

QUANTIFYING THE DISTORTION OF DISTANCE OBSERVATIONS CAUSED BY SCATTERING IN TIME-OF-FLIGHT RANGE CAMERAS

W. Karel^{a,*}, S. Ghuffar^b, N. Pfeifer^b

^aChristian Doppler Laboratory “Spatial Data from Laserscanning and Remote Sensing“ at the

^bInstitute of Photogrammetry and Remote Sensing, Vienna University of Technology,
Gusshausstraße 27-29, 1040 Vienna, Austria – {wk,sg,np}@ipf.tuwien.ac.at

KEY WORDS: Range Imaging, Range Camera, Photonic Mixer Device, Systematic Error, Scattering, Internal Reflection

ABSTRACT:

Time-of-flight range cameras simultaneously gather object distances for all pixels of a focal plane array by evaluating the round-trip time of an emitted signal. In contrast to competing techniques, cameras combining continuously emitted, amplitude modulated signals and Photonic Mixer Devices (PMD, lock-in pixels) to derive signal phase shifts and hence object distances, have reached mass production and are available at low costs. While ranging precisions typically amount to some centimetres, accuracies may be worse by an order of magnitude. Systematic distortion factors of the ranging system can be grouped into local and non-local errors. While local distortions affect the pixels individually, non-local ones contaminate larger areas of the sensor. ‘Scattering’ denotes one of these non-local errors, meaning the spreading of portions of the incident light over the sensor due to multiple reflections between the sensor, lens, and optical filter. The present contribution analyses this phenomenon with respect to various capture parameters, with the objective of a better understanding and a validation of assumptions.

1. INTRODUCTION

Time-of-flight (ToF) range cameras simultaneously gather object distances for all pixels of a focal plane array by evaluating the round-trip time of an emitted signal. In contrast to competing techniques (Leonardi et al., 2009; Niclass et al., 2008), cameras combining continuously emitted, sinusoidally amplitude modulated signals (AM-CW) and Photonic Mixer Devices (PMD, lock-in pixels) to derive signal phase shifts and hence object distances (Lange et al., 1999), have reached mass production and are available at low costs. PMD cameras provide up to 25 frames per second, sensor array sizes exceeding 176×144 pixel², measurement ranges of up to tens of metres, and deliver signal amplitude data in addition to range observations. While ranging precisions typically amount to some centimetres, accuracies may be worse by an order of magnitude.

PMD cameras combine the advantages of well-established 3D measurement techniques like image triangulation and Laserscanning, meaning the simultaneous capture of data on a solid array and the direct range determination using the time-of-flight, and are already used in applications with rather low demands on data quality. However, distance observations have been reported to be affected systematically by several local distortion factors, including the object distance itself (non-linearly), the signal amplitude, the integration time, and the position on the sensor. Being observed or known quantities, correction models have been developed that express the distortions explicitly (Karel and Pfeifer, 2009; Lichti and Rouzaud, 2009; Lindner and Kolb, 2007). In addition to these local, pixel-wise influences, two effects have been identified that affect range observations in a possibly large neighbourhood of sensor elements.

First, emitted light may be reflected multiply in object space (‘multipath’) and thus may superimpose and distort directly reflected parts of the signal (Guðmundsson et al., 2007). This

effect may only be present if surfaces are arranged appropriately in object space, e.g. when observing the corner of a room.

As a second non-local effect, the echo of the optical signal emitted by the illumination unit is scattered to some extent over the sensor due to multiple reflections within the camera i.e. between the lens, the optical filter, and the sensor (‘scattering’; also called ‘lens flare’ in conventional photography). As a result, the incident light observed by each pixel is a mixture of the light returned from the geometrically corresponding pixel footprint on the object (‘focused light’), and the parasitic signal reflected at other pixels and thus corresponding to other parts of the object (‘scattered light’), see fig. 1. While the impact on observed signal amplitudes may be negligible, phase angle measurements and hence derived object distances may be affected largely in images with high amplitude and depth contrast, which is favoured by active illumination.




Figure 1: Illustration of the scattering phenomenon. Right: 3 targets at different distances from the camera produce echoes with different phase angles and amplitudes. Portions of these echoes are reflected back to the lens, and back again to different locations on the sensor (shown for target 1). The scattered light superimposes the focused light from the other targets, which corresponds to an addition in the complex plane (left), when assuming a strictly sinusoidal signal.

* Corresponding author.

1.1 Related Work

Mure-Dubois and Hügli (2007) assume a point spread function (PSF) that is constant over the image plane, why scattering may be expressed as a 2-dimensional convolution operation with a constant kernel. By visual inspection of its efficiency, they estimate the optimal inverse filter, which is then convolved with the observed image in order to compensate for scattering. The inverse filter consists of 2-dimensional Gaussian functions, which are separated into 1-dimensional kernels in order to reduce computational complexity. However, the authors conclude that the assumption of spatial invariance of scattering may not hold.

Kavli et al. (2008) empirically derive local PSFs for various positions on the sensor. This estimation is performed using a planar, dark background in front of which the camera is mounted such that the optical axis results to be normal to the plane (normal case). On this background plane, a bright, circular target is placed at various positions, having a size such that it approximates an unresolved scattering point source. By subtracting images with the target being present from another one without (background subtraction), and rescaling to unit size, the empirical PSFs are obtained. As the target lies in the background plane, the difference in phase is zero i.e. the PSFs are real-valued. The PSFs result to be asymmetric and are modelled non-parametrically. In order to avoid the difficult deconvolution with a spatially variant, non-parametric model, they apply an iterative image restoration algorithm to compensate for scattering, which allows the PSFs to be applied in a forward-mode. Based on the observation that high-amplitude image regions affect lower-amplitude regions more than vice versa, the scatter from the brightest image regions is estimated and subtracted, using the PSF for the nearest image position. This procedure is repeated for the next brightest regions, until the scattering for the whole image has been compensated. Applied to real scenes, the approach proves to efficiently compensate for scattering distortions, even though the compensation notably overshoots in certain configurations.

In contrast to the aforementioned approaches, the present contribution aims at investigating the nature of the scattering phenomenon, with the fewest assumptions possible, and without the immediate goal of modelling or compensating. For this purpose, various capture parameters are varied, and their impact on scattering is studied.

2. EXPERIMENTAL SETUP

In order to observe scattering phenomena, images of the background without the foreground are subtracted from images in which foreground is present while keeping the camera's orientation constant. The experimental setup consists of a planar, black, diffusely reflecting paper serving as background and planar, white, circular targets serving as foreground. The targets feature radii of 20, 30, and 40mm and are made of 2mm thick cardboard. They are mounted on a tripod through a long cylindrical stick of about 5mm diameter whose surface is covered with black tape to minimize its effect on scattering. All experiments are conducted using a Swissranger SR-3000, manufactured by MESA Imaging AG.

2.1 Temporal Variation

In order to minimize noise, hundreds of frames of the same scene are averaged over time. To find the optimal number of frames to be averaged, continuous images are acquired for

several minutes. To achieve accurate background subtracted images it is crucial to use the appropriate number of frames. The plot of the mean values of amplitude and distance of all pixels in the image against the number of frames shows a response of the camera after a change in integration time. The analyses of these plots reveal the significance of using the optimal number of frames and start-up time, which cannot be neglected in order to obtain accurate camera measurements.

Therefore an experiment was performed to determine time response of the camera. In this experiment, frames were continuously captured for approximately 30 minutes keeping the imaged space constant. The integration time was changed during run time from 10 to 100. Figure 2 shows the response of the camera after every change in the integration time. The analysis of this data shows that the camera exhibits two types of temporal variations. First is the initial or start up transience during which there is significant amount of variations in the mean distance and amplitude measurements. The second temporal variation is of shorter time period spreading over the whole sequence of frames.




Figure 2: Response of the camera to changes in integration time: frame-wise mean distance (blue), and amplitude (red). The integration time is indicated in green.

The initial transient response of the camera depends on the amount of change in the 8 bit integration time value which is measured to be approximately 2 minutes for a 20 units' step, after this time the mean distance values become stable and show a periodic variation of a few millimetres. During the initial transience the mean values of the distance image change by as much as 4 cm. For a step size of 90 units the change in the mean distance is about 6 cm.

Figure 3 shows the short term variations in the camera readings. The distance measurements show a mean value of 1.689 m and a standard deviation of 2.3mm. The curve fitting of the data was done to determine the time period of these temporal variations. A fitted sine wave shows similar residuals for all the experimental data with different integration times and foreground. Hence it is imperative to use a number of frames which corresponds to a time that encompasses integer multiples of this sine wave, in order to produce accurate background subtracted images. The comparison of this temporal variation with the internal camera temperature shows a direct correspondence between the camera temperature and the observations. Hence we can infer that these variations are caused by changes in the internal temperature of the camera.

2.2 Considerations on Setup

The experiments are performed in a sufficiently large room in order to avoid any multipath effect. Severe distortions in the range and amplitude images have been observed due to objects placed just outside the field of view (FOV) of the camera. Therefore, the FOV of the camera is restricted within ample distance from the boundaries of the background to minimize any effect of objects just outside the FOV of the camera.

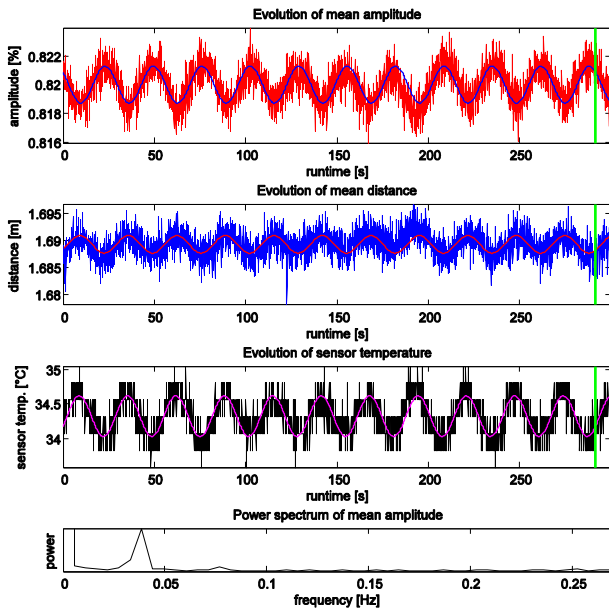


Figure 3: Short term periodic variations of the frame-wise mean of amplitude (red) and distance observations (blue), and the sensor temperature (black), which show a strong relation. The image sequence is truncated (green) at a multiple of the period, in order to compute unbiased means of observations. The power spectrum of the amplitude signal (bottom) indicates only 1 dominant frequency.

While designing the experiments, placing the target in the near-field of the illumination unit was avoided to minimize the half-shadowing of the background area neighbouring the target as shown in figure 4. The closest distance between the target and the camera during the experiments was about 75cm. Placing the target too close to the camera would cause image blur because of the fixed focus. During the experiments some horizontal line artefacts were observed in the background subtracted images, as shown in figure 5, whose magnitudes increase with integration time. Therefore, the integration time was adapted to a lower value to minimize this effect. The cause of these artefacts has not been investigated and corresponding image regions have been masked where present and disregarded during evaluation. The histograms of the rows with these artefacts, as shown in figure 6, indicate that these artefacts do not originate from outliers in the data.

2.3 Experiments

To observe the effects of different capture parameters, following experiments are performed. In all experiments, the orientation of the camera with respect to the background is kept constant, at a normal distance of 1.46m. Unless otherwise stated, the scene is captured with the integration time set to 30 units, and the target serving as foreground has a radius of 20mm, is positioned at a distance of 1.15m from the camera, and is centred at the optical axis, why it is imaged at the principal point. The camera's interior orientation is taken from Karel (2008). Work is performed in the dark, at room temperature.

The aim of the first experiment is to analyze the effect of *integration time* on scattering. The integration time is set to 30, 60, and 90 units.

The second experiment aims to analyze the effect of *target size* on scattering. Image sequences are acquired for three different targets of 20, 30 and 40mm radius.




Figure 4: Scattering of amplitudes (left) and distances (right), computed by subtracting the amplitudes and distances of fore- and background images separately. The region (here 1.3px wide) surrounding the target's image may (also) be affected by a half shadow on the background caused by the target and the extended, two-dimensional light source.




Figure 5: The same scene and plots as in fig. 4, but captured with a longer integration time. Note the distorted rows covered by the target and the colour mappings that are different from the ones in fig. 4.

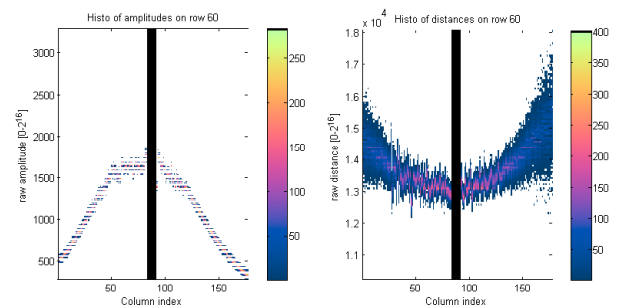


Figure 6: Histogram of raw observations of the foreground sequence used for fig. 5, for a row which features artefacts. Left: amplitude. Right: distance, which shows that the observations on each column do not feature outliers, but are systematically distorted. Columns covered by the target are masked black.

In the third experiment, both the target size and the ratio of the distance between the camera and the target to the distance between the camera and the background are changed. Targets of 40, 30 and 20mm are placed at distance ratios of 4/4, 3/4 and 2/4, which results in the targets being imaged with the same size in all three cases. Hence the effect of the *target distance* is studied independently from the target's size in the image.

The aim of the fourth experiment is to study the influence on scattering of the *target's position in the image*. Again, the target with 20mm radius is placed at a distance of 1.15m from the camera, and the centre of the target is aligned with the optical axis of the camera. Afterwards, the target is placed at 8 different positions on a circle perpendicular to the optical axis going through the last position, with an angular difference of 45° , see fig. 7. The results of this experiment help to understand the symmetry of the scattering phenomenon.




Figure 7: Experiment to observe the effect on scattering of the target's position in the image. The central point corresponds with the principal point. As all target positions lie in a plane perpendicular to the optical axis, all surrounding points are located at the same distance.

3. EVALUATION

As a preliminary step in evaluation, pixel-wise mean observations are computed for all image sequences, based on the data of the respective sequence being truncated at the maximum multiple of the period of short-term variations (see green lines in fig. 3).

The experiments described in subsection 2.3 assume a nominal positions of the targets in the image plane, which are realized with some imprecision. To account for these deviations, target positions are detected in the mean amplitude images using the method presented by Otepka (2004), and are considered via bilinear resampling where appropriate. Furthermore, the areas covered by the targets, together with those covered by the stick that the targets are mounted on, are masked and these areas are disregarded in evaluations.

AM-CW ToF cameras observe the phase angle between the emitted and returned signals, which is linearly related to the modulation wave length and the object distance. For the sake of expressiveness, object distances are given in the following instead of phase angles.

As mentioned in sec. 2, background subtraction is used to isolate the effect of scattering. Two variants of subtraction are used: (1) separate subtraction of amplitude and distance observations, which results in the actual distortion of observations, and (2) assuming a sinusoidal signal, complex subtraction of the signals (conf. fig. 1), which yields the scattered light i.e. the distortion signal.

3.1 Variation of the Integration Time

Changes of the integration time do not affect the optical signal received by the sensor. However, the amplitude observations reconstructed from the signal are linear in integration time, as may be derived from e.g. Lange et al. (1999). Therefore, changes of the integration time affect the background subtracted amplitude images, while corresponding distance images are not (see fig. 8).

Division of the background subtracted amplitudes by the integration time effectively eliminates its influence (as can be seen in figure 9), which demonstrates the linear relation mentioned above.




Figure 8: Differences of separately background subtracted images captured with integration times of 90 and 30 units, respectively. While the difference of amplitude scattering shows a notable influence (left), the corresponding image for the distance does not (right). Areas covered by the mounting stick and distorted rows (conf. fig. 5) are masked in magenta. Note that these images are the result of subtracting background subtracted images, why 4 observation variances add up to a significant amount of noise, especially towards the image corners, where the intensity of incident light and hence the signal amplitude decreases due to vignetting and illumination fall-off.

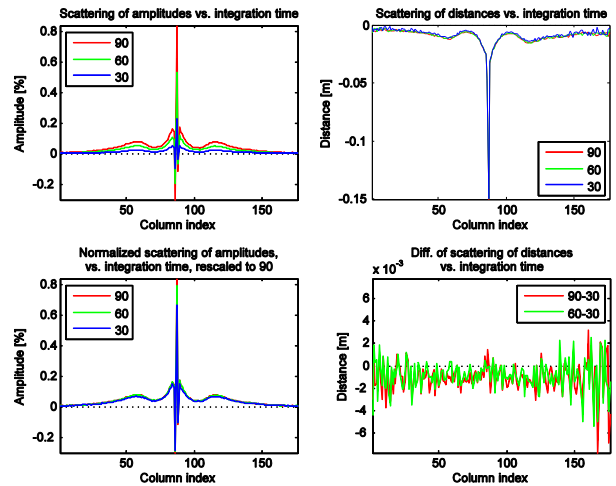


Figure 9: Top: the arithmetic mean of the separately subtracted observations on the 2 rows just above and below the area covered by the target, for the integration times of 30, 60, and 90. Left: amplitudes. Right: distance. While the integration time notably affects scattering of observed amplitudes, distances proof to be unaffected. Bottom left: mean scatter of amplitudes, with the influence of the integration time on amplitude observations eliminated, and rescaled to the integration time of 90 for better comparability. Bottom right: Differences of pairs of graphs plotted above: 90-30, and 60-30, considered as being random.

3.2 Variation of the Target Size

Increasing the target size while keeping the other capture parameters constant leads to an increase of the magnitude of the scattering halo: the maximum of the side lobes for distances increases from about 2cm for a radius of 20mm to about 8cm for a radius of 40mm. See fig. 10, which shows the difference of fore- and background images, with amplitude and distance images subtracted separately. For fig. 11, complex background subtraction has been applied. The scattering signal for distances is constant for all columns and unaffected by the target size. This is understood as a proof that the amplitude modulation of

the optical signal closely forms a sine wave. Unaffected by the target size, the distance between the target centre and the maxima of the amplitude scatters' side lobes stays constant. Note, however, that all target sizes fit into the scattering halo. Finally, the ratio of the maxima at the side lobes to the squares of the target radii is practically constant for all target sizes i.e. proportional to the target area. This conforms to the model of scattering being linear with the signal amplitude, as used e.g. by Kavli et al. (2008).




Figure 10: Separate background subtraction: scattering of amplitudes (left) and distances (right) for target sizes of 20mm (top), and 40mm (bottom).

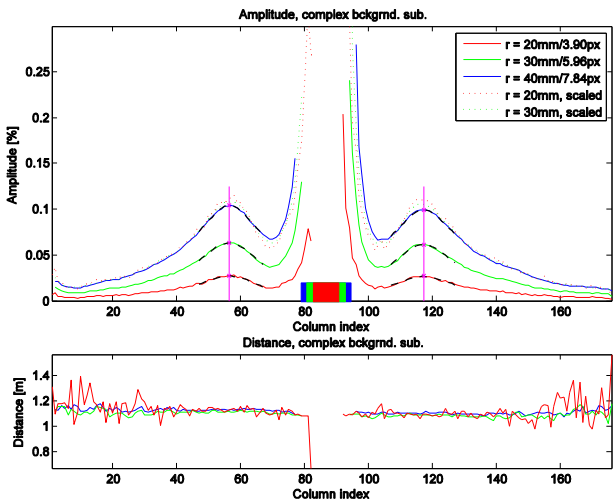


Figure 11: Mean scattering along the 5 rows next to the target centre, same data as for figure 10, but complex background subtraction: amplitudes (top), distances (bottom). The distance signal is constant for all columns and all target sizes, reflecting the narrow range of object distances to the target. For all target sizes, the maxima of the amplitude signals' side lobes (magenta crosses), determined via locally adjusting hyperbolae (black, dashed) are at the same distance from the target centre. Scaling the amplitude signals (solid) to the radius of 40mm (dotted) reveals that they flatten with increasing radius. The target image diameters are presented by the width of the rectangles, filled with the respective colour.

3.3 Variation of the distance to the target, keeping the target image size constant

Fig. 12 (top) shows the separately subtracted images of the 40mm-target. As it lies in the background plane, the introduction of the target does not affect the distance observations on the background. Fig. 12 (bottom) shows the 20mm-target at half distance between camera and the background with both distance and amplitude distortions. Complex background subtraction reveals more interesting information (see fig. 13): the scatter of distances is constant for all columns, and the mean value for each combination of object distance / target size reflects well the nominal distance. The distance from the target centre to the maxima of the side lobes, as seen in the plot of amplitude scatter in fig. 13, is constant. This indicates that the observed phenomenon is truly an internal effect, which is further substantiated by the plots of scaled amplitude scatters, which overlap closely.




Figure 12: Separate background subtraction of amplitudes (left) and distances (right) for a target with a radius of 40mm lying in the background plane (top, no distortion of distances present), and for a target with a radius of 20mm, located at half the distance, which is thus imaged with the same size (bottom).




Figure 13: Complex background subtraction, same data as for fig. 12. The distances d are specified as ratios of the distance to the target and the distance to the background.

3.4 Position with Respect to the Principal Point

The plots of the amplitudes resulting from complex background subtraction of images with the target at different positions in the FoV (see fig. 14) show that scattering is obviously not invariant w.r.t image space. However, scattering has at least mirror symmetry about the principal point. For each pair of images lying opposite to each other w.r.t the principal point, fig. 15 shows the difference, with one of the images being mirrored horizontally and/or vertically before subtraction. These differences are smaller by one magnitude, and seem random.




Figure 14: Background subtraction in the complex domain, resulting amplitudes for various positions of the foreground object, which are reflected by the position of the image within the figure: once for the target at the principal point (central image). For the surrounding images, the target centres lie on a circle around the principal point, at equal angles from each other, in steps of 45° (see fig. 7). Masked pixels are coloured magenta.

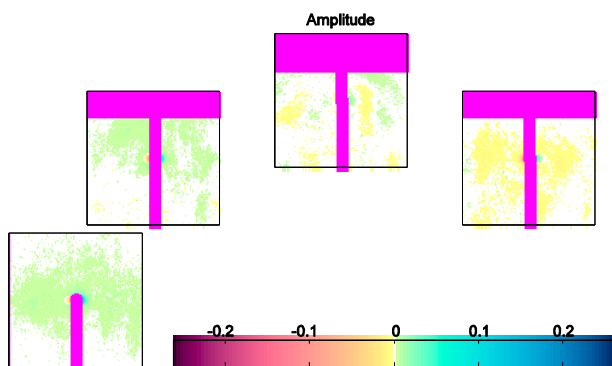


Figure 15: Difference of amplitude scatters shown in fig. 14: images lying opposite to each other w.r.t the principal point have been mirrored horizontally and / or vertically, and subtracted.

4. CONCLUSIONS

This contribution presents methods to gather precise scattering data. The results indicate that scattering is an additive, linear

phenomenon. The influence of integration time can be eliminated completely. Modelling the modulation of the optical signal as a strict sine wave seems to be a good approximation, as the phase angles / distances of the scattering signal result as constant all over the image plane, corresponding to the (mean) distance to the foreground object. However, for a proper modelling and compensation of scattering, further studies are necessary.

REFERENCES

- Guðmundsson, S., Aanas, H. and Larsen, R., 2007: *Environmental effects on measurement uncertainties of time-of-flight cameras*. In: *proc. International Symposium on Signals, Circuits and Systems*, pp.1-4.
- Karel, W., 2008. *Integrated range camera calibration using image sequences from hand-held operation*. In: *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Beijing, China, Vol. XXXVII, Part B5*, pp. 945-952.
- Karel, W., Pfeifer, N., 2009. *Range camera calibration based on image sequences and dense comprehensive error statistics*. In: *Three Dimensional Imaging Metrology, San José, USA, proc. SPIE Vol. 7239-12*.
- Kavli, T., Kirkhus, T., Thielemann, J., Jagielski, B. 2008. *Modelling and compensating measurement errors caused by scattering in time-of-flight cameras*. In: *Two- and Three-Dimensional Methods for Inspection and Metrology VI, San Diego, USA, proc. SPIE Vol. 7066-4*.
- Lange, R., Seitz, P., Biber, A., Schwarte, R., 1999. *Time-of-flight range imaging with a custom solid state image sensor*. In: *Laser Metrology and Inspection, proc. SPIE Vol. 3823*, pp. 180-191.
- Leonardi, F., Covi, D., Petri, D. and Stoppa, D., 2009. *Accuracy performance of a time-of-flight CMOS range image sensor system*. In: *IEEE Transactions on Instrumentation and Measurement*, 58(5), 1563–1570.
- Lichti, D., Rouzaud, D., 2009. *Surface-dependent 3D range camera self-calibration*. In: *Three Dimensional Imaging Metrology, San José, USA, proc. SPIE Vol. 7239*.
- Lindner, M., Kolb, A., 2007: *Calibration of the intensity-related distance error of the PMD TOF-camera*. In: *Intelligent Robots and Computer Vision XXV, proc. SPIE Vol. 6764/1*.
- Mure-Dubois, J., Hügli, H., 2007: *Real-time scattering compensation for time-of-flight camera*. In: *proc. ICVS Workshop on Camera Calibration Methods for Computer Vision Systems, Applied Computer Science Group, Bielefeld University, Germany*
- Niclass, C., Favi, C., Kluter, T., Gersbach, M. and Charbon, E., 2008. *A 128x128 single-photon imager with on-chip column-level 10b time-to-digital converter array capable of 97ps resolution*. In: *proc. IEEE International Solid-State Circuits Conference*, pp. 44–46.
- Otepka, J., 2004: *Precision target mensuration in vision metrology. Dissertation at the Institute of Photogrammetry and Remote Sensing, Vienna University of Technology, Austria*.