IN SITU DETERMINATION OF RANGE CAMERA QUALITY PARAMETERS BY SEGMENTATION

Wilfried KAREL[1], Peter DORNINGER[2], Norbert PFEIFER[2]
[1]Christian Doppler Laboratory for "Spatial Data from Laser Scanning and Remote Sensing" at the Institute of Photogrammetry and Remote Sensing (IPF)
Vienna University of Technology, Austria
wk@ipf.tuwien.ac.at
[2]Institute of Photogrammetry and Remote Sensing (IPF)
Vienna University of Technology, Austria
{pdo,np}@ipf.tuwien.ac.at

KEY WORDS: Range Imaging, Range Camera, Quality, Segmentation, In Situ.

ABSTRACT
The present article documents the derivation of quality parameters from raw range imaging data captured under simplified real-world conditions. A sequence of thousand images showing cuboids from constantly changing points of view and in different camera attitudes serves for evaluation. The data are assigned to plane segments by means of a highly robust segmentation algorithm, and without user interaction. The according residuals are evaluated statistically with respect to known error factors. These are the observed distance, the signal intensity, the position in the field of view, and the angle of incidence.

1 INTRODUCTION

Range Imaging (RIM) augments the variety of optical 3-d measurement techniques being a rather new method with remarkable characteristics. A bundle of distances is determined simultaneously at every pixel of a two-dimensional sensor array of a range camera. These ranges are deduced from the time it takes an emitted signal to return to the device. As opposed to the self-evident approach to register the time-of-flight directly using light pulses, the indirect way employing continuously modulated signals has already achieved a feasible development state. By sampling the return at every quarter of the modulation period, the phase shift and hence the object distance are deduced.
RIM combines and supplements properties of other commonly used optical 3-d measurement techniques (i.e. digital imaging and laser scanning). A unique benefit can be stated concerning the registration also referred to as the reconstruction of the exterior orientation. Contrary to interferometry and triangulation by image matching or projected light patterns, RIM does not require multiple observation rays to determine a point's position and thus minimizes the number of device setups necessary to capture complex objects. Unlike laser scanning, the area-based gathering of point data by range cameras allows for their application to kinematical scenes and for a mobile utilization without the demand for inertial and / or positioning devices. The compact dimensions of range cameras further increase their fitness for mobile use. Moreover, RIM supersedes the need for time-consuming image correlations. Holding an array sensor featuring frame rates of some fifty Hertz, RIM offers high data rates. Finally, the purchase costs of range cameras are comparable to the ones of digital amateur SLR cameras and thus far below those of laser scanners.
Besides these promising advantages, severe drawbacks have to be mentioned for current RIM systems. In addition to the partially proved low temporal instrument stability, primarily the distance measurement precision and accuracy are still limited due to large systematic and random errors.

Moreover, as current continuous-wave modulating range cameras merely employ a single modulation frequency and as the emitted radiation ought to be eye-safe, the range of (unambiguously) determinable distances is limited at the upper end by half the modulation wave length (in general smaller than ten meters) and by the maximum admissible illumination power, and at the lower end by reason of physical restrictions of current sensor designs [1].
RIM is on the way to the integration in off-the-shelf products in the domain of applications with lower demands on quality [2]. However, its viability for more sophisticated tasks has not been investigated fully yet.

1.1 Related Work: Determination and Separation of Error Sources in the Laboratory

The straightforward way to detect and quantify RIM error characteristics is to set up laboratory experiments according to the device behavior predicted by theory. Employing intensity images of symmetrical targets with high image contrast (sharp edges), the calibration of the camera objective may be done prior to and separated from the inspection of the sensor characteristics. Camera objectives have been widely investigated, resulting in different approaches for calibration. For the RIM camera SR-3000, [3] reports a highly dynamic principal point. [4] describes the promising way to calibrate a RIM camera objective employing intensity images of a test field equipped with active NIR LEDs.
Applying the resulting objective calibration parameters, the bare sensor characteristics may be inspected using specific experiment setups. Recent studies discriminate the impact of

- nominal distance
- internal temperature / self-induced heating
- external temperature
- integration time
- angle of incidence
- position in the field of view (FOV), and
- target reflectivity

on the observed distances [4-8]. The experiments carried out under known laboratory conditions result in sophisticated calibration functions that allow for the reduction of systematic effects and for the estimation of observation qualities. The investigations show up drifting effects and complex correlations between the error factors. Up to now, the effects of these error sources could however not be separated clearly from each other. Furthermore, the long-term stability of these effects is not reported in literature.

1.2 Objective

It is an open question whether high measurement qualities in real-world applications can be achieved deploying the calibration from a laboratory, as this assumes the predictability i.e. long-term stability of the camera properties.
In the domain of photogrammetry it is common practice to determine systematic effects of unstable camera parameters by means of an extended adjustment of in situ observations. This way, the interior orientation together with the unknown object points and all other parameters are identified at the same time. From the set of possible parameters a subset is selected for use according to the data characteristics. This approach has the advantage that temporally volatile device parameters are determined for the time of data capture.
The present article examines RIM data from this point of view – considering RIM cameras as devices of dynamic nature. It concentrates on the determination and quantification of a selected set of known error sources of RIM employing a data set captured under simplified real-world conditions. This is done by means of residuals gained by assigning the data points to planes in object space. The residuals are evaluated with respect to the error factors under investigation. Neither the exterior orientations of the range images are computed, nor is the tracking of object

planes employed. As each frame must hence be evaluated independently, the set of effects that may be discernable is limited. However, the results will point out the potentials of RIM self-calibration.
The data set being investigated was captured using the amplitude modulating Range Camera SwissRanger™ 3000 (SR-3000) fabricated by Centre Suisse d'Electronique et de Microtechnique SA (CSEM).
Before the presentation of the observed error statistics in section 3, section 2 outlines the required input, consisting of a description of the data set, and the determination of the reference data. The results are discussed in section 4.

2 DATA

2.1 Range Video

In order to facilitate the subsequent interpretation of data deviations, the data set to be examined was captured under conditions that suppress several error sources, according to the observations documented in [4-7]:

- the warm-up period was awaited, after which the self-induced heating of the sensor and the device cooling system have reached an equilibrium and cease to vary distance measurements;
- data were captured in a room with constant temperature, again avoiding variations of the sensor temperature;
- the integration time was kept constant for all images, making them comparable this way;
- objects within the FOV held similar reflection properties;
- data were captured in a darkened room, thus with constant, low background noise;

For the present paper, a sequence of 1000 range images of three cuboids in front of a wall was captured. In the course of the video, the camera wanders around along a random path, showing the cuboids from a constantly changing point of view and camera attitude. Figure 1 shows an exemplary range (left) and the according intensity image (right).

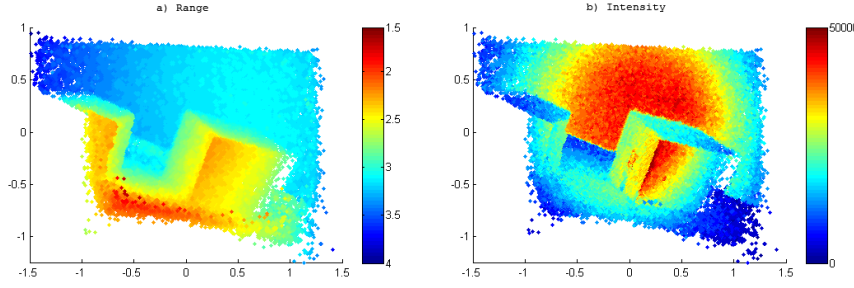

Figure 1: Orthographic projection of the data points from an exemplary range image (left), and the corresponding intensity image (right).

2.2 Extraction of Reference Planes

As a consequence of the huge amount of data, the determination of reference planes for the subsequent derivation of point residuals demands an automated segmentation approach. According to the characteristics of the raw RIM data, several requirements are introduced for this segmentation. In addition to fast performance, these are the capability to handle an extremely high noise level (the measurement noise is much greater than the distance between neighboring points), robustness at sharp surface features (planes intersecting at a common edge), and mode seeking behavior due to the ignorance of the error distribution (e.g. in laser scanning it is often asymmetric).

The segmentation is based on highly robust local regression planes determined for each point. The robustness is achieved by means of the Fast Minimum Covariance Determinant approach [9]. For performance reasons, heuristics have been introduced. Subsequently, seed-clusters are determined by global histogram analyses in the space defined by the local regression plane parameters. A subsequent region growing by means of this parameter space and a distance threshold in object space is performed. The final plane segments are determined by iterative L1-norm plane adjustments. Finally, (disconnected) plane components are separated through respective analyses. The algorithm is a fully three-dimensional approach which is independent from the definition of the coordinate system. A detailed description of the segmentation algorithm can be found in [10].
To increase the reliability of the segmentation, the range images are averaged using a time window comprising the chronologically preceding and succeeding images. This averaging introduces further systematic errors but reduces measurement noise. Naturally, the original images are employed however for the calculation of the residuals used in section 3. In total, about 17 million points are assigned to plane segments, which are 67% of all data. Figure 2 (right) shows the segmentation result of the range image depicted in figure 1. The discriminative power of the algorithm may be deduced from a look at the shading of the raw data (left). Sensor pixels capturing range discontinuities yield averaged, grossly erroneous measurements (‘tail of a comet’). To minimize their influence, points featuring angles of incidence larger than 85 degrees are discarded.

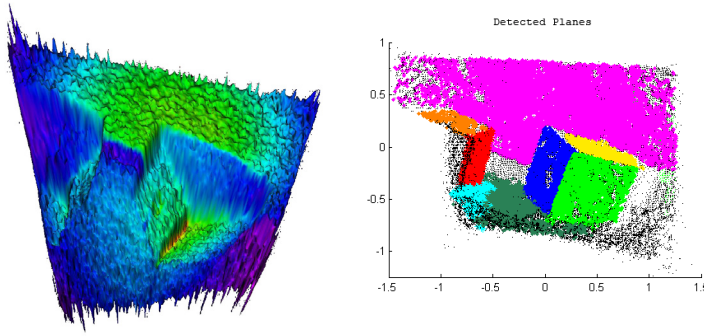

Figure 2: Shaded triangulation of the matrix of the range image shown in figure 1, in a perspective view from above and color coded according to the observed signal intensities (left). Right: Planes detected by means of automatic segmentation of the same range image. Planes are depicted in different colors: the right cuboid in yellow, blue, and light green, the left cuboid in orange and red, the floor in cyan and dark green, and the wall behind in magenta.

3 RESIDUAL STATISTICS

The residuals evaluated in this section are defined as the normal distances between each classified point of the range video and the adjusting plane of the respective plane segment. They are examined with respect to those RIM error sources stated in [4-7] that have not been suppressed during data capture (cf. section 2.1). Primarily, this is done by histogram analyses in the one- or two-dimensional domains of the respective error sources.

3.1 Position in the Field of View

Figure 3 (left) shows the number of classified points i.e. evaluable residuals, accumulated for every image pixel. Obviously, the numbers of residuals are not uniformly distributed, but the floor and wall of the captured room are apparent, which are imaged in these sensor areas throughout large portions of the range video. Considering this, the standard deviation per pixel given in figure 3 (center) may be interpreted more easily: in general, precision increases towards the image center.

However, the cuboids captured primarily in the central sensor area affect precision negatively due to an elevated rate of misclassifications along the segment borders. Figure 3 (right) furthermore depicts the mean residuals per pixel and a rather unexpected image: the upper central area of the sensor indicates a large cluster of positive mean residuals, while the areas on the left and right sides present the contrary. Again this is the sensor area where primarily the wall was captured. [4,6] compute a distance offset depending on the position in the FOV, which radially decreases with the distance from the image center and is caused by the accordingly decreasing illumination of object space by the SR-3000 (cf. Figure 1, right). Most likely, the oscillation of mean residuals in the upper image area is the remainder of this effect having adjusted planes across large fractions of the image width.

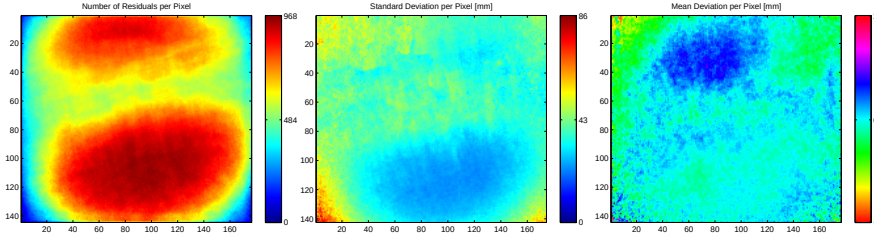

Figure 3: For each pixel of the range sensor, the number of classified points is given (left). The image therefore has the size of the RIM sensor: 144px by 176px. Furthermore, the standard deviations (center), and the mean residuals per pixel (right) are shown.

3.2 Frame Number

The frame-wise evaluation of the residuals and hence their temporal evolution is not expected to point out the drifting effects due to self-induced heating reported in [4-7], as the warm-up period had been awaited before the capture of data, and anyway the reference planes are adjusted frame-wise. However, correlations between the standard deviations, mean distances, intensities and angles of incidence become apparent. Figure 4 shows the plot of the frame-wise averaged parameters along the time line of the range video. The corresponding correlation coefficients are presented in Table 1. They are all positive and highly significant. As expected, the highest correlation results for the standard deviation and the mean distance.

The remaining coefficients are only plausible when considering the test data. The data errors in lateral direction originating from a deficient projection model (camera objective) and the pixel misalignments in the sensor matrix have been proven to be very low compared to the errors in distance of raw SR-3000 data [3]. While the prevailing errors in distance wholly affect point residuals on planes featuring angles of incidence near zero, their impact decreases with this angle. The however positive correlation between the angle of incidence and the standard deviation is explained by one of the dominating planes: the mean angle of incidence mostly reflects the portion of the sensor covered by the wall, which is captured at low angles of incidence and yields small residuals. By contrast, the inferiorly classified cuboid planes were imaged mainly at large angles.

Likewise, the positive correlation between the mean angle of incidence and the mean intensity must be explicated considering the peculiarities of the data set: the floor and the cuboid planes were captured at large angles. While the floor was imaged at short and intermediate distances, the cuboid planes appeared from nearby throughout the video, at accordingly highest intensities.

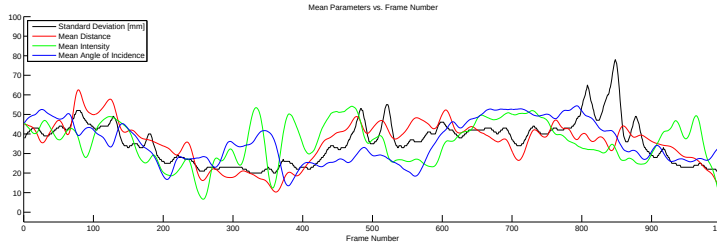

Figure 4: Mean values for each frame are given along the time line of the range video: standard deviation, object distance (scaled), intensity (scaled), and angle of incidence (scaled). All graphs are smoothed with a moving average.

R	Standard Deviation	Distance	Intensity	Angle of Incidence
Standard Deviation	1.00	0.57	0.15	0.42
Distance	0.57	1.00	0.29	0.24
Intensity	0.15	0.29	1.00	0.32
Angle of Incidence	0.42	0.24	0.32	1.00

Table 1: Correlation coefficients among the four parameters plotted in figure 4.

3.3 Distance

Figure 5 depicts the number of residuals (left) and their mean and standard deviation (right) with respect to the distance between the respective data points and the focal point. The left histogram features a dominant maximum around 2.5 meters, while distances below one and above four meters are underrepresented. The significance of the values given on the right must be taken into account accordingly. Nevertheless, the systematic effect that precision decreases with distance can certainly be stated.

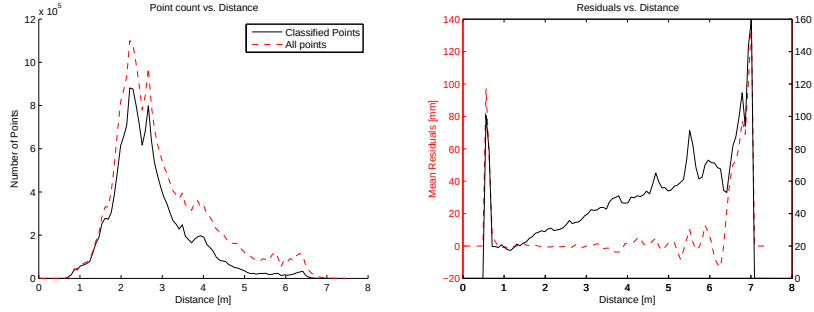

Figure 5: The number of classified points (left, black), the number of all data points (left, red), and mean residuals (right, red) and standard deviations (right, black) for the classified points with respect to the observed distance.

3.4 Intensity

Analogical to figure 5, figure 6 depicts the statistics with respect to the observed signal intensity, which is distributed more evenly. The tenor that precision improves with intensity is noticed up to intensities of about 45k. Above 55k, saturation effects obviously boost the standard deviation again.

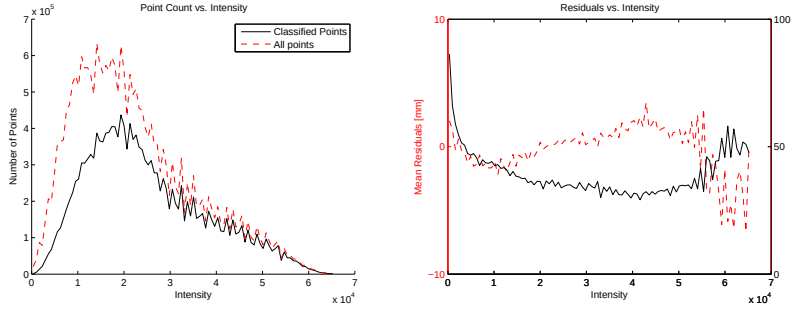

Figure 6: The number of classified points (left, black), the number of all data points (left, red), and the mean residuals (right, red) and standard deviations (right, black) for the classified points with respect to the observed signal intensity.

3.5 Angle of Incidence

Finally, figure 7 shows the statistics regarding the angle of incidence. As remarked in section 2.2, no angles above 85 degrees are evaluated. The histogram of point counts features two modes: the smaller one originating from the wall, the other from the floor. According to the standard deviation given on the right, an impact of the angle of incidence on measurement precision can be inferred at least for angles above 70 degrees, above which the better lateral measurement precision becomes prevailing (cf. section 3.2).

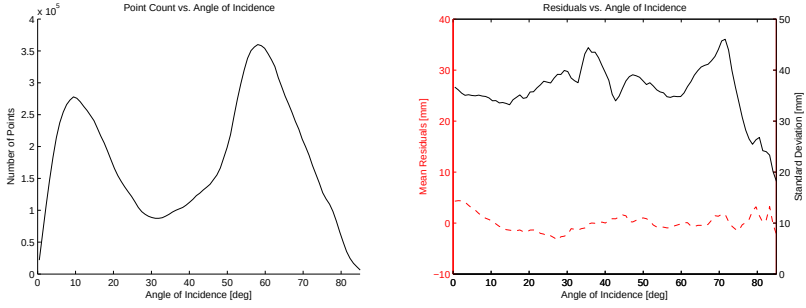

Figure 7: The number of classified points (left, black), mean residuals (right, red) and standard deviations (right, black) with respect to the angle of incidence.

4 DISCUSSION

Despite the high correlation between the sources of systematic errors and the exceptionally high level of measurement noise in the set of raw in situ RIM data being examined, multiple systematic effects can be observed by means of statistical analyses. Considering only the sensor areas where mainly the large planes of the room's floor and wall were captured, the deterioration in precision towards the image borders and the decrease of the distance offset are apparent. Apart from the mean distances, the frame-wise evaluation of mean residuals only reflects the errors being of interest when strongly regarding the captured objects, and may otherwise lead to misinterpretations. The studies of residuals with respect to polar distances and signal intensities result in definite conclusions, among which most interesting are the strong relation between observed distance and measurement precision, and the deterioration of measurement precision yet above 55k intensity. The examination regarding the angle of incidence does not permit overall inferences, but an improvement of precision for large angles can be stated.

5 CONCLUSIONS AND OUTLOOK

The segmentation algorithm applied to reconstruct the captured planes in object space and gain the needed point residuals performs well, having averaged consecutive range images. The consideration of the large difference between distance and angle measurement precisions, and furthermore of the large effects of the true distance, the target reflectivity and the angle of incidence on the observed distance would certainly enhance the results further. Detecting erroneous pixels at discontinuities on the object ('tail of a comet') by means of plane segmentation and checking of the angle of incidence seems to be a viable approach.

The evaluation of residuals of raw RIM data points assigned to adjusting plane segments allows for the deduction of the general error performance of the distance measurement system with respect to the position in the field of view, the observed distance, the signal intensity and partly of the angle of incidence. However, the investigations also highlight the risks of statistical induction, presenting highly significant correlations between variables that must be reasoned using exterior information. The tracking of planes across multiple frames and their separated evaluation is expected to further improve the presented approach. As soon as the impacts of target reflectivity on distance measurements are known, the error factors can be separated more clearly. Even the modeling of non-Lambertian reflectors employing intensities and distances captured at diverse angles of incidence seems to be in the realms of possibility.

REFERENCES

1. Büttgen, B., Oggier, T., Lehmann, M.: CCD/CMOS Lock-in pixel for range imaging: challenges, limitations and state-of-the-art. In proceedings: 1st range imaging research day, pp.21-32, Ingensand / Kahlmann (eds.), Zurich, 2005.
2. Oggier, T., Büttgen, B., Lustenberger, F.: SwissRanger SR3000 and first experiences based on miniaturized 3d-TOF cameras. In proceedings: 1st range imaging research day, pp.97-108, Ingensand / Kahlmann (eds.), Zurich, 2005.
3. Westfeld, P.: Ansätze zur Kalibrierung des Range-Imaging-Sensors SR-3000 unter simultaner Verwendung von Intensitäts- und Entfernungsbildern. Photogrammetrie - Laserscanning - Optische 3D-Messtechnik. Beiträge der Oldenburger 3D-Tage 2007, Luhmann / Müller (eds.), Wichmann, 2007
4. Kahlmann, T,. Remondino, F., and Ingensand, H.: Calibration for increased accuracy of the range imaging camera Swissranger™". International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVI, part 5, pp. 136-141, 2006.
5. Kahlmann, T., Ingensand, H.: Investigation of the 3d-range-imaging-camera Swissranger™. In proceedings Optical 3-D Measurement Techniques VII/I, pp. 227-236, Grün and Kahmen (eds.), Vienna, 2005.
6. Kahlmann, T., Remondino, F., and Guillaume, S.: Range imaging technology: New developments and applications for people identification and tracking. SPIE proceedings, Vol. 6491, 2007.
7. Kahlmann, T., Ingensand, H.: Range imaging sensor properties and calibration. In proceedings: 1st range imaging research day, pp.71-80, Ingensand / Kahlmann (eds.), Zurich, 2005.
8. Steitz, A., Pannekamp, J.: Systematic investigation of properties of PMD-Sensors. In proceedings: 1st range imaging research day, pp. 59-69, Ingensand / Kahlmann (eds.), Zurich, 2005.
9. Rousseeuw, P. J., van Driessen, K.: A fast algorithm for the minimum covariance determinant estimator. Technometrics, Vol. 41, No. 3, pp. 212-223, 1999.
10. Dorninger, P, Nothegger, C.: 3d segmentation of unstructured point clouds for building modeling. In proceedings: Photogrammetric image analysis, Munich, to appear, 2007