington, A., Carnegie, D. and Cree, M., (2006). "Toward 1-mm depth precision with a solid state full-field range imaging system," *Proc. SPIE - Sensors, Cameras, and Systems for Scientific/Industrial Applications VII*, vol. 6068, pp. 60680K.

Dorrington, A. A., (2004). "Range Sensing System," Patent Application No. PCT/NZ2004/000070, filed 6 April 2004.

Dorrington, A. A., Cree, M. J. and Carnegie, D. A., (2005a). "The Importance of CCD Readout Smear in Heterodyne Imaging Phase Detection Applications," *Proc. Image and Vision Computing New Zealand (IVCNZ)*, pp. 73-78.

Dorrington, A. A., Cree, M. J., Carnegie, D. A. and Payne, A. D., (2005b). "Selecting Signal Frequencies for Best Performance of Fourier-based Phase Detection," *Proc. Twelfth New Zealand Electronics Conference (ENZCON)*, pp. 189–193.

Dorrington, A. A., Cree, M. J., Carnegie, D. A., Payne, A. D. and Conroy, R. M., (2007a). "Heterodyne range imaging as an alternative to photogrammetry," *Proc. SPIE - Videometrics IX*, vol. 6491, pp. 64910D.

Dorrington, A. A., Cree, M. J., Payne, A. D., Conroy, R. M. and Carnegie, D. A., (2007b). "Achieving sub-millimetre precision with a solid-state full-field heterodyning range imaging camera," *Measurement Science & Technology*, vol. 18, pp. 2809-2816.

Elder, A. D., Frank, J. H., Swartling, J., Dai, X. and Kaminski, C. F., (2006). "Calibration of a wide-field frequency-domain fluorescence lifetime microscopy system using light emitting diodes as light sources," *Journal of Microscopy-Oxford*, vol. 224, pp. 166-180.

Ens, J. and Lawrence, P., (1993). "An Investigation of Methods for Determining Depth from Focus," *Proc. IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15 (2), pp. 97-108.

Fairchild Semiconductor, (1997). "74LCX74 Low Voltage Dual D-Type Positive Edge-Triggered Flip-Flop with 5V Tolerant Inputs", Technical Datasheet available from http://www.fairchildsemi.com/ds/74/74LCX74.pdf, accessed June 2004.

Fraden, J., (2004). *Handbook of Modern Sensors: Physics, Designs, and Applications 3rd ed.*, 3rd edition, Springer, New York.

Freischlad, K. and Koliopoulos, C. L., (1990). "Fourier Description of Digital Phase-Measuring Interferometry," *Journal of the Optical Society of America A - Optics Image Science and Vision*, vol. 7 (4), pp. 542-551.

Frenkel, A., Sartor, M. A. and Wlodawski, M. S., (1997). "Photon-noise-limited operation of intensified CCD cameras," *Applied Optics*, vol. 36 (22), pp. 5288-5297.

Gökstorp, M., (1994). "Computing depth from out-of-focus blur using a local frequency representation," *Proc. 12th IAPR Int. Conf. Pattern Recognition, Conference A: Computer Vision & Image Processing*, vol. 1, pp. 153-158.

Gokturk, S. B., Yalcin, H. and Bamji, C., (2004). "A Time-Of-Flight Depth Sensor - System Description, Issues and Solutions," *Proc. 2004 IEEE Conf. Computer Vision and Pattern Recognition Workshop*, vol. 3, pp. 35.

Goldberg, B. G., (1999). *Digital frequency synthesis demystified : DDS and fractional-N PLLs*, LLH Technology, Eagle Rock, VA.

Gonzalez-Banos, H. and Davis, J., (2004). "Computing depth under ambient illumination using multi-shuttered light," *Proc. IEEE Computer Vision and Pattern Recognition, 2004.*, vol. 2, pp. 234-241.

Gulden, P., Vossiek, M., Heide, P. and Schwarte, R., (2002). "Novel opportunities for optical level gauging and 3-D-imaging with the photoelectronic mixing device," *IEEE Transactions on Instrumentation and Measurement*, vol. 51 (4), pp. 679-684.

Halmos, M., Jack, M., Asbrock, J., Anderson, C., Bailey, S., Chapman, G., Gordon, E., Herning, P., Kalisher, M. and Klaras, L., (2003). "3D flash ladar at Raytheon," *Proc. SPIE - Laser Radar Technology and Applications VI*, vol. 4377, pp. 84-97.

Hammond, P., Hutton, T. J., Allanson, J. E., Campbell, L. E., Hennekam, R. C. M., Holden, S., Patton, M. A., Shaw, A., Temple, I. K., Trotter, M., Murphy, K. C. and Winter, R. M., (2004). "3D analysis of facial morphology," *American Journal of Medical Genetics Part A*, vol. 126A (4), pp. 339-348.

Hariharan, P., (1985). *Optical Interferometry*, Academic Press, Sydney.

Harris, F. J., (1978). "Use of Windows for Harmonic-Analysis with Discrete Fourier-Transform," *Proc. IEEE*, vol. 66 (1), pp. 51-83.

Henderson, A. and Schulmeister, K., (2004). *Laser Safety*, Institute of Physics Publishing, Bristol, Philadelphia.

Henderson, A. R., (1997). *A Guide to Laser Safety*, Chapman & Hall, London.

Hosking, R., (2002). "Direct Digital Synthesizer using the AD9854 quadrature DDS", available from http://members.iinet.net.au/~richardh/Direct+Digital+Synthesizer+using+the+AD9854+quadrature+DDS.htm, accessed June 2004.

Höß, P., Fleder, K. and Ehrhardt, J., (1998). "Subnanosecond optical gating using coax cable input microchannel plate image intensifier," *Optical Engineering*, vol. 37 (8), pp. 2213-2216.

Howard, I. P. and Rogers, B. J., (1995). *Binocular vision and stereopsis*, Oxford University Press, USA.

Hsu, S., Acharya, S., Rafii, A. and New, R., (2006). "Performance of a Time-of-Flight Range Camera for Intelligent Vehicle Safety Applications," *Advanced Microsystems for Automotive Applications 2006*, pp. 205-214.

Hu, M. H., Nguyen, H. K., Song, K. C., Li, Y. B., Visovsky, N. J., Liu, X. S., Nishiyama, N., Coleman, S., Hughes, L. C., Gollier, J., Miller, W., Bhat, R. and Zah, C. E., (2006). "High-power high-modulation-speed 1060-nm DBR lasers for green-light emission," *IEEE Photonics Technology Letters*, vol. 18 (4), pp. 616-618.

Iddan, G. J. and Yahav, G., (2001). "3-D imaging in the studio and elsewhere," *Proc. SPIE - Three-Dimensional Image Capture and Applications IV*, vol. 4298, pp. 48-55.

Iizuka, K., (2007). Personal email communication, 14 April 2007.

Izhal, A. H., Ushinaga, T., Sawada, T., Homma, M., Maeda, Y. and Kawahito, S., (2005). "A CMOS time-of-flight range image sensor with gates on field oxide structure," *IEEE Sensors*, pp. 141 - 144.

Jelalian, A., (1992). *Laser Radar Systems*, Artech House, Boston.

Jimenez, A. R., Ceres, R. and Pons, J. L., (2000). "A vision system based on a laser range-finder applied to robotic fruit harvesting," *Machine Vision and Applications*, vol. 11 (6), pp. 321-329.

Juberts, M. and Barbera, A., (2004). "Status report on next-generation LADAR for driving unmanned ground vehicles," *Proc. SPIE - Mobile Robots XVII*, vol. 5609, pp. 1-12.

Kahlmann, T. and Ingensand, H., (2005). "Calibration and improvements of the high-resolution range-imaging camera SwissRanger," *Proc. SPIE - Videometrics VIII*, vol. 5665, pp. 144-155.

Kahlmann, T. and Ingensand, H., (2006). "Calibration of the fast range imaging camera SwissRanger for use in the surveillance of the environment," *Proc. SPIE - Electro-Optical Remote Sensing II*, vol. 6396, pp. 639605.

Kahlmann, T., Remondino, F. and Ingensand, H., (2006). "Calibration for Increased Accuracy of the Range Imaging Camera SwissRanger," *Proc. ISPRS Commission V Symposium - Image Engineering and Vision Metrology*, pp. 136-141.

Kawakita, M., Iizuka, K., Aida, T., Kikuchi, H., Fujikake, H., Yonai, J. and Takizawa, K., (2000). "Axi-Vision Camera (real-time distance-mapping camera)," *Applied Optics*, vol. 39 (22), pp. 3931-3939.

Kawakita, M., Iizuka, K., Aida, T., Kurita, T. and Kikuchi, H., (2004a). "Real-time three-dimensional video image composition by depth information," *IEICE Electronics Express*, vol. 1 (9), pp. 237-242.

Kawakita, M., Iizuka, K., Iino, Y., Kikuchi, H., Fujikake, H. and Aida, T., (2006). "Real-time depth-mapping three-dimension TV camera (Axi-Vision camera)," *Systems and Computers in Japan*, vol. 37 (8), pp. 77-89.

Kawakita, M., Iizuka, K., Iwama, R., Takizawa, K., Kikuchi, H. and Sato, F., (2004b). "Gain-modulated Axi-Vision Camera (high speed high-accuracy depth-mapping camera)," *Optics Express*, vol. 12 (22), pp. 5336-5344.

Kawakita, M., Iizuka, K., Nakamura, H., Mizuno, I., Kurita, T., Aida, T., Yamanouchi, Y., Mitsumine, H., Fukaya, T., Kikuchi, H. and Sato, F., (2004c). "High-definition real-time depth-mapping TV camera: HDTV axi-vision camera," *Optics Express*, vol. 12 (12), pp. 2781-2794.

Koechner, W., (1999). *Solid-State Laser Engineering*, Springer, New York.

Kolski, S., Ferguson, D., Bellino, M. and Siegwart, R., (2006). "Autonomous Driving in Structured and Unstructured Environments," *Proc. 2006 IEEE Intelligent Vehicles Symposium*, pp. 558-563.

Lacoste, R., (2001). "DDS-GEN-Part 2: The Generator," *Circuit Cellar* (130), pp. 50-59

Lakowicz, J. R. and Berndt, K. W., (1991). "Lifetime-Selective Fluorescence Imaging Using an Rf Phase-Sensitive Camera," *Review of Scientific Instruments*, vol. 62 (7), pp. 1727-1734.

Lange, R., (2000). "3D Time-of-flight distance measurement with custom solid-state image sensors in CMOS/CCD-technology," *Ph.D. dissertation, University of Siegen*.

Lange, R. and Seitz, P., (2001). "Solid-state time-of-flight range camera," *IEEE Journal of Quantum Electronics*, vol. 37 (3), pp. 390-397.

Lange, R., Seitz, P., Biber, A. and Lauxtermann, S., (2000). "Demodulation Pixels in CCD and CMOS Technologies for Time-of-Flight Ranging," *Proc. SPIE - Sensors and Camera Systems for Scientific, Industrial, and Digital Photography Applications*, vol. 3965, pp. 177-189.

Lau, K. Y., (1988). "Gain Switching of Semiconductor Injection-Lasers," *Applied Physics Letters*, vol. 52 (4), pp. 257-259.

Lee, T. P. and Dentai, A. G., (1978). "Power and Modulation Bandwidth of GaAs-AlGaAs High-Radiance Led's for Optical Communication Systems," *IEEE Journal of Quantum Electronics*, vol. 14 (3), pp. 150-159.

Marino, R. M. and Davis, W. R., (2005). "Jigsaw: A Foliage-Penetrating 3D Imaging Laser Radar System," *MIT Lincoln Lab Journal*, vol. 15, pp. 23-36.

Medina, A., Gaya, F. and del Pozo, F., (2006). "Compact laser radar and three-dimensional camera," *Journal of the Optical Society of America A-Optics Image Science and Vision*, vol. 23 (4), pp. 800-805.

Mitsubishi Electric, (2004). "Mitsubishi Laser Diodes ML1XX21 Series Datasheet", available from http://pagesperso-orange.fr/redlum.xohp/electronics/data/DS_1X21_Jan04.PDF, accessed Feb 2008.

Möller, T., Kraft, H., Frey, J., Albrecht, M., Lange, R., GmbH, P. and Eichenhang, A., (2005). "Robust 3D Measurement with PMD Sensors," *Proc. 1st Range Imaging Research Day*, pp. Supplement to the Proceedings.

Monson, T. C., Grantham, J. W., Childress, S. W., Sackos, J. T., Nellums, R. O. and Lebien, S. M., (1999). "Characterization of scannerless ladar," *Proc. SPIE - Laser Radar Technology and Applications IV*, vol. 3707, pp. 409-420.

Mynbaev, D. and Scheiner, L., (2001). *Fiber-optic communications technology*, Prentice Hall Upper Saddle River, NJ.

Nagihara, S., Mulligan, K. R. and Xiong, W., (2004). "Use of a three-dimensional laser scanner to digitally capture the topography of sand dunes in high spatial resolution," *Earth Surface Processes and Landforms*, vol. 29 (3), pp. 391-398.

Negreiros, M., Souza Jr, A., Carro, L. and Susin, A., (2007). "RF Digital Signal Generation Beyond Nyquist," *Proc. 25th IEEE VLSI Test Symmposium*, pp. 15-22.

Nellums, R. O., Habbit, R. D., Heying, M. R., Pitts, T. A. and Sandusky, J. V., (2006). "3D scannerless LADAR for Orbiter inspection," *Proc. SPIE - Spaceborne Sensors III*, vol. 6220, pp. 62200G.

Niclass, C., Rochas, A., Besse, P. A. and Charbon, E., (2005). "Design and characterization of a CMOS 3-D image sensor based on single photon avalanche diodes," *IEEE Journal of Solid-State Circuits*, vol. 40 (9), pp. 1847-1854.

Oggier, T., Lehmann, M., Kaufmann, R., Schweizer, M., Richter, M., Metzler, P., Lang, G., Lustenberger, F. and Blanc, N., (2004). "An all-solid-state optical range camera for 3D real-time imaging with sub-centimeter depth resolution (SwissRanger)," *Proc. SPIE - Optical Design and Engineering*, vol. 5249, pp. 534-545.

Payne, A. D., (2007). "Harmonic rejection sampling methodology," Provisional Patent Application No. 562739, filed 19 October 2007.

Payne, A. D., Carnegie, D. A., Dorrington, A. A. and Cree, M. J., (2005). "A Synchronised Direct Digital Synthesiser," *Proc. First International Conference on Sensing Technology (ICST)*, pp. 174-179.

Payne, A. D., Carnegie, D. A., Dorrington, A. A. and Cree, M. J., (2006a). "Full Field Image Ranger Hardware," *Proc. Third IEEE International Workshop on Electronic Design, Test and Applications (DELTA'06)*, pp. 263-268.

Payne, A. D., Dorrington, A. A., Cree, M. J. and Carnegie, D. A., (2006b). "Image Intensifier Characterisation," *Proc. Image and Vision Computing New Zealand 2006 (IVCNZ'06)*, pp. 487-492.

Payne, A. D., Dorrington, A. A., Cree, M. J. and Carnegie, D. A., (2008a). "Characterizing an image intensifier in a full-field image ranging system," *IEEE Sensors Journal, accepted for publication*.

Payne, A. D., Dorrington, A. A., Cree, M. J. and Carnegie, D. A., (2008b). "Improved Linearity Using Harmonic Error Rejection in a Full-Field Range Imaging System," *Proc. SPIE - 3D Image Capture and Applications VII*, vol. 6805, pp. 68050D.

Photek, (2002). "Test Data Summary" datasheet supplied with image intensifier.

Photek, (2003). "MCP Image Intensifiers", available from http://www.photek.com/datasheets/Image%20Intensifiers.pdf, accessed February 2006.

PMD Technologies, (2006). "PMD [vision] 19k", Product brochure available from http://www.pmdtec.com/inhalt/download/documents/PMDvision19k_000.pdf, accessed April 2007.

Princeton Instruments, (2007). "Comparison of Lens-Coupled and Fibreoptic-Coupled ICCD Cameras", available from http://content.piacton.com/Uploads/Princeton/Documents/Library/UpdatedLibrary/Comparison_of_Coupled_ICCD_Cameras_PI.pdf, accessed March 2008.

Razavi, B., (1996). "Challenges in portable RF transceiver design," *IEEE Circuits and Devices Magazine*, vol. 12 (5), pp. 12-25.

Salvi, J., Pages, J. and Batlle, J., (2004). "Pattern codification strategies in structured light systems," *Pattern Recognition*, vol. 37 (4), pp. 827-849.

Saul, R. H., Lee, T. P. and Burrus, C. A., (1985). "Light-Emitting Diode Device Design" In: *Semiconductors and Semimetals*. Eds, R. K. Willardson and A. C. Bear, Academic, New York, vol. 22 Part C, pp. 193-237.

Schrey, O. M., Elkhalili, O., Mengel, P., Petermann, M., Brockherde, W. and Hosticka, B. J., (2004). "A 4 X 64 pixel CMOS image sensor for 3-D measurement applications," *IEEE Journal of Solid-State Circuits*, vol. 39 (7), pp. 1208-1212.

Schroeder, W., Forgber, E. and Estable, S., (1999). "Scannerless laser range camera," *Sensor Review*, vol. 19 (4), pp. 285-291.

Schubert, E., (2006). *Light-Emitting Diodes*, 2nd edition, Cambridge University Press, Cambridge, UK.

Schwarte, R., (2004). "Breakthrough in multichannel laser-radar technology providing thousands of high-sensitive lidar receivers on a chip," *Proc. SPIE - Laser Radar Techniques for Atmospheric Sensing*, vol. 5575, pp. 126-136.

Sirohi, R. and Chau, F., (1999). *Optical Methods of Measurement: Wholefield Techniques*, Marcel Dekker, New York.

Spirig, T., Marley, M. and Seitz, P., (1997). "The multitap lock-in CCD with offset subtraction," *IEEE Transactions on Electron Devices*, vol. 44 (10), pp. 1643-1647.

Standards New Zealand, (2004). *AS/NZS 2211.1:2004 Safety of laser products Part 1: Equipment classification, requirements and user's guide*.

Stann, B., Giza, M. and Robinson, D., (1999). "A scannerless imaging ladar using a laser diode illuminator and FM/CW radar principles," *Proc. SPIE - Laser Radar Technology and Applications IV*, vol. 3707, pp. 421-431.

Stettner, R., Bailey, H. and Richmond, R., (2004). "Eye-safe laser radar 3 D imaging," *Proc. SPIE - Laser Radar Technology and Applications IX*, vol. 5412, pp. 111-116.

Strand, T. C., (1985). "Optical three-dimensional sensing for machine vision," *Optical engineering*, vol. 24 (1), pp. 33-40.

Suhara, T., (2004). *Semiconductor Laser Fundamentals*, Marcel Dekker, New York.

Swartz, B., (1994). "Laser range gate underwater imaging advances," *Proc. IEEE OCEANS'94 - Oceans Engineering for Today's Technology and Tomorrow's Preservation.*, vol. 2, pp. 722-727.

Thomas, M., (1994). "Fast optical gating using planar-lead MCPIIS and linear microstrip impedance transformers," *Proc. SPIE - Ultrahigh- and High-Speed Photography, Videography, and Photonics '94*, vol. 2273, pp. 214-225.

Thomas, M., Zagarino, P. and Yates, G., (1993). "Dynamic modulation transfer function measurements in gated microchannel-plate image intensifiers," *Proc. SPIE - Ultrahigh- and High-Speed Photography, Videography, and Photonics*, vol. 1757, pp. 32-39.

Thomas, S. W., Shimkunas, A. R. and Mauger, P. E., (1991). "Subnanosecond intensifier gating using heavy and mesh cathode underlays," *Proc. SPIE - 19th Intl Congress on High-Speed Photography and Photonics*, vol. 1358, pp. 91-99.

Tomkins, P. and Lyons, A., (1999). "Properties of Low-light-level Intensified Cameras" In: *Fluorescent and luminescent probes for biological activity: a practical guide to technology for quantitative real-time analysis*. Ed, W. Mason, Academic Press Ltd., U.K., pp. 491-506.

Ushinaga, T., Halin, I. A., Sawada, T., Kawahito, S., Homma, M. and Maeda, Y., (2006). "A QVGA-size CMOS time-of-flight range image sensor with background light charge draining structure," *Proc. SPIE - Three-Dimensional Image Capture and Applications VII*, vol. 6056, pp. 34-41.

Vankka, J. and Halonen, K., (2001). *Direct Digital Synthesizers: Theory, Design and Applications*, Kluwer Academic Publishers, Boston.

Vaucher, C. S., (2002). *Architectures for RF Frequency Synthesizers*, Kluwer Academic Publishers.

Viarani, L., Stoppa, D., Gonzo, L., Gottardi, M. and Simoni, A., (2004). "A CMOS smart pixel for active 3-D vision applications," *IEEE Sensors Journal*, vol. 4 (1), pp. 145-152.

Webb, C. E. and Jones, J. D. C., (2004). *Handbook of Laser Technology and Applications*, Institute of Physics, Bristol.

Webster, J., (1999). *The Measurement, Instrumentation, and Sensors Handbook*, CRC Press, Boca Raton, FL.

Webster, J. G., (2004). *Electrical measurement, signal processing, and displays*, CRC Press, Boca Raton, FL.

Weldon, J. A., Narayanaswami, R. S., Rudell, J. C., Lin, L., Otsuka, M., Dedieu, S., Tee, L., Tsai, K. C., Lee, C. W. and Gray, P. R., (2001). "A 1.75-GHz highly integrated narrow-band CMOS transmitter with harmonic-rejection mixers," *IEEE Journal of Solid-State Circuits*, vol. 36 (12), pp. 2003-2015.

Whitaker, J., (2002). *The RF Transmission Systems Handbook*, CRC Press, Boca Raton, Fl.

Wolf, K., Roller, D. and Schafer, D., (2000). "An approach to computer-aided quality control based on 3D coordinate metrology," *Journal of Materials Processing Technology*, vol. 107 (1-3), pp. 96-110.

Worthington, P. and Hancock, E., (2000). "Histogram-based object recognition using shape-from-shading," *Proc. IEEE Computer Vision and Pattern Recognition 2000*, vol. 1, pp. 643-648.

Yang, R. K., Han, X. G., Hao, Y. and Sun, Z. Y., (2007). "Propagation characteristics of infrared pulse waves through windblown sand and dust atmosphere," *International Journal of Infrared and Millimeter Waves*, vol. 28 (2), pp. 181-189.