Performance Assessment and Calibration of the Kinect 2.0 Time-of-Flight Range Camera for Use in Motion Capture Applications

Jeremy STEWARD, Dr. Derek LICHTI, Dr. Jacky CHOW, Dr. Reed FERBER, and Sean OSIS, Canada

Key words: Low cost technology; Photogrammetry; Range camera; Motion capture

SUMMARY

Robust, three-dimensional (3D) geometric information is a powerful analytical tool, and it is of interest to determine the size, shape, and geometric properties of objects in the real world when performing pre-mission surveys, deformation analyses, or motion capture analyses. Traditional photogrammetric reconstruction techniques require multiple sensors and retro-reflective markers or targets to be placed on objects of interest during data acquisition. The Microsoft Kinect 2.0 sensor provides an onboard time-of-flight (ToF) ranging sensor based on the Canesta technology. Given the price for the Kinect 2.0 is $200 USD, it shows potential to become a cost-effective, single-sensor solution for capturing full 3D geometric information in place of costly, multi-sensor techniques requiring invasive or otherwise difficult to place markers. This study examines the performance characteristics and calibration of the Kinect 2.0 sensor in order to determine the feasibility of its use in 3D imaging applications; particularly that of human motion capture.

The Kinect 2.0 sensor was tested under controlled conditions in order to determine the warm-up time, distance measurement precision, target reflectivity dependencies, residual systematic errors, and the quality of human body reconstruction when compared to a device of known quality. The sensor in question proved promising, showing similar precision to other ToF imaging systems at a mere fraction of the price. Over the course of this testing, it was found that negligible warm-up time is required before the geometric measurement performance stabilizes. Furthermore, a distance measurement precision of approximately 1.5mm is achievable when imaging highly reflective, diffuse target surfaces. Beyond the performance characteristics of the sensor itself, a self-calibration of the sensor for un-modelled lens distortions improved image measurement residuals by an average of 88%, and likewise improved the range measurement precision by 81%.

Despite these results, factors beyond the user’s control such as scene-dependent distortions, and inhomogeneity in depth accuracy across the image plane limit the potential performance of the sensor. Thus, the following “best practice” guidelines were put forth: 1) Only the inner 300x300 pixels about the centre of the sensor should be used, due to loss in signal strength near the periphery of the image; 2) ensure that the object of interest is within the foreground of the scene, ideally at a range approximately 1-2.5m away from the sensor; and 3) highly-reflective, diffuse objects should be preferred to darker or shiny objects in the captured scene.

Performance Assessment and Calibration of the Kinect 2.0 Time-of-Flight Range Camera for Use in Motion Capture Applications

Jeremy STEWARD, Dr. Derek LICHTI, Dr. Jacky CHOW, Dr. Reed FERBER, and Sean OSIS, Canada

1. INTRODUCTION

The acquisition of three-dimensional (3D) geometric information is paramount to many types of practical surveys that exist today. Subsequently, systems for the capture and analysis of such data prove to be important tools in a variety of fields, from airborne surveys to close-range human motion capture. With specific regard to motion capture applications, time-of-flight (ToF) range cameras offer several advantages in comparison to alternative sensors, such as traditional photogrammetric and terrestrial laser scanning (TLS) systems. ToF cameras, which operate on the principles described by Lange & Seitz (2001) provide a single, active sensor solution which can directly measure the 3D coordinates of a scene in a single shot, without requiring a scanning mechanism. In contrast, traditional stereo-photogrammetric techniques require multiple sensors or sensor locations and often require retro-reflective targets or markers, in order to be able to solve the correspondence problem and register images in a common system. Alternatively, such systems can rely on pattern projection or natural texture; however, in the context of human motion applications, these pose a problem since skin does not have a clearly defined texture between persons, and projecting patterns over moving humans prove difficult, given that their visibility depends highly on the contrast between the patter provided and clothing / skin. Moreover, Fraser (1984) suggests that convergent imagery is necessary when performing highly accurate non-topographical surveys, which may prove to be a constraining factor in many clinical or home-based applications where robust 3D geometric information is needed. TLS based systems are similar to range cameras in that they are active imaging systems that do not require retro-reflective targets; however, TLS systems are unsuitable in applications where large amounts of change or motion are observed, given that subjects of interest in the scene may move considerably between scan passes.

One such sensor that employs ToF camera technology is the Kinect 2.0 sensor, which is produced by Microsoft, primarily for use in home gaming applications. The sensor itself is based on the Canesta technology. The Kinect 2.0 itself costs approximately $200 USD, and features a combination of several sensors: a red-green-blue (RGB) digital camera; a microphone; and 3D ToF range camera. For the purposes of this paper, we solely focus on the 3D ToF sensor and its use in the context of human motion capture applications and close-range surveying. The ToF camera sensor is large (512 x 424 pixels) compared to other ToF range cameras currently on market. Moreover, the Kinect 2.0 is much more cost-effective compared to not only other competing range camera devices, but likewise is generally an order of magnitude or two less expensive than TLS systems. The benefits of full-frame, 3D capture from a single sensor, as well as both the cost-effectiveness and compactness of the Kinect 2.0 sensor therefore make it extremely desirable to use in close-range surveys and motion capture applications.

Unfortunately, the sensor was not developed for surveying and motion capture applications, and as such, the Kinect 2.0 technical documentation does not provide a sensor model which describes the physical characteristics of the sensor (Microsoft, 2015), nor does the sensor have a suitable calibration procedure for modelling its systematic errors. Put simply, although the Kinect 2.0 sensor can produce fully 3D data out-of-the-box with a single sensor, the quality of such data is unknown. Thus the purpose of this paper is to explore and quantify the common error sources within the Kinect 2.0 sensor and relate these to the feasibility of its use in motion capture applications. A brief overview of the operating principles and common error sources will be given, after which a description into the methodology behind quantifying these error sources will commence. The paper concludes with an examination of the results from the Kinect 2.0 sensor, and evaluates some of the best practices as well as the limitations that the sensor maintains.

2. RANGE IMAGING & ERROR SOURCES

A complete description of the operating principles of ToF range cameras such as the Kinect 2.0 can be found in Lange & Seitz (2001). However, a brief description is provided here for clarity. A cone of amplitude-modulated, near-infrared light is projected over the scene of interest. The back-scattered light is focused onto a CMOS/CCD detector array and demodulated at every pixel location (detector site). Four cross-correlation measurements acquired from four separate integration periods are made in this way. In effect, four images are taken, and the signal response at each detector site is cross-correlated to determine the phase difference (from which the range signal is derived) and amplitude (from which we construct intensity images). Thus, collocated X, Y, Z, and amplitude values are obtained for every pixel within the image.

While ToF range cameras provide a number of advantages compared to other systems, they are affected by various error sources. The performance of the sensor and its corresponding error models must be understood in order to maximize the use of the sensor in high-accuracy application. Lichti & Kim (2011) categorized these errors into four distinct groups: random errors; scene-dependent artefacts; scene-independent artefacts; and errors that depend on the camera’s operating conditions. The first of which, random errors, are largely attributable to what is referred to as shot and dark noise (Lange & Seitz, 2001). For the most part, these errors cannot be strictly removed as they are not systematic, but can be reduced by averaging multiple frames over a period of time. Secondly, the scene-dependent distortions, consist of systematic effects which comprise ambient imaging conditions, including external temperature effects (Kahlmann et al., 2006), and internal scattering artefacts, which appear as mixed pixels or a range bias present in background objects within the scene (Mure-Dubois & Hügli, 2007). Scene-independent artifacts include lens-distortions, the range-finder offset, range scale error, periodic errors, and latency errors (Lichti et al., 2010). By integrating a digital camera self-calibration (Fraser, 1997) with the additional parameters and models given by Lichti et al. (2010), it is possible to model some of these completely and remove their effects from the data.

Finally, errors depending on the operating conditions of the camera, as described by Chiabrando et al. (2009) and Kahlmann & Ingensand (2008), include factors such as the camera's warm-up time as well as the sensor's integration time, which can have a direct correlation with the rangefinder offset parameter. Unfortunately, one of the shortcomings of the Microsoft Kinect 2.0 is that it does not currently support modification of the integration time. Consequently, this makes it extremely difficult, if not impossible, to quantify this effect in any meaningful way. For this reason, the effects of changing the integration time in the camera were not examined, as Microsoft has yet to provide a way to make these changes; in addition, many of the issues regarding the correlation between the range-finder offset and integration time are less relevant provided the camera is properly self-calibrated.

3. METHODOLOGY

3.1 Kinect 2.0 Specifications

The Kinect 2.0 comprises multiple sensors, including a red-green-blue (RGB) digital camera, a microphone, and most pertinently, a 3D ToF range camera sensor, which can provide range (X, Y, Z) and amplitude images. Some basic specifications are listed in Table 1:

Table 1: Specifications for Kinect 2.0 system and sensors (Microsoft, 2015)

Specification	Value
RGB camera pixel resolution	1920 x 1080
ToF camera pixel resolution	512 x 424
Framerate (all sensors)	30 FPS
Depth range of ToF sensor	0.5m – 4.5 m
Dimensions of sensor (cm)	25 x 6.5 x 6.5
Nominal principal distance of ToF camera (pixels)	364.5731

Aside from the above information, very little is provided about the sensor, both in the technical documentation (Microsoft, 2015) as well as given how new the sensor is. The sensor provides a framerate slightly above that of modern video systems, and operates within the range of approximately 0.5m to 4.5m. Unfortunately, no information is provided regarding the modulation frequency or pixel size were found within the technical documentation, so these parameters cannot be listed above.

3.2 Basic Performance Testing

3.2.1 <u>Warm-up Time</u>

Time-of-flight range cameras typically require some time for the internal components of the camera to align and stabilize due to the internal temperature gradient across the sensor (Chiabrando et al., 2009). Understanding how this warm-up time impacts the accuracy of the sensor, given that it may be undesirable or infeasible to use sensors that require long wait times before actual data acquisition can be performed.

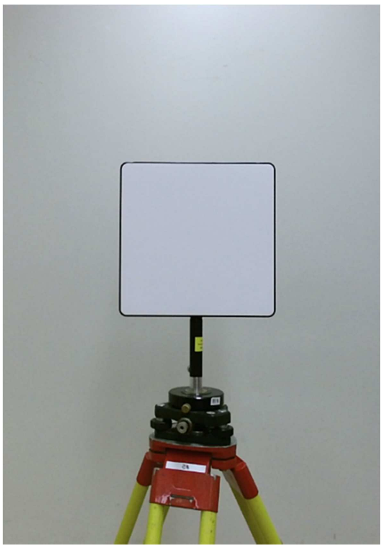

Figure 1: Spectralon planar target with 99% albedo (reflectivity)

This can be quantified by monitoring the temporal behavior of the distance to a known plane over the course of a few hours. In the absence of warm-up effects, the distance should remain constant over time, except for random variations due to shot noise. If there are warm-up effects, the distance will change but any transient behavior should disappear once the temperature gradient within the device has stabilized.

The warm-up test was conducted by setting up the Kinect 2.0 sensor to image a white Spectralon target with 99% albedo (Figure 1) at a nominal distance of 1m and at normal incidence in a temperature-controlled room at the University of Calgary. The Spectralon target was used because it acts as a diffuse surface which reflects 99% of incoming electromagnetic radiation, which is an important factor that can change the level of noise present in the image (Lange & Seitz, 2001). The test was performed for 3 hours, with 30 frames of data collected every 5 minutes. Points lying on the Spectralon target were extracted from the data and a least-squares plane was performed using the methods described in Shakarji (1998). The distance *d* between the camera origin and the plane centre was computed using the standard plane equation:

$$d = -(Ax + By + Cz)$$

The mean and standard deviation of the distances were computed from each set of 30 frames to quantify the sensor performance over time under near-ideal conditions.

3.2.2 Distance and Reflectivity

The ranging precision of a ToF range camera can depend on several factors, two of which are the distance between the camera and object and the reflectivity of the object being imaged. To quantify the impact of these factors, the Kinect 2.0 sensor was used to collect data from both white and black Spectralon targets (99% and 5% albedo, respectively) at varying distances.

The distances were nominally measured, starting at approximately 1m from the Spectralon target in 0.5 m increments up to 4.5 m, whereupon the power of the backscattered signal was too low for the Kinect 2.0 to collect measurements. The data were then processed as in the warm-up experiment.

3.2.3 Vignetting

Vignetting can cause significant power loss near the periphery of the image. The aim of this experiment was to evaluate the impact of this power loss on the measurement precision. A flat wall with homogeneous colour and reflectivity spanning the Kinect 2.0's full field-of-view and was imaged. Both the amplitude and range images were studied for their variations relative to the neighbouring pixels.

3.3 Photogrammetric Self-Calibration with Bundle Adjustment

The photogrammetric bundle adjustment with self-calibration is a popular and effective means for estimating camera specific parameters, also referred to as the internal orientation parameters (IOPs) (Fraser, 1997). With minor modifications, it can likewise be used to estimate IOPs for range cameras (Lichti et al., 2010). For such a calibration, a large target field, such as the one in Figure 2, is imaged so that many targets are visible to the camera when it is placed in various orientations and positions about the target field. For this experiment, images were captured from 36 stations at nine nominal positions in a temperature-controlled room at the University of Calgary (Figure 3). At each if these nine positions, images were acquired from four stations, one each for viewing the target field directly, viewing the target field rotated by a 90^{o} shift, viewing the target field with the camera raised by approximately 40 cm, and viewing the target field with the camera raised by approximately 40 cm and rotated by a 90^{o} shift.

The object-space coordinates of the targets were determined using a high-precision survey-grade Leica HDS6100 terrestrial laser scanner and 3D circle fitting. The target centres in the Kinect 2.0 range imagery were located, and then the data was run through a standard self-calibrating bundle adjustment. To detect if any outliers or erroneous observations were made, Baarda's data snooping was performed, and the residuals were analyzed.

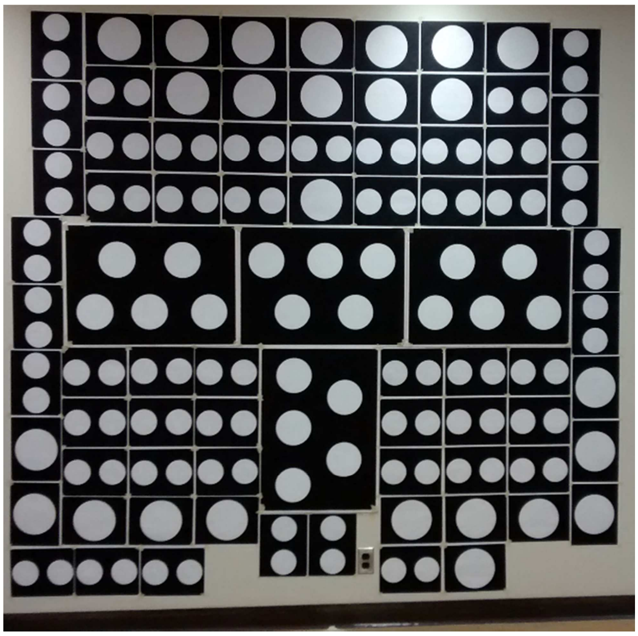

Figure 2: Target field used for self-calibrating bundle adjustment

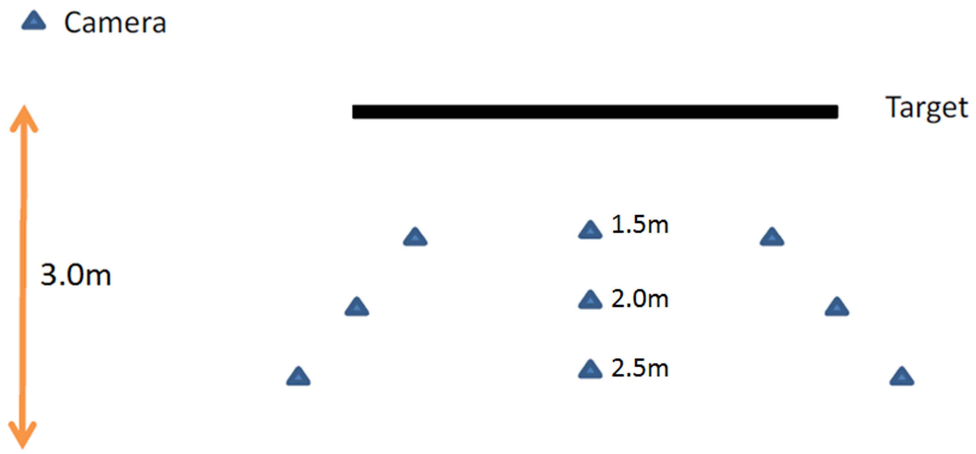

Figure 3: Plan view of camera setup locations. At each setup four sets of images are captured (36 in total).

3.4 Torso Reconstruction In Static Environment

One of the intended applications for the Kinect 2.0 sensor is human motion capture. This experiment examines how accurately a static the human body can be reconstructed. For this purpose, a mannequin was set up in a controlled environment, and several datasets were captured from various positions around the object.

Data for the reconstruction was collected with both the Kinect 2.0 and the aforementioned Leica HDS6100. The datasets two were registered together to the same coordinate system using the 6-inch retro-reflective targets seen in Figure 4. A comparison of the regsitered point cloud data was made to evaluate the differences in accuracy and overall quality of the reconstruction.

Figure 4: Setup of mannequin and retro-reflective targets for registration in static environment

4. EXPERIMENTAL RESULTS

4.1 Basic Performance Testing

4.1.1 Warm-up Time

The temporal behavior of the camera-Spectralon distance is shown in Figure 5. Compared to other range cameras the Kinect 2.0 shows only a 10 minute warm-up period and the overall change in distance is only approximately 3mm over the 3 hour experiment duration. After one hour the variations are less than ±1mm. However, observing such a long warm-up period is likely unnecessary in most scenarios due to the small distance variations that are comparably lower than those reported for other ToF range cameras (Chiabrando et al., 2009) (Lahamy & Lichti, 2012).

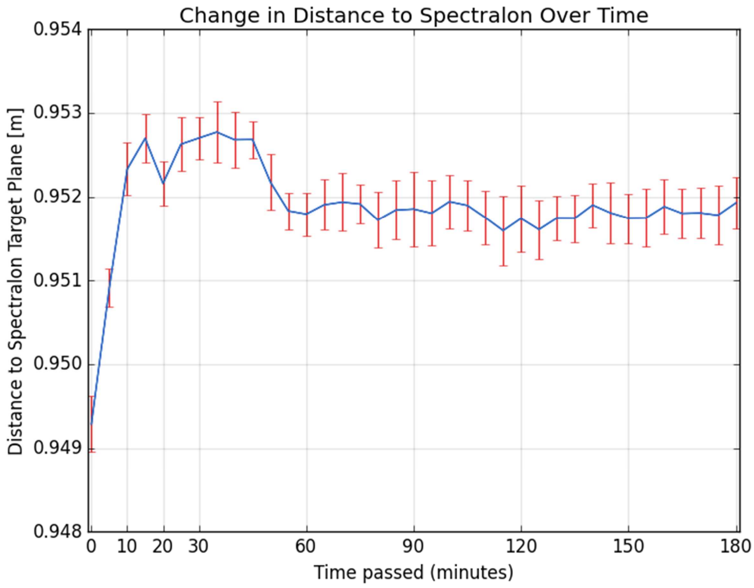

Figure 5:

The reson why the Kinect 2.0 performs as it does with respect to warm-up time may potentially be attributed to the fans mounted to the exterior as part of the cooling system of the Kinect 2.0 sensor. Given that the Kinect 2.0 was originally designed for controlling home entertainment systems, this quick stabilization and long-term consistency could likely have been a part of its engineering design.

4.1.2 Distance and Reflectivity

The results from testing the effect of distance between the camera and object, as well as reflectivity of the object, can be seen in Figures 6a and 6b. Overall, the white Spectralon performed better in terms of plane fit precision as expected, but interestingly had a much more flat trend than that of the black Spectralon comparison. This is to be expected, as there is less power received by the backscattered signal; therefore, reconstructing the phase and amplitude from cross-correlation as outlined by Lange & Seitz (2001) is much more difficult, and is determined with greater noise because the signal-to-noise ratio is much lower.

In general, the measurement precision was between 1.0 and 3.5mm over the entire measurement range. The optimal distance between the object and camera is between 1.0m and 2.5m, as even for dark targets with low overall reflectance, less than 1cm of overall error can

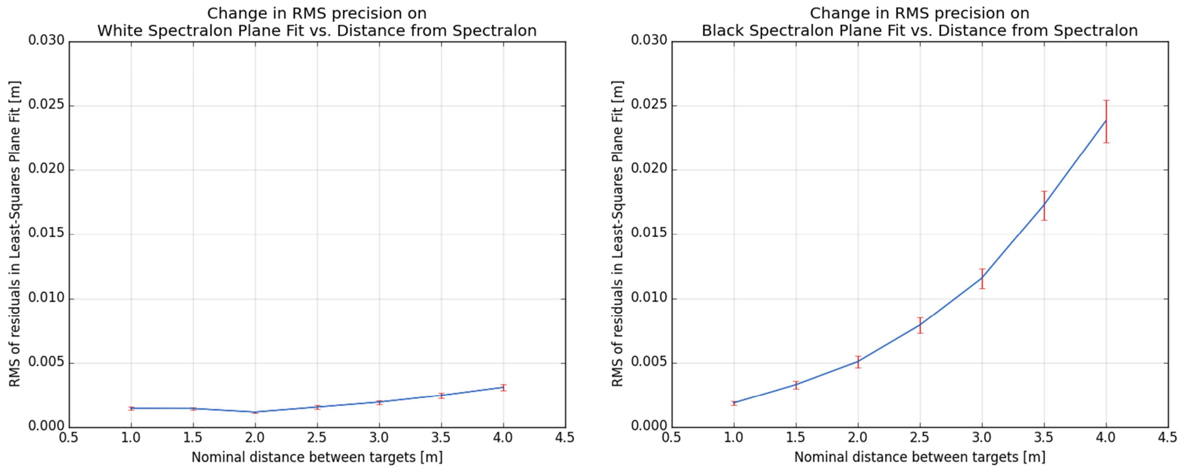

Figure 6: a) Change in RMS precision of plane fit as a function of distance using white Spectralon (left); b) Change in RMS of plane fit as a function of distance using black Spectralon (right)

be observed. Beyond that, the errors grow for darker objects until about 4.5m where the signal strength is insufficient to produce a reliable image from the Kinect 2.0. This is in line with the minimum motion tracking distance of 1.4m recommended by Microsoft, which was likely suggested to prevent both detector saturation and to ensure that a strong enough backscattered signal can be collected.

4.1.3 Vignetting of Backscattered Power

In images captured by the Kinect 2.0, it is possible to show that the power of the returned signal onto the sensor is non-uniform. Unfortunately, the amplitude images are too dark to show this effectively. Instead, the general effect of vignetting can be seen by looking at the cross sections of the amplitude and range residual images. The plots shown in Figure 7a and

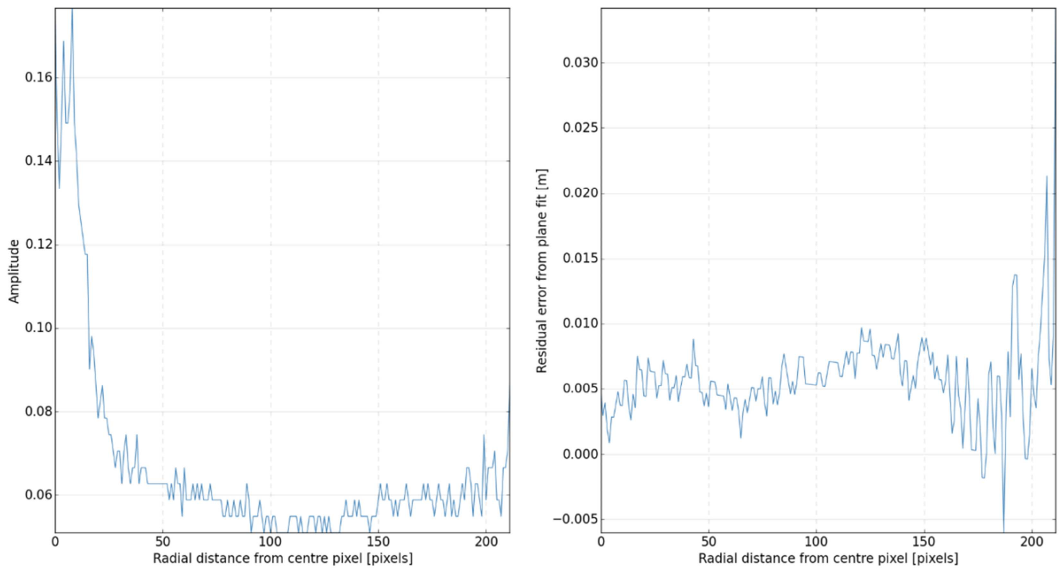

Figure 7: Diagonal cross-section of a) an amplitude image (left); and b) corresponding range residuals (right)

7b show these cross sections of the amplitude and range residual images as a function of radial distance from the centre pixel.

The effect of vignetting on range data quality can be seen in Figure 8b. Though data quality, as indicated by the plane-fit residuals, is homogeneous between the image centre and a radial distance up to about 150 pixels, it degrades markedly thereafter toward the image periphery. In addition, small quantization errors can be seen in the results above; however, they are not very significant in comparison to the other errors seen thus far as they are on the order of about ±1mm. Therefore, it is recommended as a general precaution to image the object of interest within the central 300 x 300 pixels the sensor.

4.2 Photogrammetric Self-Calibration with Bundle Adjustment

Over the course of the self-calibrating bundle adjustment, it was found that only three additional parameters (APs) in addition to the principal point and principal distance parameters were necessary in order to model the systematic effects within the camera. Specifically, these were the radial distortion (k_1, k_2) and rangefinder offset (d_0) parameters as outlined by Lichti et al. (2010). The values and standard errors of these parameters are listed in Table 2:

Once corrected, the overall magnitude of residuals was reduced by up to 89%, as specified in Table 3. Overall the x and y residuals saw the most improvement, at an average of approximately 88% improvement in RMS, while the range residuals improved at a slightly lower percentage of 81%.

Table 2: Estimated Interior Orientation Parameters of Kinect 2.0 and corresponding standard deviations

Parameter	Value	Std. Deviation
x_p [pix]	-4.74	0.15
y_p [pix]	-3.48	0.14
c [pix]	366.45	0.23
k_1 [pix^{-2}]	6.518×10^{-7}	9.536×10^{-9}
k_2 [pix^{-4}]	-1.226×10^{-11}	9.562×10^{-14}
d_0 [mm]	-16.9	1.3

Table 3: Residual RMS improvement before and after calibration

	Before Calibration	After Calibration	Percent Improvement
x [pix]	1.98	0.26	87%
y [pix]	2.43	0.26	89%
range [mm]	66.1	12.6	81%

Overall, individual system calibration significantly reduced the residuals in terms of x and y image measurements and in range. Because the improvement is significant, it is recommended that individual camera calibration be done before use of the device, as the built in model for the sensor does not appear to sufficiently model all the errors present within the data.

4.3 Torso Reconstruction in Static Environment

The geometric quality of the Kinect 2.0 for human body was assessed by comparing the Kinect 2.0 point cloud to a reference cloud taken by a survey-grade laser scanner. Figure 8 depicts a 3D mannequin captured by the Kinect 2.0. On the left of Figure 8 shows the relative differences between the Kinect 2.0 using the manufacturer's default settings (i.e. no additional model for systematic errors). On the right, the same mannequin is shown, however this data was first corrected using the parameters derived in *Section 4.3*.

The scale on the right of the figure denotes the difference between the surface as described by the Leica HDS6100 scanner and the Kinect 2.0 sensor. Ideally, the Kinect 2.0 would match the model of the mannequin described by the Leica scanner; however, this is not the case. More importantly, while the errors between the Kinect 2.0 and the scanner vary between ±3cm, the results visibly improve after calibration of the camera as shown on the right hand side.

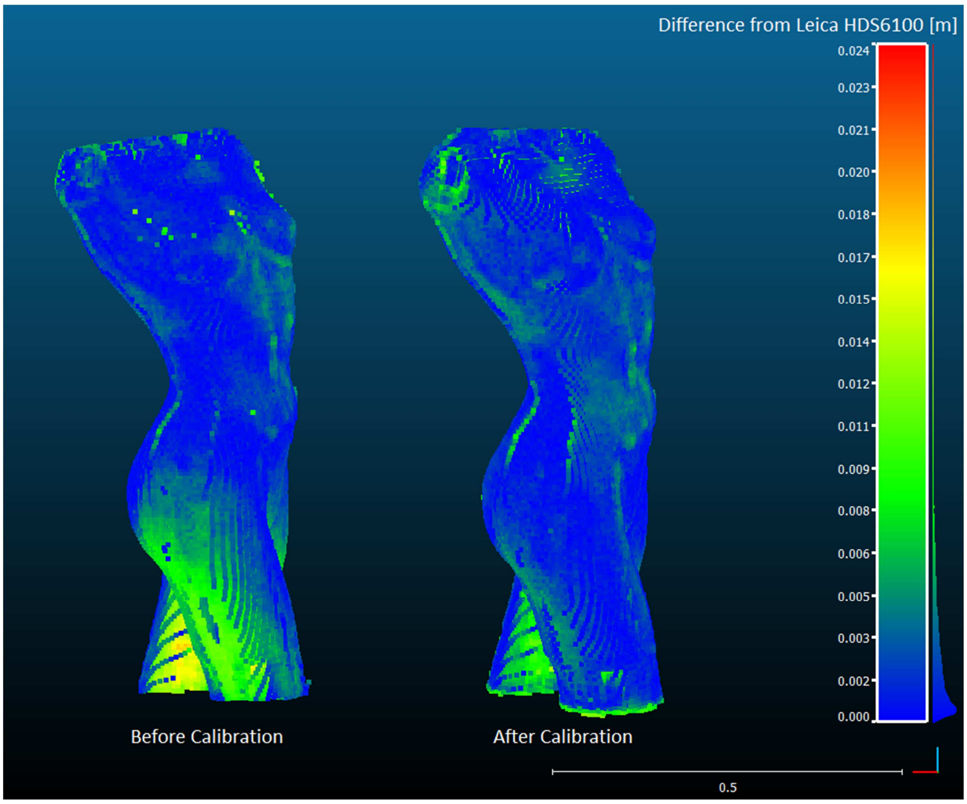

Figure 8: Results of torso reconstruction before calibration (left) and after calibration (right) of Kinect 2.0 data in reference to Leica HDS6100 laser scanner.

These results prove promising for future work in the field of human motion capture, as many of the observations in the registered torso have errors less than 1cm. In general the results were quite good, with mean and standard deviation of cloud-to-cloud distances listed in Table 4. As the mean and standard deviation are quite good (on the order of ±3mm), we can therefore conclude that the on average many of the measurements from the Kinect are quite good, and are generally on the order of centimetre level accuracy or better. Although the above statistics do not change very much after calibration, it is important to note that this is expected, as the mannequin generally had high, diffuse reflectivity, and was contained within the centre of the extent of the image. These conditions are near-ideal, but the results do show that the Kinect 2.0 is capable of achieving very high measurement accuracy.

Table 4: Mean and standard deviation of cloud-to-cloud (C2C) distances before and after calibration

	Mean C2C Distance (m)	**Standard Deviation (m)**
Before Calibration	0.003	0.003
After Calibration	0.003	0.002

5. CONCLUSION

This paper examined the major systematic effects that exist within ToF range camera data, within the context of the new Microsoft Kinect 2.0 sensor. Overall, the camera was shown to have negligible warm-up time, strong dependence on object-camera distance in scenarios where the reflectivity of the object of interest was low, and a decrease in range quality near the periphery of the image due to vignetting. A self-calibration was performed, which improved the observed residuals by approximately 88% for the x and y observations and 81% for the range observations. Finally, reconstruction of the human torso was successful, showing results very consistent to that of the survey-grade Leica TLS used as ground truth in a static scenario.

Several limitations of the sensor were discovered over the course of these tests, which lead to a list of supposed “best practices” when using the sensor. First and foremost, it appears that the manufacturer’s laboratory calibration can be improved upon to remove systematic errors; therefore, individual system calibration is recommended. Beyond that, ensure that only the inner 300 x 300 pixels within the image are used for depth acquisition, due to a loss of signal strength within the image. Moreover, it is desirable that the person or object of interest within the image is at the foreground of the scene, ideally at about 1 – 2.5m away from the Kinect 2.0. Lastly, if possible, ensure that the object or person of interest within the scene is either highly reflective (with diffuse reflectivity), or is wearing brighter clothing in preference to dark clothing, as the Kinect 2.0 measurement quality degrades when measuring objects with low reflectivity.

ACKNOWLEDGEMENTS

This work was supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada.

REFERENCES

Chiabrando, F., Chiabrando, R., Piatti, D., & Rinaudo, F. (2009). Sensors for 3D Imaging: Metric Evaluation and Calibration of a CCD/CMOS Time-of-Flight Camera. *Sensors*, *9*(12), 10080–10096. doi:10.3390/s91210080

Fraser, C. S. (1984). Network Design Considerations for Non-Topographic Photogrammetry. *Photogrammetric Engineering and Remote Sensing*, *50*(8), 1115 – 1126.

Fraser, C. S. (1997). Digital camera self-calibration. *{ISPRS} Journal of Photogrammetry and Remote Sensing*, *52*(4), 149 – 159. doi:http://dx.doi.org/10.1016/S0924-2716(97)00005-1

Kahlmann, T., & Ingensand, H. (2008). Calibration and development for increased accuracy of 3D range imaging cameras. *Journal of Applied Geodesy*, *2*(1), 1–11. doi:10.1515/JAG.2008.001

Kahlmann, T., Remondino, F., & Ingensand, H. (2006). Calibration for Increased Accuracy of the Range Imaging Camera SwissRanger. *ISPRS Commission V Symposium "Image Engineering and Vision Metrology," XXXVI*(5), 136–141.

Lahamy, H., & Lichti, D. (2012). Evaluation of Real-Time Hand Motion Tracking Using A Range Camera and the Mean-Shift Algorithm. *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.*, *XXXVIII-5/W12*, 139–144. doi:10.5194/isprsarchives-XXXVIII-5-W12-139-2011

Lange, R., & Seitz, P. (2001). Solid-State Time-of-Flight Range Camera. *IEEE Journal of Quantum Electronics*, *37*(3), 390–397.

Lichti, D. D., & Kim, C. (2011). A Comparison of Three Geometric Self-Calibration Methods for Range Cameras. *Remote Sensing*, *3*(5), 1014–1028. doi:10.3390/rs3051014

Lichti, D. D., Kim, C., & Jamtsho, S. (2010). An integrated bundle adjustment approach to range camera geometric self-calibration. *{ISPRS} Journal of Photogrammetry and Remote Sensing*, *65*(4), 360 – 368. doi:http://dx.doi.org/10.1016/j.isprsjprs.2010.04.002

Microsoft. (2015, February). Technical documentation and tools. Retrieved February 22, 2015, from https://www.microsoft.com/en-us/kinectforwindows/develop/downloads-docs.aspx

Mure-Dubois, J. C. C., & Hügli, H. (2007). Optimized scattering compensation for time-of-flight camera. In *Proceedings of the SPIE Three-Dimensional Methodes for Inspection and Metrology V* (Vol. 6762). Boston, MA, USA.

Shakarji, C. M. (1998). Least-Squares Fitting Algorithms of the NIST Algorithm Testing System. *Journal of Research of the National Institute of Standards and Technology*, *103*(6), 633–641.

BIOGRAPHICAL NOTES

Jeremy Steward is currently a graduate student studying for his MS.c in Geomatics Engineering at the University of Calgary. He completed his undergraduate degree in Geomatics Engineering at the same institution, where he became involved in research in optical and range-imaging sensors. His research program continues to explore the use of optical metrology methods to help quantify human motion, and structural deformation in both clinical and industrial environments.

Dr. Derek Lichti is currently Professor and Head of the Department of Geomatics Engineering at the University of Calgary, Canada, and Editor-in-Chief of the ISPRS Journal of Photogrammetry and Remote Sensing. His research program is focused on developing solutions for the exploitation of optical and range-imaging sensors for the automated creation of accurate, 3D models for measuring human motion, monitoring structural deformation and compiling asset inventories of the built environment.

CONTACTS

Mr. Jeremy Steward
Department of Geomatics Engineering
The University of Calgary
2500 University Dr. NW
Calgary AB T2N 1N4
CANADA
Tel: +1 403 210 7140
Email: stewardj@ucalgary.ca

Prof. Derek Lichti
Department of Geomatics Engineering
The University of Calgary
2500 University Dr. NW
Calgary AB T2N 1N4
CANADA
Tel: +1 403 210 9495
Fax: +1 403 284 1980
Email; ddlichti@ucalgary.ca
Web site: http://www.ucalgary.ca/lichti/