[10] J. Shamma and D. Whitney: *A Method for Inverse Robot Calibration*, ASME Journ. of Dyn. Syst., Meas., and Control (109), March 1987, pp. 36 – 43.

[11] E. Trucco and R. Fisher: *Experiments in Curvature-Based Segmentation of Range Data*, to appear in IEEE Trans. Patt. Anal. Mach. Intell., 1994.

[12] S. J. White: *High Speed Scanning Method and Apparatus.* U.S. Patent No. 4,498,778, Technical Arts Corp., 1985.

[13] T. Wohlers: *3-D Digitizers*, Computer Graphics World, July 1992, pp. 73 – 77.

surface type	δz	e	σ_e	e_a	σ_{e_a}
matt black,50	0.0013	-0.160	0.095	0.167	0.095
red paint,50	0.0043	0.126	0.106	0.154	0.061
anodized black,100	0.0236	-0.096	0.079	0.100	0.074
polished alum.,100	0.0029	0.066	0.147	0.132	0.093

Table 1: test results (see text). δz is given in mm/quanta.

our method for direct calibration of small-workspace sensors proves simple, practical, and capable of supporting satisfactory accuracies. Second, the consistency tests can improve dramatically range measurements in the presence of highly reflective surfaces and holes, and eliminate most of the wrong measurements arising from spurious reflections.

Acknowledgement
We gratefully acknowledge the invaluable assistance of the staff of the Department of Artificial Intelligence Workshop, directed by Mr. David Wyse. Thanks to Forest Zhang for providing useful references. This work was supported by ACME under Grant GR/H/86905 "Model Acquisition from Multiple Range Views".

References

[1] P. Besl: *Active, Optical Imaging Sensors*, Machine Vision and Applications, 1988, pp. 127 – 152.

[2] H. R. Everett: *Survey of Collision Avoidance and Ranging Sensors for Mobile Robots*, Robotics and Autonomous Systems (5), 1989, pp. 5 – 67.

[3] P. Saint-Marc, J.-L. Jezouin and G. Medioni: *A Versatile PC-Based Range Finding System*, IEEE Trans. Robot. Autom. (7), no. 2, April 1991, pp. 250 –256.

[4] R. Fisher: *The Design of the IMAGINE2 Scene Analysis Program*, in 3D Model Recognition from Stereoscopic Cues, J. Mayhew and J. Frisby (eds.), MIT Press, 1991.

[5] R. Fisher, D. K. Naidu and D. Singhal: *Rejection of Spurious Reflections in Structured Illumination Range Finders*, patent pending, 1993.

[6] R. Fisher, D. K. Naidu and D. Singhal: *Rejection of Spurious Reflections in Structured Illumination Range Finders*, Proc. 2nd Conf. on Optical 3-D Measurement Techniques, Zurich, October, 1993, pp. 467 – 474.

[7] R. Fisher, E. Trucco and A. Fitzgibbon: *Direct Calibration of Small-Workspace, Triangulation-Based Range Finders*, patent pending, 1993.

[8] D. K. Naidu and R. Fisher: *A Comparative Analysis of Algorithms for Determining the Peak Position of a Stripe to Subpixel Accuracy*, Proc. British Machine Vision Conf., 1991, pp. 217 – 225.

[9] R. Fisher, E. Trucco, M. Waite, A. Fitzgibbon, M. Orr: *IMAGINE: a 3-D Vision System*, Sensor Review (2), no. 13, April 1993, pp. 23 – 26.

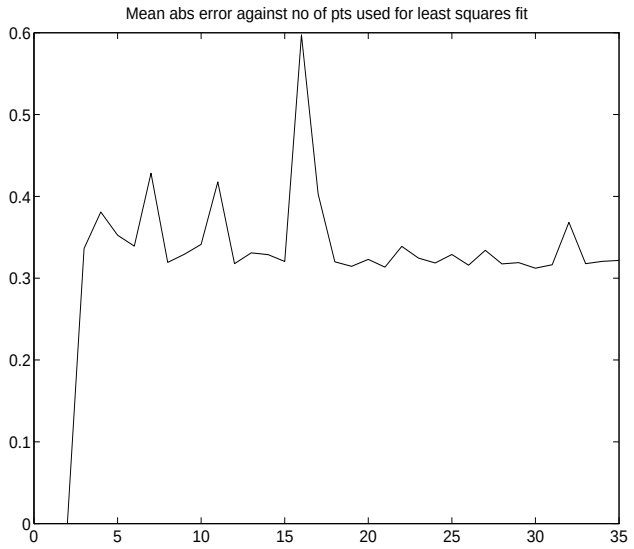

Figure 6: Mean absolute error of least squares plane fit against number of points used.

Figure 7 shows a range image, with consistency tests enforced, of the polished aluminium block with holes which caused the spurious surfaces in Figure 1. The larger holes' diameters are 18mm, the smaller ones' 14mm; depths varied between 9 and 13mm. The dramatic rejection of spurious range values is evident. Some of the true range points have also been eliminated, which has caused a slightly more ragged appearance to the object surface; notice however that the height of the remaining range points has been correctly estimated. In spite of the strong reflections, there are also enough data to estimate the real depth of all holes.

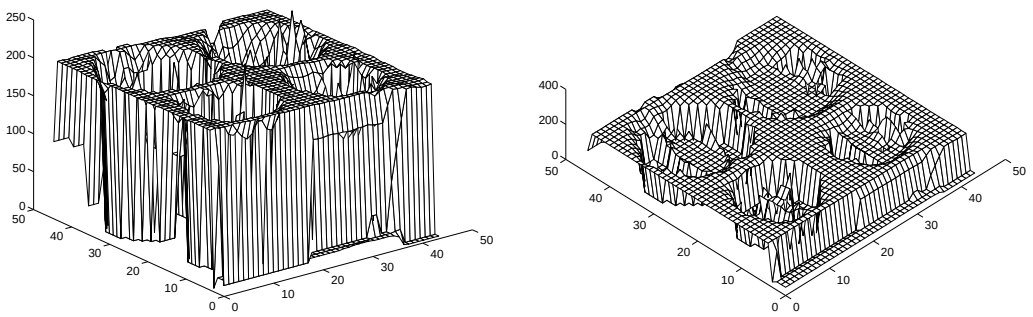

Figure 7: Range image of block with holes with consistency tests enforced.

6 Conclusions

We have described a triangulation-based range scanner built with off-the-shelf, low-cost components, and presented a direct, local calibration technique. We have also sketched some consistency tests which proved very helpful in acquiring images of highly reflective surfaces. To demonstrate the validity of the techniques, we have reported concisely the results of several experiments. We believe this paper offers two main contributions. First,

Illumination direction constraint. Given that the direction of the plane of laser light is fixed, it is not possible for a ray of light to intersect the surface twice. When more than one point is observed, all points are eliminated.

Observable surface constraint. One constraint that eliminates many spurious range surfaces is the requirement that the visible surface portions must face the observing sensor (otherwise the surface could not have been seen). Hence, any local surface point $\vec{P}_\alpha(t)$ whose normal $\vec{n}_\alpha(t)$ satisfies

$$\vec{n}_\alpha(t) \circ (\vec{P}_\alpha(t) - \vec{O}_\alpha) > 0$$

where $\vec{O}_\alpha$ is the origin of the camera reference frame, should be rejected. This constraint is independent of the number of cameras used.

Consistent surface constraint. If a true point is observed by both cameras, then the range values $Z_L(t)$ and $Z_R(t)$ from both cameras should be the same within a tolerance τ_d, which is chosen on the basis of the noise statistics. In addition, the surface normals of the surface as observed from the left and right cameras should be the same, again within a tolerance τ_n chosen on the basis of the noise statistics of true range images. Notice that τ_n requires careful setting, since surface normals are related to the first-order derivatives of the data and thus are more affected by noise.

Unobscured-once-viewed constraint. If a point was visible by only one camera, there must be a valid point seen by the other camera that obscures the first point. Hence, any points that are visible to one camera and are not obscured relative to the other camera, yet were not observed, are likely to be spurious and are removed.

5 Experimental results

The measurements reported in Table 1 are indicative of the accuracy of our striper using direct calibration. The table gives the z quantisation step δz (using 256-level resolution), the mean error, e, its standard deviation, σ_e, the mean absolute error, e_a, and its standard deviation, σ_{e_a}, all in mm, measured using both cameras and accurately known planes of different materials placed at different heights (material and height, in mm, are specified in the leftmost column). Comparable accuracies are obtained using each camera individually; however, no-data regions appear more frequently, as the visible portion of the object surface is in general smaller for one single camera.

In another series of tests, we measured sloping planes spanning the whole sensor workspace, again finding accuracies similar to those reported above. When the slope angle could not be accurately determined, we found the best fitting plane (in the least square sense) for an increasing number of data points organised in regular grids. We then considered the plane associated with the minimum mean absolute error, and used that to compute error statistics. We mention only one representative example. Figure 6 shows the mean absolute error (in mm) plotted against the square root of the number of points used for the fit for a black-painted slope with $\alpha \cong 30$ deg. The minimum-error fit occurred using 30x30=900 points, but it can be seen that the variation is nearly always small given the sensor's estimated accuracy. The corresponding mean absolute error was 0.318mm (standard deviation 0.011mm); the mean error (signed) was 0.023mm (standard deviation 0.016mm).

We noticed that the error magnitude remains constant throughout the field of view, whereas it tended to increase towards the periphery (considering pixels falling in the same image regions) with a model-based calibration technique adopted previously.

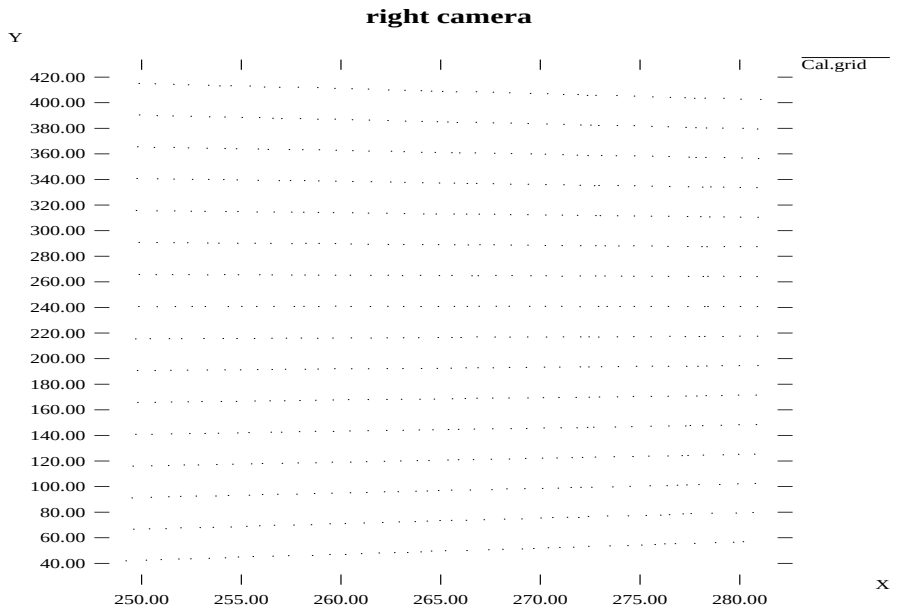

Figure 4: Calibration grid for 80 height levels (steps).

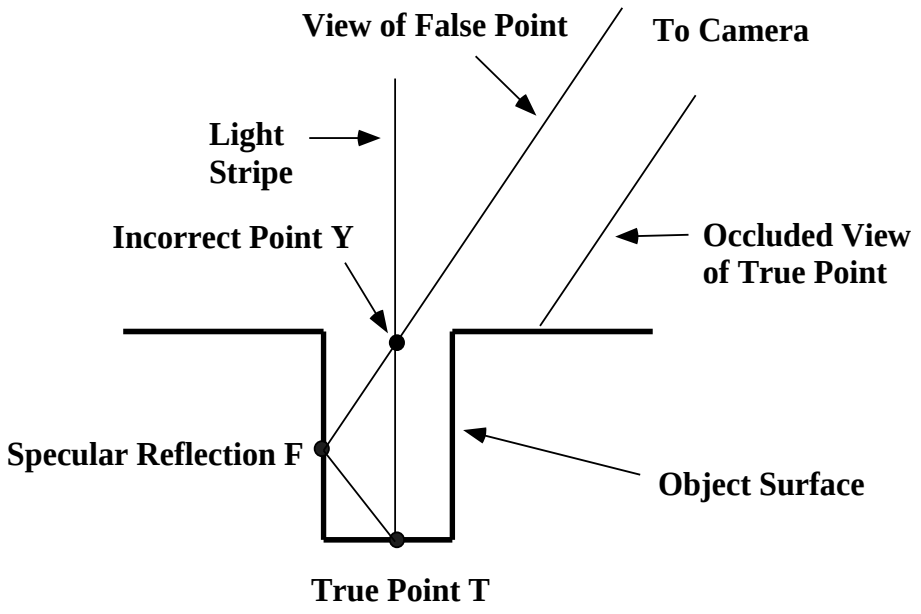

Figure 5: How specular reflections cause false range values.

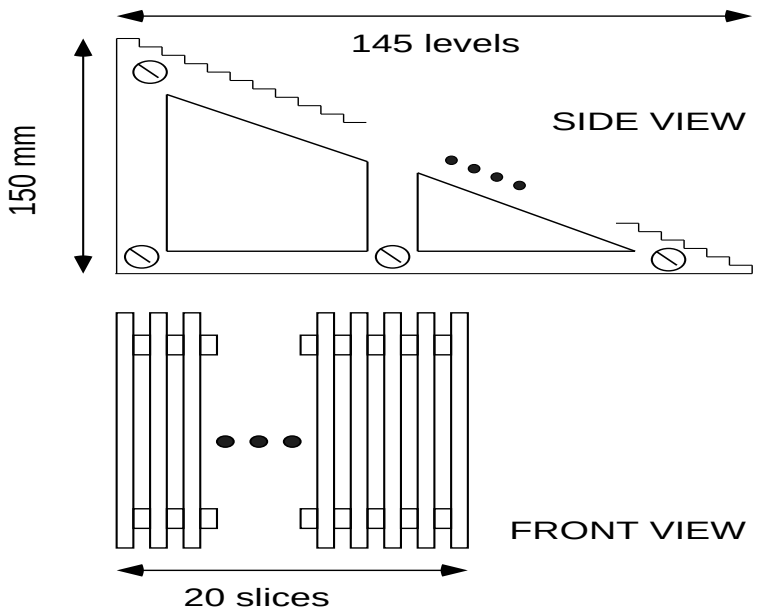

Figure 3: The calibration block.

Stage 2: building image-world maps. In the second stage, the calibration grids are inverted and interpolated to obtain a complete look-up table for each camera. Using a least-square linear interpolation with 5 calibrated points, *each* image pixel is associated to a 3-D point within the calibrated workspace.

Our DC procedure is simple, automatic and fast (currently about 30 minutes for 140 height levels). It also allows a simple range measurement algorithm. When acquiring a range image, the bounding-integer pixel coordinates obtained from the subpixel coordinates of each stripe pixel are used to index in the lookup table of each camera. The position of the 3-D point corresponding to the observed subpixel position is computed by linear interpolation between the bounding pixels. Although speed was not a research objective, the acquisition rate is a few stripes per second - not a despicable one given the equipment used. Speed could be greatly improved with the use of a synch generator, analogue stripe detection hardware and a DSP chip for the range value calculations.

4 Discarding inconsistent points

How do reflections from specular surfaces cause spurious range values? Figure 5 shows a cross section, taken perpendicularly to the light plane, through a rectangular hole in the object surface. Suppose that the light stripe is observed after reflection from the specular surface of the hole: the specular reflection at point **F** is observed rather than the true point **T**. The false point might be chosen because it is brighter (often possible on specular surfaces) or because the true point is hidden. Since all observed points must lie in the plane of laser light, the height of point **Y** is incorrectly recorded.

The false range surface shown in Figure 1 resulted from this phenomenon occurring at many positions along each of many stripes. The tilting false surface arises because, as the stripe moves away from the wall, the triangulated false position moves further away from the true surface. This simple false-surface pattern arises from the simple rectangular hole geometry; more complex holes or combinations of specular surfaces produce more complex artifacts.

The rejection of false range values is based on the consistency constraints sketched below [5, 6]. Any points that do not satisfy the constraints are eliminated.

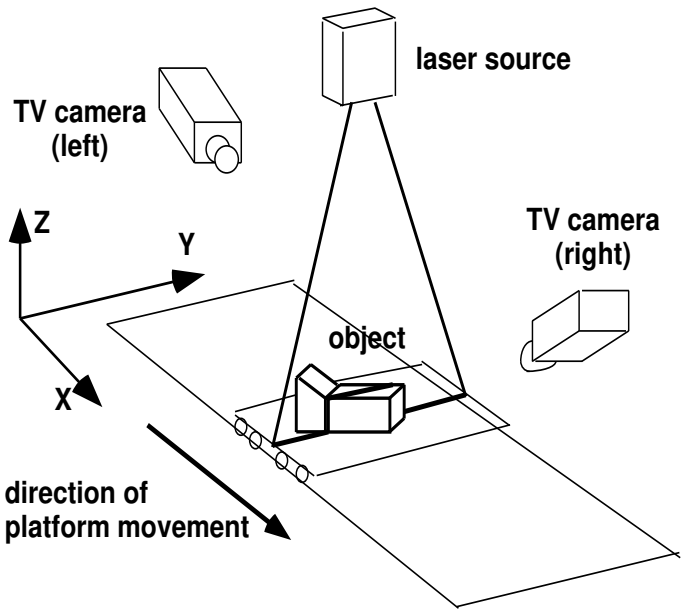

Figure 2: Architecture of the range sensor.

are stored in a Datacube as 512x512 frames. A difference of one millimeter in the vertical direction corresponds roughly to a one-pixel difference in the images. Several parameters can be controlled by the user, including image scaling factors, depth quantisation, image resolution, and the depth interval to be scanned.

3 Direct calibration

Our DC method is based on a simple idea. If the camera coordinates of a sufficiently dense grid of workspace points (called *calibration grid*) can be measured accurately, then the position of any point in the workspace can be obtained by inverting the resulting world-to-camera maps and interpolating between surrounding points [7]. We have implemented this idea in a two-stage procedure.

Stage 1: building the calibration grids. In the first step, a calibration grid is built for each camera. We have designed and built a calibration block (Figure 3) consisting of 145 steps, each 2mm in length and 1mm in height. In order to detect calibration points in the Y direction (refer to Figure 2), the block is formed by 20 parallel slices spaced regularly. The only operator intervention required is to place the block on the striper's platform so that the laser stripe falls entirely on the top surface of the lower step. The block is then advanced automatically 2mm at a time, so that the stripe is observed by both cameras on each of the 145 steps. For each position, the stripe appears as a linear sequence of segments (corresponding to the top surfaces of the step's slices), and the position of the segments' centers is detected to subpixel accuracy [8] and recorded. The block's size is designed so that every observable point in the stripe plane lies no farther than 1mm in range (z axis) and 4mm along the stripe (y axis) from the nearest calibrated point. The slope of the block allows simultaneous calibration of both cameras of our sensor. An example of the resulting grid of calibration points, measured for 80 heights, is shown in Figure 4. The x axis shows the image position of the calibration points, which depends on the height of the block level on which the stripe impinges. The y axis shows the points' image y position. The slight irregularity in the x direction, largely compensated for by second stage 2 (interpolation), is a result of irregular changes in the shape of the stripe profile as the stripe sweeps the CCD sensor.

confused with the primary signal, in which case false range values will result. This effect can render the sensor unusable. For instance, range images of shiny surfaces with holes may contain spurious peaks or whole surfaces protruding from the holes. Figure 1 illustrates this effect in two range images of a polished-aluminium block with holes. The images were acquired using two opposing cameras independently (see Section 2). To obviate this problem, some users of industrial scanners simply coat surfaces with a matt white substance, e.g. tempera paint, which can rinse clean from most parts [13]. This may however be unacceptable, e.g. whenever very high accuracies ($100\mu m$ or less) are required.

We sketch in Section 4 some *consistency tests* that eliminate most of the spurious range values [5, 6]. The key observation is that specular reflections produce range values depending on camera position. Hence, the range values obtained from each camera can be compared, and points leading to inconsistent range values eliminated.

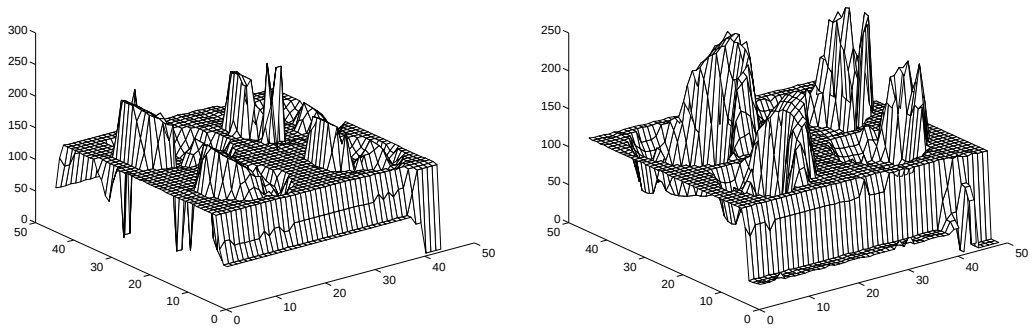

Figure 1: Two range images of polished aluminium block with holes. Specular reflections make spurious surfaces appear.

Our direct calibration and measurement consistency techniques have been implemented and tested using the laser striper we built in the framework of the IMAGINE research project, aimed at recognition of complex 3-D objects from range data [4, 9]. The striper was meant primarily to support surface-based segmentation [11] by achieving good accuracies with low-cost, off-the-shelf components, a feature of major interest in itself. Indeed the accuracy and repeatability of our system is currently in the region of 0.15mm, a very good result for such a low-cost system (*cf.* for instance the similar sensor described in [3], with a reported accuracy of 0.25mm). Section 5 illustrates briefly the sensor's performance after direct calibration and demonstrates the effect of our consistency tests.

2 Sensor operation principles

The architecture of our range sensor is sketched in Figure 2. The object to be scanned sits on a platform moved on a linear rail by microstepper motors under computer (Sun3) control through a Compumotor CX interface with in-built RS-232C interface (in part, to simulate the data acquired from a part passing conveyor belt under a range sensor). One microstep is $6.6\mu m$, and the nominal positioning accuracy and repeatability is 2μ steps. Objects must be contained in a parallelepipedal workspace about 15cm each side. The object is moved through the path of a planar laser beam (output power 2mW at 632.8nm). The curve (stripe) resulting from the intersection of the laser beam with the object surface is observed by two opposing cameras (off-the-shelf 577x581 Panasonic BL200/B) mounted about one meter from the platform. This camera arrangement limits the occlusion problem and is essential for some of our measurement consistency constraints. The acquired images

Direct Calibration and Data Consistency in 3-D Laser Scanning

E. Trucco‡, *R. B. Fisher, A. W. Fitzgibbon*
Department of Artificial Intelligence, University of Edinburgh,
Edinburgh, Scotland
‡Department of Computing and Electrical Engineering, Heriot-Watt University,
Edinburgh, Scotland

Abstract

This paper addresses two aspects of triangulation-based range sensors using structured laser light: calibration and measurements consistency. We present a *direct calibration* technique which does not require modelling any specific sensor component or phenomena, therefore is not limited in accuracy by the inability to model error sources. We also sketch some *consistency tests* based on two-camera geometry which make it possible to acquire satisfactory range images of highly reflective surfaces with holes. Experimental results indicating the validity of the methods are reported.

1 Introduction

This paper addresses two aspects of the acquisition of range data with the popular triangulation-based range sensors using structured laser light [1, 2, 3, 12]: *direct calibration* and *measurement consistency*.

The commonly adopted *model-based calibration* estimates the parameters of the geometric transformation that back-projects any points of the image plane of each camera onto the laser plane [1, 2]. This requires a valid closed-form model of the sensor components and basic phenomena, including at least the position, orientation and intrinsic parameters of the cameras, and the position of the light plane. The higher the measurement accuracy required, however, the more phenomena the model must include (e.g. lens distortion, image center, scale factor), thus becoming significantly complicated. Some phenomena may always remain elusive. Finally, model-based calibration procedures can be long and tedious.

After some experience with model-based calibration, we devised an alternative method called *direct calibration* (henceforth DC), reminiscent of the "black-box" inverse calibration of robotic manipulators [10]. The method consists in measuring the image coordinates of a grid of known workspace 3-D points, then building lookup tables for the whole image by interpolation. An immediate advantage is that there is no need to model *any* phenomena, since all phenomena are *implicitly* accounted for. The overall accuracy of the method is therefore limited only by the repeatability of the equipment and of the stripe detection algorithm, not by shortcomings of the model. The only assumption is that the characteristics of the scanner do not vary between calibration and acquisition, as it is to be expected of any calibrated device. We have devised a simple, fast and automatic procedure implementing DC on our range sensor. The sensor itself is sketched in Section 2. Section 3 describes our calibration technique.

Many potential applications of range finders are in industrial settings, where object surfaces are often made of polished metal or plastic, and are likely to reflect laser light specularly. When surfaces have a specular component, noisy reflections of the main laser stripe may appear in the images observed by the cameras. These reflections can be easily