

Heterodyne range imaging as an alternative to photogrammetry

Adrian Dorrington^{*1}, Michael Cree¹, Dale Carnegie², Andrew Payne¹, and Richard Conroy¹.

¹Department of Engineering, University of Waikato, Private Bag 3105, Hamilton, New Zealand.

²School of Chemical and Physical Sciences, Victoria University of Wellington,
PO Box 600, Wellington, New Zealand.

ABSTRACT

Solid-state full-field range imaging technology, capable of determining the distance to objects in a scene simultaneously for every pixel in an image, has recently achieved sub-millimeter distance measurement precision. With this level of precision, it is becoming practical to use this technology for high precision three-dimensional metrology applications. Compared to photogrammetry, range imaging has the advantages of requiring only one viewing angle, a relatively short measurement time, and simplistic fast data processing. In this paper we first review the range imaging technology, then describe an experiment comparing both photogrammetric and range imaging measurements of a calibration block with attached retro-reflective targets. The results show that the range imaging approach exhibits errors of approximately 0.5 mm in-plane and almost 5 mm out-of-plane; however, these errors appear to be mostly systematic. We then proceed to examine the physical nature and characteristics of the image ranging technology and discuss the possible causes of these systematic errors. Also discussed is the potential for further system characterization and calibration to compensate for the range determination and other errors, which could possibly lead to three-dimensional measurement precision approaching that of photogrammetry.

Keywords: Range, distance, imaging, heterodyne, photogrammetry, metrology

1. INTRODUCTION

Range Imaging is a technique used to produce a digital photograph like output that contains not only color and/or intensity information, but also distance or range information for each pixel. Such measurements are useful for determining the shape, size, location, and surface profile of objects.

There are numerous techniques available to implement a range imaging system, all of which can be divided into the four broad categories of laser scanning, stereo vision, structured light, and imaging lidar. In this paper, we concentrate on the imaging lidar implementation because it has the advantages of no moving parts, and only requiring one illumination and view angle. Although in its relative infancy compared to the other techniques, recent demonstrations of imaging lidar systems with sub-millimeter ranging precisions has shown their potential for high quality metrology applications [1,2].

More specially, we discuss a heterodyne type range imaging system, developed in the Department of Engineering at the University of Waikato (Hamilton, New Zealand), and its implementation as a three-dimensional (3D) metrology system. We first review the system's principle of operation, then describe the metrology of a calibration and test object including resection, calibration, and 3D data projection algorithms. To evaluate its performance and potential we compare these measurements to a photogrammetric survey. Our results show that although the range imaging measurement currently exhibits some systematic errors, with further calibration the performance of the imaging lidar system will start to approach that of photogrammetry.

* a.dorrington@waikato.ac.nz; phone +64-7-858-5062; fax +64-7-838-4835

2. SOLID-STATE FULL-FIELD HETERODYNE RANGE IMAGING

Range imaging systems use an indirect time-of-flight method to measure the distance to the objects in the scene for every pixel simultaneously across the field of view. The three primary techniques used are pulsed, homodyne, and heterodyne based methods. Each technique has advantages and disadvantages, but for the purposes of this paper we will discuss only the heterodyne method because of its higher ranging precision performance that is critical to high precision 3D metrology. For a review of these different range imaging techniques and a more in depth discussion of imaging lidar technologies, see Dorrington et al [1].

The heterodyne image ranging system functions by illuminating a scene with intensity modulated light and collecting the scattered light with a gain modulated camera system, as shown in figure 1. The camera's gain modulation is typically performed with an external (non-mechanical) shuttering device. Both the illumination and gain modulation are in the region of 10 to 100 MHz, and in a heterodyne system, these modulation frequencies are offset slightly. The result is a mixing effect in the camera system and a beat signal that appears as a flashing image at a rate equal to the difference in the modulation frequencies [3].

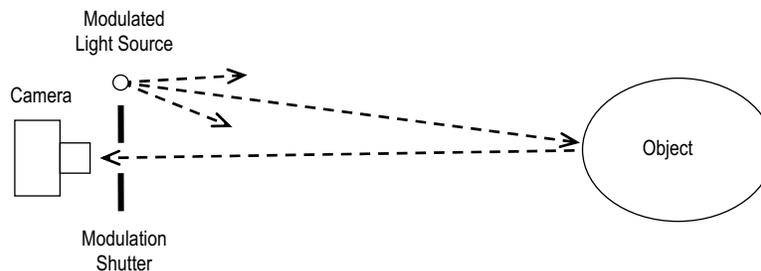


Figure 1 – Imaging lidar operating principle

The propagation of the illumination from the light source to the object and back to the camera causes a delay in the modulation envelope exhibited as a phase change in the modulation signal. The mixing effect in the camera is effectively a frequency down-conversion process and the resulting beat signal contains the relative phase and intensity information between the two modulation signals. This results in the propagation delay of the illumination being encoded in the beat signal phase, hence analysis of the time varying intensity of the individual image pixels will reveal the propagation delay, and with knowledge of the speed of light, the distance to the object.

In our implementation, this illumination is provided by a bank of four laser diodes fiber-optically coupled to an illumination head. This fiber-optic delivery allows the found illumination sources to be arranged in a circular fashion around the imaging lens to provide co-axial illumination (on average) and reduce the effects of shadowing. In addition, the fiber delivery provides a mode scrambling function to laser light for more homogeneous and circular illumination patterns than is achievable directly from the laser diode output. Furthermore, small lenses placed at the end of the fibers allow for optimal control of the illumination divergence to match the imaging lens.

The camera gain modulation is performed with an image intensifier utilizing photocathode modulation. Image intensifiers have a number of drawbacks including a reduction in imaging quality and requiring high voltage power supplies, but for an off-the-shelf solution they are the only practical means of image modulation at 100 MHz. Our image intensifier housing includes a standard Nikon F-mount lens adaptor for flexible lens selection and replacement. Most of our imaging to date is performed using an 80 to 200 mm f/# 4.5 zoom lens. The output of the image intensifier is lens coupled to a high quality digital video camera with a short working distance (15 mm) f/# 1.0 lens. This short working distance and large aperture lens provide good optical coupling efficiency between the image intensifier and the video camera. Even though fiber-optic coupling is a yet more efficient option, we chose lens coupling for the additional flexibility it provides in the laboratory. A more detailed explanation of the hardware set-up can be found in the literature [1,4].

A video sequence of the beating image is acquired on the personal computer and processed to find phase and amplitude information for each pixel. The phase information is then further processed to find distance. However, because of the cyclic nature of the phase measurement, ambiguity arises in the actual distance measured. We typically operate our system with modulation frequencies around 80 MHz (with a 1 Hz beat signal), which means phase measurements are repeated for every 1.875 m (half of the wavelength). This ambiguity can be resolved by performing a second measurement at a slightly difference frequency, 82 MHz for example, and comparing the results. Because the ambiguity distance is different at the second modulation frequency, the true range can be determined [5]. All measurements performed in this paper are ambiguity resolved in this way.

3. CALIBRATION BLOCK & MEASUREMENT METHODS

3.1 The Calibration Block

For the purposes of calibrating the range imager and evaluating its measurement precision, a 3D calibration block was constructed, shown in figure 2. This block measures approximately 800 mm cubed and is shaped to have four separate planar surfaces resembling a small staircase. It has been manufactured by a local cabinet maker from plastic laminated custom-wood with a light gray diffuse finish. The surfaces measure approximately 800 mm by 200 mm, and are separated by approximately 200 mm. Two rows of 8 (30 mm diameter) retro-reflective targets are attached to each surface in a regular grid pattern spaced at approximately 100 mm intervals.

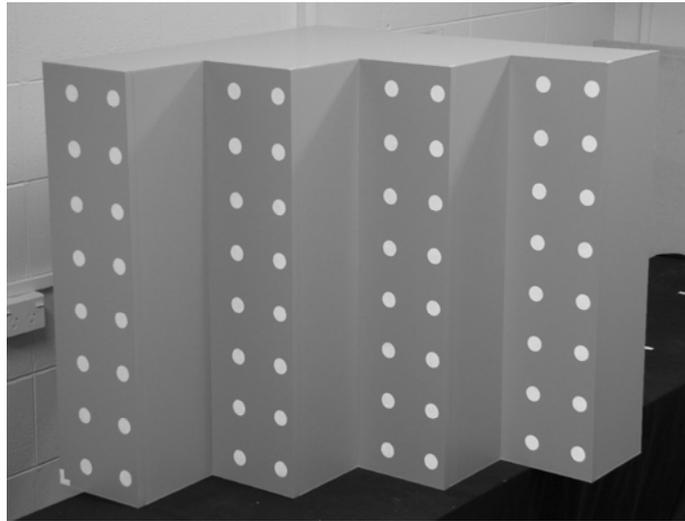


Figure 2 – The 3D target block used to calibrate and evaluate the image ranger used for 3D metrology.

Although it would be preferable to construct the calibration block from a material that is more rigid and temperature stable, this construction is adequate for the purposes of this experiment. To avoid the potential of shape changes from thermal or physical disturbance, both the photogrammetric and image ranger surveys (discussed below) were performed in short order without moving the block.

For reference and calibration purposes the 3D locations for the retro-reflective targets on the calibration block were required. These locations were found with a photogrammetric survey using the PhotoModeler™ software package and photographs from a Canon 300D digital SLR camera with a 28 mm lens. A total of 13 photographs were used, and measurement produced average target precisions of 0.011 mm in-plane and 0.024 mm out-of-plane. The bottom left target (as viewed in figure 2) was designated as the origin, with the x-axis through the target immediately to the right, and the y-axis through the target at the top left.

3.2 Image Ranger Survey

Only one 10 second measurement was required to acquire the 3D target locations of the calibration block using the image ranger. Because of the limiting resolution of the image intensifier currently in use, the digital video camera was set to 2 by 2 binning mode providing an x-y imaging resolution of 512 by 512 pixels. The camera frame rate was set to 29 Hz, providing a total of 290 samples of the 1 Hz beat signal for analysis. With the lens set to a focal length of approximately 80 mm, the calibration block filled the field of view at around 3.5 m distance from the imager. Note that because of the physical nature of the image intensifier, the imaged field of view is circular with approximately a 500 pixel diameter, and the corners of the images from the camera contain no useful data. To provide the best response from the retro-reflective targets and to avoid any occlusions, the calibration block was arranged so the targeted surfaces were approximately perpendicular to the optical axis of the range imager.

The captured video data were processed to derive both intensity and range data. Figure 3 shows an image of intensity data produced by the image ranger (gamma corrected for increased visibility). The approximate location of the target centers was found manually with a graphics editing program, and square regions of 25 by 25 pixels centered at these locations were defined for processing of individual targets. Centroiding was then performed to find the locations of the targets on the image plane as discussed below.

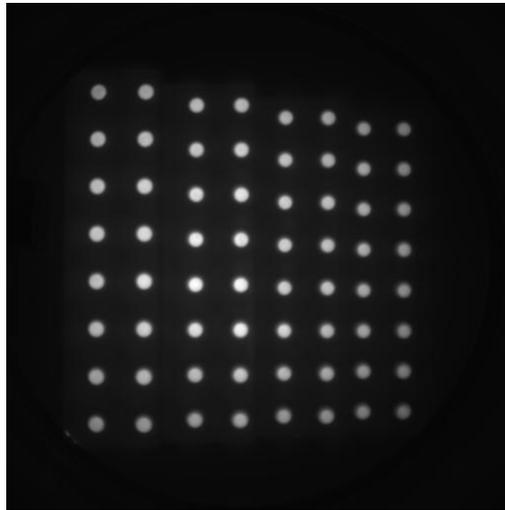


Figure 3 – Intensity data from range imager (gamma corrected $\gamma=2$).

The range data for each target was also processed to improve range determination precision by averaging range data for pixels selected by thresholding the intensity data. An estimate of the single pixel ranging precision can be found by least-squares fitting a plane to the range data and finding the standard deviation of the residuals. Because the target range data are an average of a number of pixels, the target ranging precision improves by the square-root of the number of pixels averaged. The average individual pixel one-sigma ranging precision was found to be 0.48 mm, and the target ranging precision 0.071 mm. Close inspection of the range data for a given target showed some systematic errors, discussed below, so further ranging calibration could easily improve these precision figures.

3.3 Centroiding, Resection, and Lens Calibration

Our calibration method is a dual-component approach similar to one suggested by Kahlmann and Ingensand [6] where the imaging distortions are calibrated independently from the distance measurement distortions. In contrast to the one-step “black box” approach, the dual-component calibration coefficients can be related to real world physical properties of the system. Useful information can be derived by examining these coefficients in an effort to determine the physical reasons for the dominant causes of error, thereby allowing future system refinement to be targeted at the most important areas.

A simple weighted average grey scale centroiding algorithm [7] was performed on the intensity data to locate the sub-pixel target locations on the image plane. Note that the return from the targets in the far corners is almost a factor of 5

below those in the center of the image, but this does not adversely affect the centroiding precision because the intensity data is not quantized in the usual way. These intensity data are generated with an inner-product phase determination algorithm (single bin Fourier transform type calculation) [5] from a sequence of 290 video frames. This process has an averaging type effect providing quantization levels much smaller than those of a single image frame from our 12-bit camera.

At this point the target locations were split into two sets, one for calibration and the other for performance evaluation. A checker-board type pattern was used splitting off the odd numbered rows of targets in the odd numbered columns and the even numbered targets in the even numbered columns.

The centroid values along with the known object target locations for the calibration data set were used in a Direct Linear Transform (DLT) [7] resection process to find the exterior camera parameters of the range imager, which are shown in table 1. For the 3D projection reconstruction procedure (explained below) an additional “distance offset” calibration parameter is required. This arises because the image ranger has delays in both the electronic and optical signal paths that cause the measured distance to appear greater than it actually is. The distance offset parameter is used to compensate for this effect by providing a simple subtractive correction. It is found by subtracting the distance to the DLT camera location from the measured distance to target at the origin.

Table 1 – Exterior camera parameter results of resection relevant for 3D range image data projection.

Parameter	Value
Principle point	256.38×250.21 pixels
Principle distance	84.10 mm
Camera location	$369.92 \times 192.72 \times -3450.76$ mm
Distance Offset	634.44 mm

Using these camera parameters, the 3D photogrammetric reference data were re-projected to the image plane to find the expected centroid locations for comparison. These projected locations and 100 times scale residual vectors are shown in figure 4. The residuals have a standard deviation of 0.17 pixels, and generally have a random pattern with a tendency to point radially outward from the principle point indicating some slight radial lens distortion. However, because the residuals are of the same order as the expected centroiding precision, and because the majority of the 3D measurement errors are in the z direction (explained below) no lens distortion correction was performed.

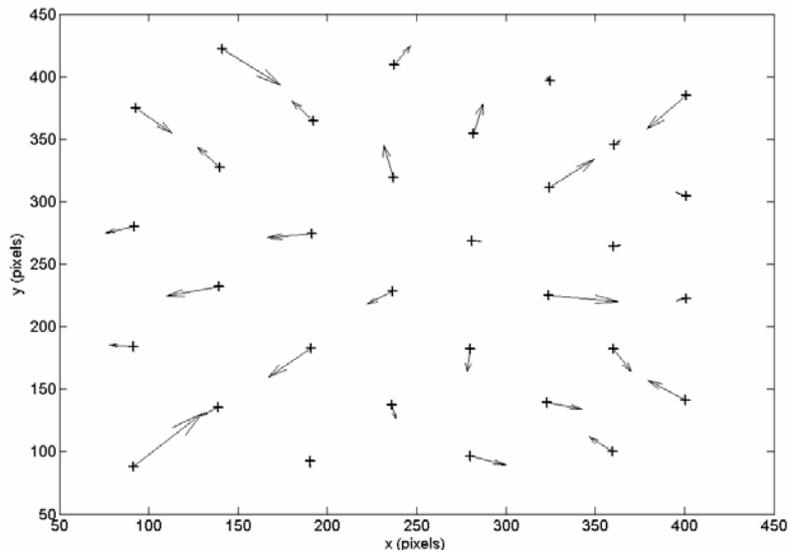


Figure 4 – Projected image plane target locations and residual errors (100 times scale) to centroided targets.

3.4 3D Projection and Reconstruction

With the calibration completed, the second performance evaluation target set was used to find the 3D locations of the targets for comparison to the known locations. The projection of image plane centroid locations to 3D object space locations is a relatively straight forward process, and is illustrated in figure 5. First we define the perspective center as the origin, and the optical axis of the camera to lie on the z-axis with positive z in the direction of viewing. For each centroid location a “centroid vector” can be calculated from the 3D centroid location in object space to the perspective center or origin. Because of the geometry we have defined, the object space coordinates of the centroids are easily calculated as the image plane centroid location in millimeters in x and y, and the negative of the principle distance in z. The object space location of the target can be calculated as the centroid unit vector (located at the origin) multiplied by the distance value for the target. Note that the true target distance is the measured range value minus the sum of the centroid vector length and the distance offset calibration value.

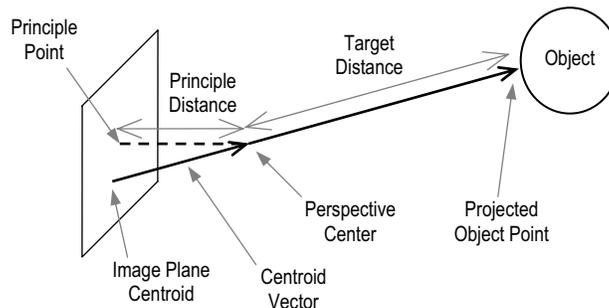


Figure 5 – Representation of the 3D projection process to calculate object locations from image ranger data.

4. MEASUREMENT COMPARISON

Because of our chosen geometry for the 3D projection reconstruction, the measured target locations have a different orientation and location compared to the photogrammetrically measured reference locations. To make a meaningful comparison, the image ranger measurements were translated and rotated using a least-squares minimization algorithm to match the points back to the reference location. The 3D residuals were then found as the difference between the locations in the two sets of data and have been plotted in figure 6. The standard deviation of these residuals are 0.54 mm, 0.53 mm, and 4.54 mm in the x, y, and z dimensions respectively.

The precision of approximately half a millimeter in the x and y dimensions is quite reasonable given the low resolution of the sensor and the lack of radial lens calibration. This could be improved further with a higher x-y resolution system, improvements in the centroiding algorithm, and implementation of lens calibration. For the purposes of this paper however, these topics have not been pursued because the errors in the z direction are significantly larger and more dominant.

Close examination of the residuals plotted in figure 6 shows a pattern in the z direction. Generally, their distance measurements tend to be too long in the center of the imaging area and too short at the extents. This is confirmed by plotting the z-axis residuals against the target centroid radius on the image place (distance from the centroid to the principle point), as shown in figure 7, and noting that there is a relationship. Because of the predictable nature of this error, there is a good chance it can be corrected with some form of calibration. The physical mechanism behind these errors is an effect in the image intensifier called “irising”. Irising arises because of the physical layout and the electrical properties of the photocathode. The photocathode electrode, to which the modulation signal is applied, is connected around the circumference photocathode itself to avoid interference with the photo-sensitive imaging area. Because of the electrical properties of the photocathode, there is a particular speed at which the modulation signal propagates to its center, which results in a phase delay of the modulation signal [9]. Without compensation, this phase delay manifests itself as a propagation delay of the illumination and interferes with the distance measurements.

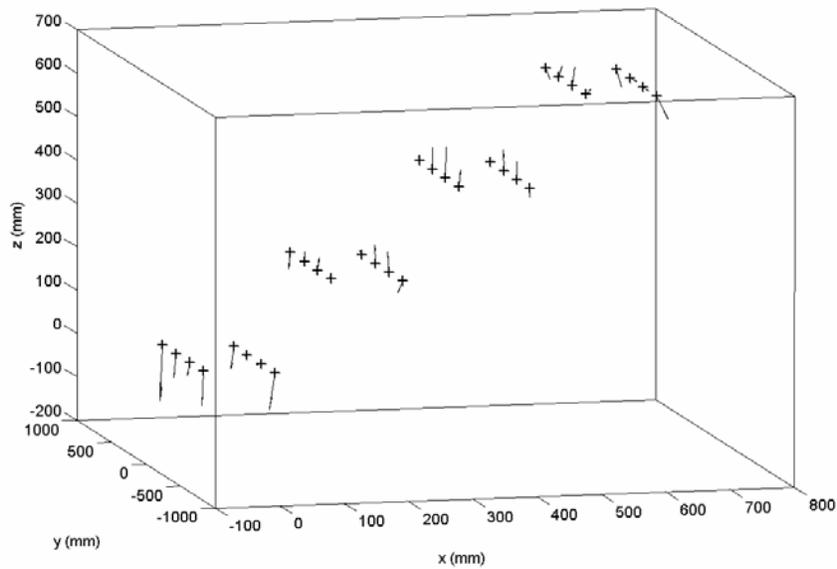


Figure 6 – The 3D reference locations of the targets and the residual vectors (10 times scale).

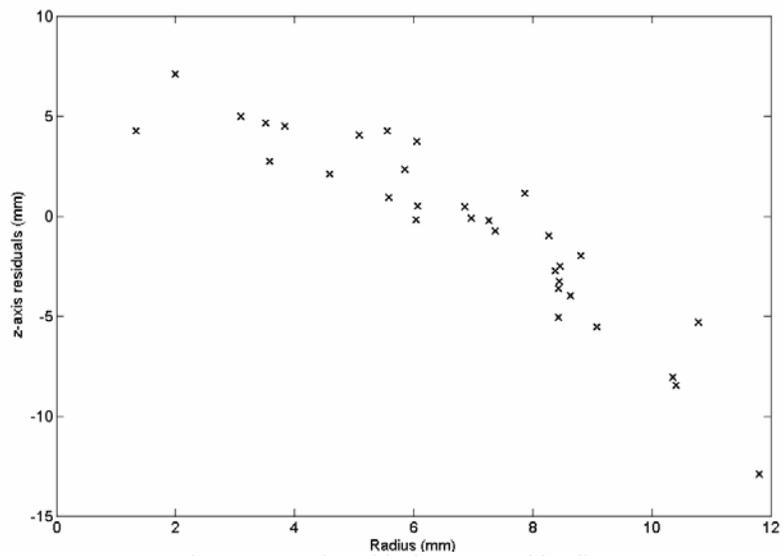


Figure 7 – z-axis residuals vs. centroid radius.

A further, but less significant, error arises due to a small amount of intensity to distance cross-talk in the range imager. This cross-talk causes a distance error related to intensity. Close examination of the range data for one retro-reflective target (shown in figure 8) shows curvature on what should be a planar surface. This curvature is due to the distance errors caused by varying intensity across the target. Again, it is probable that these errors could be corrected by characterizing the relationship between returned light intensity to distance errors.

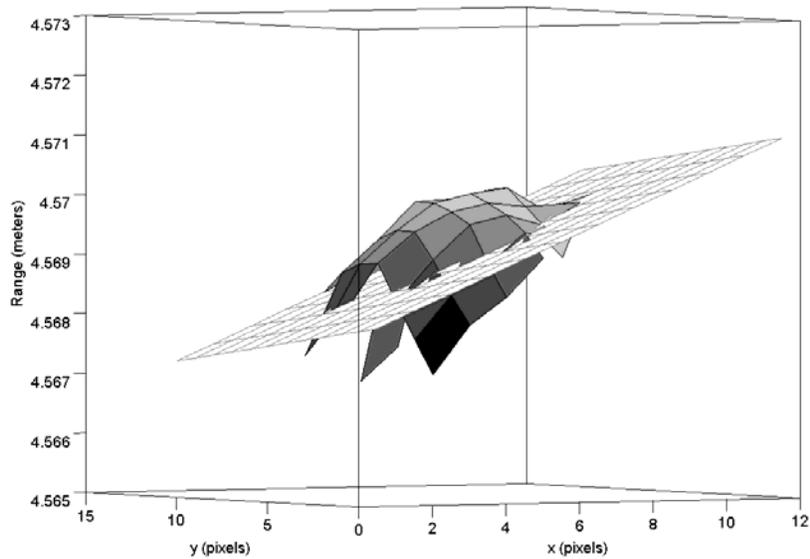


Figure 8 – Plot of range data for an individual target showing distance dependency on intensity.

5. CONCLUSION

We have shown a comparison of three-dimensional measurements from both standard photogrammetry and a solid-state full-field range imaging system. Although the range imaging system exhibits several errors, these errors are systematic in nature and could be calibrated out with further work. Our results show the potential for image ranging apparatus in high quality metrology applications. We do not claim the image ranging is a replacement for photogrammetry or other metrology techniques, simply an alternative where the application can make good use of its advantages. The primary advantages are the need for only one viewing angle, the presence of scale information inherent in the measurement (with a calibrated field of view), fast image acquisition (around 10 seconds), and simple and fast on-line data processing. An additional advantage, compared to photogrammetry, is the ability to obtain a full-field high spatial resolution range image that provides a continuous surface profile across the field of view.

ACKNOWLEDGEMENTS

Funding for this research was provided in part by the Foundation for Research, Science and Technology (New Zealand) through the New Zealand Science and Technology Postdoctoral fellowship, and the Pre-seed Accelerator Fund.

REFERENCES

- 1 A. A. Dorrington, D. A. Carnegie, and M. J. Cree, "Towards 1 mm depth precision with a solid-state full-field range imaging system", *SPIE Vol. 6068 - Sensors, Cameras, and Systems for Scientific/Industrial Applications VIII*, Paper Number 6068-22, 2006.
- 2 J. Busck and H. Heiselberg, "Gated viewing and high –accuracy three-dimensional laser radar", *Applied Optics*, Vol. 43, pp. 4705-4710, 2004.
- 3 A. A. Dorrington, *Range sensing system*, International patent PCT/NZ2004/000070, April 6, 2004.
- 4 D. A. Carnegie, M. J. Cree, and A. A. Dorrington, "A high-resolution full-field imaging system", *Review of Scientific Instruments*, Vol. 76, CID 083702, 2005.
- 5 A. A. Dorrington, M. J. Cree, R. M. Conroy, A. D. Payne, and D. A. Carnegie, "Achieving sub-millimeter precision with a solid-state full-field heterodyning range imaging camera", submitted to *Measurement Science and technology*.
- 6 T. Kahlmann and H. Ingensand, "Calibration and improvements of the high-resolution range-imaging camera SwissRanger™", *SPIE Vol. 5665 - Videometrics VIII*, pp. 144-155, 2005.
- 7 M. R. Shortis, T. A. Clarke, and T. Short, "A comparison of some techniques for the subpixel location of discrete target images", *SPIE Vol. 2350 – Videometrics III*, pp. 239-250, 1994.
- 8 E. M. Mikhail, J. S. Bethel, and J. C. McGlone, *Introduction to modern photogrammetry*, John Wiley & Sons, New York, 2001.
- 9 A. D. Payne, A. A. Dorrington, M. J. Cree, and D. A. Carnegie, "Image Intensifier Characterisation", Accepted for *Image and Vision Computing New Zealand*, Great Barrier Island, New Zealand, November 27-29, 2006.