OPTICAL ORIENTATION MEASUREMENT FOR REMOTE SENSING SYSTEMS WITH SMALL AUXILIARY IMAGE SENSORS

Jürgen Wohlfeil[1] **and Anko Börner**[2]

[1]Humboldt-Universität zu Berlin
Department of Computer Science
Rudower Chaussee 25, 12489 Berlin, Germany

[2]German Aerospace Center
Institute of Robotics and Mechatronics
Rutherfordstr. 2, 12489 Berlin, Germany

juergen.wohlfeil@dlr.de, anko.boerner@dlr.de

KEY WORDS: Exterior Orientation Measurement, Optical Flow Navigation, Line Camera, Pushbroom Scanner, Laser Scanner, Optical Correlator, Point Feature Tracking, Real Time

ABSTRACT:

The accurate measurement of the exterior orientation of remote sensing systems is essential for the quality of the resulting imagery products. Fast orientation changes of the camera immediately before, during or after image acquisition are still serious problems in aerial and satellite based remote sensing. This is due to the fact that in many cases such orientation changes can neither be suppressed effectively nor measured adequately with conventional technology at reasonable costs. In this article, an approach for an auxiliary orientation measurement system is presented that is able to measure a remote sensing system's orientation changes with both, a very high rate and appropriate precision. Two or more auxiliary image sensors are used to measure the orientation changes on the basis of the shift in their images. It is shown that these shifts can be determined by tracking suitable point features through a series of images in real time using a standard mobile CPU for up to 480 images per second. From these shifts the orientation of the whole system can be derived and offset-corrected by conventional orientation measurements. The approach was tested on a test flight with the DLR's MFC line camera and two auxiliary high-speed CMOS cameras. The results are presented and compared to the reference measurements of a high-end INS/GPS system.

1 INTRODUCTION

The demand for cost-efficient aerial and satellite imagery with high resolution and exact georeferencing is growing steadily. The stabilization of the camera and the measurement of its exterior orientation are key problems in meeting this demand. Due to the high resolution of today's imaging systems, their exterior orientation has to be measured with high precision to ensure the quality of the imaging products. Especially the measurement of the orientation with a high angular and temporal resolution is vital for many line-camera and/or laser scanner based systems due to their increasing angular resolution and scan rates. But conventional exterior orientation measurement systems that meet these requirements are usually too expensive, too large, too heavy or need too much power to allow smart and cost-efficient solutions.

A promising approach to overcome this problems is the compensation of fast orientation changes of a line camera by means of small auxiliary frame sensors, positioned next to the main imaging sensors on the focal plane. These frame sensors are exposed at a high rate and the shift of their images' contents provides information about the motion of the imaging system. It has been shown (Janschek and Tchernykh, 2001, Tchernykh et al., 2003) that a special optical correlator hardware can determine the shifts (optical flow) within the images of these sensors in real time. Only these image shifts are saved or transmitted. This is vital because the storage or transmission of the entire auxiliary images would be unreasonably expensive, if not impossible. In a post-processing step, the image shifts can be used to correct the distortions in the line image. Especially if high frequency vibrations affect the camera's orientation during image acquisition this approach achieves very good corrections of the images.

In this contribution we present an approach that is inspired by the successful work of Prof. Janscheck and his colleagues, but suggests two basic innovations. First, we aim to use a standard mobile CPU instead of a special optical correlator hardware. In an empirical test we show that our approach allows to operate at very high measurement rates and can deal with the resulting low image quality and high data rates. Second, we present an approach to derive the exterior orientation of any remote sensing system from the determined image shifts. This way we are able to apply the results to arbitrary sensors. Hence, we also allow the auxiliary sensors to be equipped with separate optics.

In Section 2 of this article we present our approach to determine the optical flow in sequences of small images aiming both, high computational efficiency and precision. In Section 3 it is shown how the orientation changes of the imaging system can be derived from the optical flow and how these results can be offset-corrected with low-rate or small low cost exterior orientation measurement systems. The results of an empirical test of the approach is presented and discussed in Section 4 and concluded in Section 5.

2 EFFICIENT OPTICAL FLOW DETERMINATION

The auxiliary sensors are positioned on the focal plane of a main camera or behind separate optics in a way that they are oriented in different directions. Due to their very small size and a long focal distance the image content moves uniformly in the same direction in the entire image. Hence, it is sufficient to determine a mean image shift vector between every pair of consecutive images. To determine it in real time via image cross correlation requires a special hardware. Unless a special interpolation technique can be used to calculate the image shifts with sub-pixel precision (Janschek and Tchernykh, 2001), true sub-pixel based correlation can't be achieved without major modifications. Another disadvantage is that all structures that are present in a pair of images are used to determine their displacement. Linear structures, which are ubiquitous in urban areas, correlate not only at exactly one shift vector, but also at many different shifts along a line. If there are many dominant linear structures in the image, they can widely overtop the correlation values of remaining non-linear structures that correlate unambiguously.

Our approach only uses structures that correlate clearly at one single shift vector. By ignoring all other structures, not only the effects of ambiguity are reduced drastically, but also a large amount of unnecessary computation time is saved. The findings of (Tomasi and Kanade, 1991, Shi and Tomasi, 1994) provide a very useful solution for the selection and tracking of small, appropriate image regions (features) with very high computational efficiency. According to Tomasi et al. for every pixel of the image the eigenvalues of the 2×2 covariance matrix of image gradients within its 3×3 neighborhood are calculated. A large value of the smaller of the two eigenvalues indicates that the pixel lies on a corner that can be tracked precisely through a series of images. In contrary, if the pixel lies on an improper linear edge or homogeneous area, its smaller eigenvalue is relatively small. Hence, a reasonable number of corners can be selected as trackable features (see Figure 1a and 1b). Because the selection of new features requires a relatively large computational effort, it is not performed at every image, but in periodic intervals.

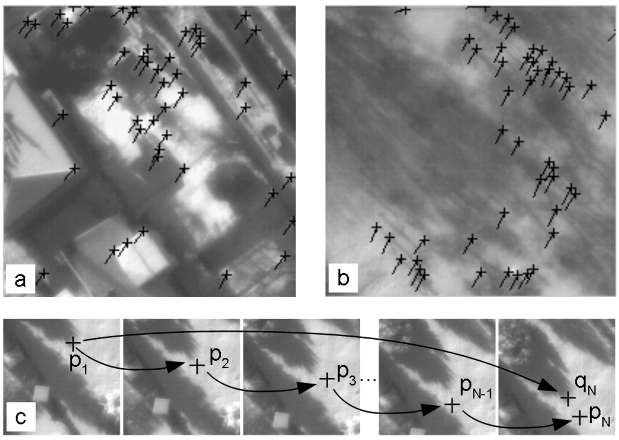

Figure 1: Selected and tracked features from the left (a) and right (b) auxiliary sensor (the trail of the features indicates their shift vectors relative to their position in the previous image). The image series (c) illustrates the technique of feature drift compensation

The selected features are tracked with sub-pixel precision through subsequent images using the pyramidal implementation of the Lucas-Kanade Feature Tracker (Bouguet, 2000). Mainly due to image noise, but also for numerical reasons small errors in tracking a feature from one image to another cannot be avoided. This error accumulates while a feature is tracked through a series of images. I causes the feature to drift away from the originally selected corner and unambiguous tracking can't be ensured anymore. To avoid this problem, we track a feature only through a small number of images and then replace it by a newly selected feature. Additionally, we compensate the remaining feature drift in a very efficient way. Every feature is tracked from its position p_1 in the first image of a small series of N images to the last image in two different ways (see Figure 1c). First, a feature is tracked sequentially through all images of the series to the position p_N. Then it is tracked directly from the first image to the position q_N in last image. While the determined position p_N is affected by the errors of $N - 1$ tracking steps, the error of the directly tracked position q_N is relatively small. By assuming a linear feature drift, q_N can not only be used to correct the final position p_N, but also the positions p_2 ... p_{N-1} in the intermediate images of the series proportionally.

As p_N is expected to be close to the feature's true position in the last image it can be used support the direct tracking step by predicting the position. This makes the direct tracking step more robust and limits the additional computational effort for the feature drift compensation to the effort of one regular tracking step each $N - 1$ images.

Incorrectly and inaccurately tracked features are identified by means of basic statistics. From the two-dimensional shift vectors of all features tracked from one image to another the mean value and the standard deviation is calculated. The features with shifts that diverge more than one standard deviation from the mean value are treated as outliers and excluded from further use. The shift vector of two subsequent images is finally defined to be the mean shift vector of the remaining features. By averaging the feature shift vectors, the tracking errors are reduced. The mean shift vector is finally stored or transmitted with the data of the main sensor.

In the following we explain how the orientation of the whole remote sensing system can be derived from the mean image shifts of the different auxiliary sensors, regardless of the approach that was used to determine them.

3 ORIENTATION DERIVATION

The shifts in the images of the auxiliary sensors are caused by both, translation and rotation of the remote sensing system they belong to. As the translation of the camera is assumed to be measured with sufficient accuracy the only unknown variable is the exact orientation of the system. Nevertheless, the orientation is known to be roughly looking down in nadir direction.

Actually also the height of the terrain affects the shifts in the images. The closer objects are to the camera the faster their projections are shifted over the focal plane while it is being translated. Strictly speaking, the optical flow isn't homogeneous if the terrain isn't completely flat. But thanks to the very small base line between two consecutive images and the large distance of the camera from the ground, small elevations (like houses, trees, etc.) don't have a significant effect on the optical flow. This allows the use of a horizontal plane T at mean terrain height as an adequate approximation for the real ground surface. Deviations of the terrain from T in the range of the altitude above ground (e.g. if flying low over mountains and valleys) may result in systematic errors in the orientation. The compensation of these errors is performed in a later processing step.

The presented approach supports a number of two or more auxiliary sensors. In the following we suppose to have the minimal number of two sensors, one directed to the left and one to the right. For both sensors the mean image shift vectors s_L and s_R are available. The difference in spacial rotation δR of the remote sensing system between the acquisition of two subsequent auxiliary image pairs [I_{L0}, I_{R0}] and [I_{L1}, I_{R1}] can be calculated as follows (see also Figure 2). The center of projection O_0 of the first image pair is known. It is assumed that the system is oriented perfectly in nadir direction. The locations of the object points P_L

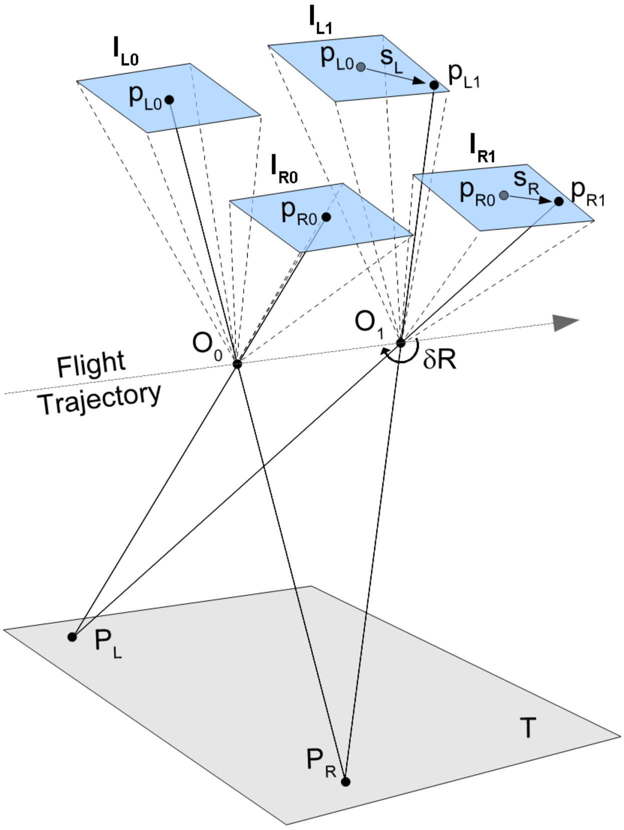

Figure 2: Illustration of the orientation calculation from the mean image shifts s_L and s_R in case of two auxiliary sensors positioned behind one single optics with the center of projection O (Not true to scale).

and P_R can be calculated by projecting the central image points p_{L0} and p_{R0} onto the horizontal plane T at mean terrain height. From the mean shift vectors s_L and s_R it is known that the central image points p_{L0} and p_{R0} have been shifted to the points p_{L1} and p_{R1} in the second image pair. As O_1 can be derived from the flight trajectory, δR can be determined by finding the parameters of the rotation (for example three Euler angles) that best satisfy the collinearity constraints with respect to their projections p_{L1} and p_{R1}. This optimization problem can be solved by a nonlinear least squares algorithm. For practical reasons we used the the Levenberg-Marquardt algorithm from the sparse bundle adjustment implementation of (Lourakis and Argyros, 2004), but a basic Gauss-Newton approach is expected to give appropriate results, too.

The absolute orientation of the system can be determined by integrating the differences in rotation over time. Of course, the resulting orientation is biased by a rotational offset R_0. This is caused by the difference between the assumed and the real orientation of the system at the beginning of the integration. Additionally, significant deviations from the assumed mean terrain height cause the orientation to drift away over time. This drift rate can be modeled adequately by a linear drift rate δR_d. The errors of the calculated shift vectors cause a non-systematic drift, which can only approximately be compensated by the linear drift rate. They are the reason for the remaining errors in orientation calculation.

Depending on the type of the remote sensing system and its absolute orientation measurement technique, the rotational offset can be determined in different ways. Remote sensing satellites typically have a very precise attitude and orbit control system (AOCS) with a relatively low measurement rate. The dynamics of a satellite during necessary movements of gyroscopic actuators, solar panels, antennas, etc. is much higher and can be detected well with the auxiliary imaging sensors. In such cases R_0 and δR_d can be calculated for each time interval between two AOCS measurements separately.

The exterior orientation measurement systems that are used in airborne systems typically consist of an inertial measurement unit (IMU) combined with a differential GPS receiver. For such systems, the bandwidth and the measurement noise of the IMU are the main limiting factors for its effective angular and temporal resolution. By averaging the orientation measurements over a floating time interval the measurement noise can be reduced. By averaging the optically determined orientation over the same interval R_0 can be obtained from the difference of these two averaged orientations. This way, the absolute orientation is provided by the IMU (which is, in turn, stabilized by the GPS and/or a global bundle adjustment (Wohfeil and Bucher, 2009)), whereas fast orientation changes are determined by the auxiliary sensors. The test results for these two variants are presented in the following Section.

4 EMPIRICAL TEST AND DISCUSSION

The presented approach was tested at a test flight over Berlin-Adlershof with the DLR's three-line MFC3-8K camera (Börner et al., 2008) as main imaging sensor using line-rate of about 435 Hz. Two Prosilica GC640 cameras were used as auxiliary sensors, each equipped with a 100 mm optics that was laterally rotated by 20° (see Figure 3). Each camera captured the central 200×200 pixels of its panchromatic CMOS sensor with a frame rate of 480 Hz. Although the extraction of the image shifts is meant to be performed online in a later state of development, at this first test flight all the camera images where recorded and processed offline. The ground sampling distances of the MFC and the auxiliary cameras were similar to allow the determination of orientation changes with an appropriate accuracy for the correction of the line images.

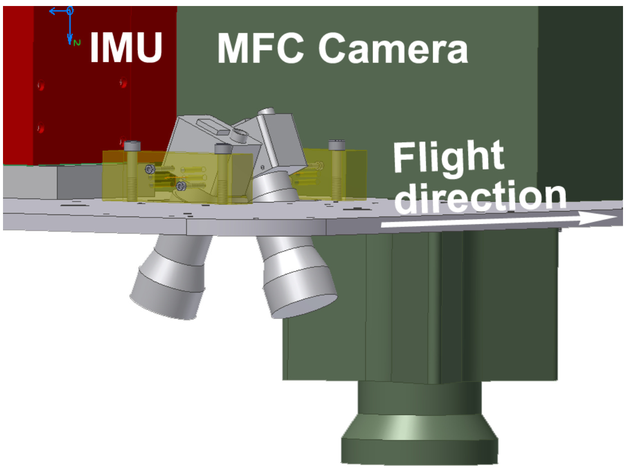

Figure 3: Test flight configuration with the MFC camera, the IMU of the IGI AEROcontrol system and the two auxiliary high-speed cameras mounted on the ground plate

To provide a reference measurement, a high-end IGI AEROcontrol DGPS/INS System was used with its IMU mounted on the ground plate next to the other cameras. The IMU has a sampling rate of 256 Hz but its effective bandwidth is known to be lower than half of this. The limit of the temporal resolution of the optically determined orientation is given by the half frame rate of the auxiliary sensors. To obtain an exact reference measurement a stabilization platform had to be used to suppress orientation changes with frequencies over 100 Hz. The existing ground control points were not usable as a reference because they couldn't be selected with an appropriate precision in the range of hundredths of a pixel.

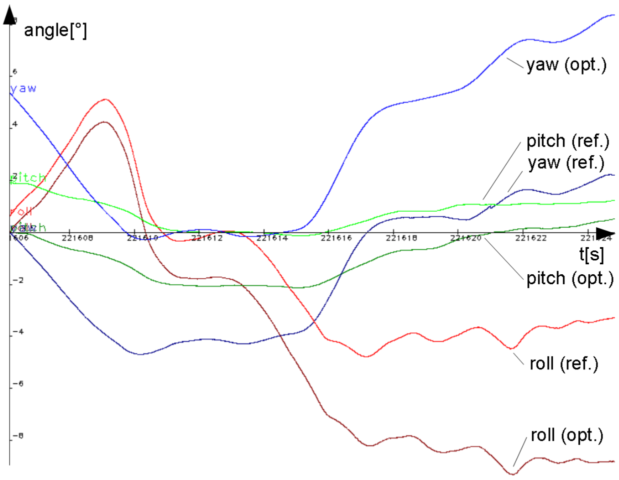

Figure 4: Overview of the orientation during the test flight (Euler angles). The reference measurement is illustrated with light lines and the uncorrected optical measurement with dark lines

A line scan with a duration of almost 20 s is used to evaluate different variants of the approach. About 9000 images were captured by each auxiliary sensor during this time. Due to the very short exposure time, the quality of the auxiliary images is poor. Although a pixel response non uniformity (PRNU) correction is performed, there remains a high amount of uncorrectable image noise. Unfortunately, this noise is highly correlated within the rows and columns of the CMOS sensor, causing challenging conditions for any optical flow detection approach. As we aim to operate with cost-efficient sensors at very high frame rates we want to be able to deal with this.

For the test, the image shifts are determined via feature tracking with drift compensation (FT) and via image cross correlation (CC). For both approaches, the two presented variants of orientation offset correction are performed with exemplary configurations. For the AOCS variant, the orientation of the reference measurement was used in intervals of 1.0 s and 0.1 s to calculate the offset and drift rate of the orientation. For the IMU variant, the reference orientation measurement was overlaid by white noise with an amplitude of 0.01° (around each of the three axes). For the offset correction an averaging interval of 0.2 s turned out to be a good choice with regard to good noise reduction and adequate drift compensation. The RMS and the maximum of differences between the optically determined orientation and the reference measurement are shown in Table 1. For a small part of the test the results are illustrated graphically in Figure 5.

The results in the upper part of Table 1 show that the feature tracking approach gives significantly better results than the cross correlation if the reference measurement interval is one second. Thanks to the feature drift compensation, long-term errors are effectively reduced in the feature tracking approach. If the correction interval is only 0.1 s, the results of both approaches are similar (center of the Table). The RMS remains in a range of a few thousandth degrees, which is a fraction of the MFC's an-

	$\Delta roll$ [°]	$\Delta pitch$ [°]	Δyaw [°]
	AOCS variant (1.0 s correction interval)		
FT	0.00405 (0.016)	0.00634 (0.026)	0.0185 (0.074)
CC	0.06680 (0.329)	0.0456 (0.222)	0.0583 (0.267)
	AOCS variant (0.1 s correction interval)		
FT	0.00099 (0.009)	0.00114 (0.014)	0.00285 (0.047)
CC	0.00188 (0.011)	0.0016 (0.010)	0.00384 (0.026)
	IMU variant		
FT	0.00163 (0.015)	0.00176 (0.017)	0.00235 (0.048)
CC	0.00203 (0.011)	0.00203 (0.008)	0.00308 (0.013)

Table 1: RMS (and maximum) of the differences between optically determined and offset corrected rotation and the reference measurement (Euler angles)

gular resolution of 0.0138°. Hence, its images are expected to be correctable accurately with the determined orientation. Also the IMU variant of orientation offset correction proves to achieve good results for both approaches (lower part of the Table).

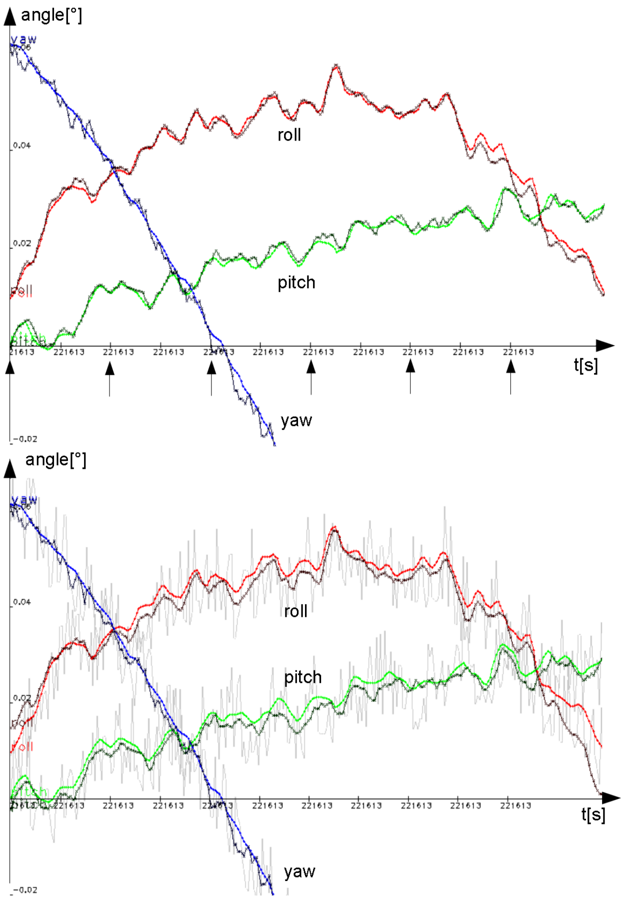

Figure 5: A detailed view of the reference measurement (light) and offset-corrected optical measurement (dark) with the AOCS (upper graph) and IMU correction variant (lower graph). The vertical arrows indicate the AOCS correction intervals. The simulated noisy IMU measurements used for the offset correction are drawn in light gray.

The main goal of the presented approach is the correction of the main sensor's images. Thus we finally applied the determined orientation to the images of the installed MFC3-8K line camera and evaluate the results. Figure 6a shows a Section of the uncorrected line images (the scanned lines are simply joined together). The image in Figure 6b was corrected with the noisy IMU measurements mentioned above. They obviously have an insufficient precision to correct it adequately. For the image in Figure 6c the optically determined orientation was used (feature tracking and IMU variant of offset correction). The image corrected with the reference measurement is not illustrated because there are barely visible deviations from the image in Figure 6c. If no orientation offset correction is performed, there remain only marginal perspective distortions due to the small orientation offset of a few degrees.

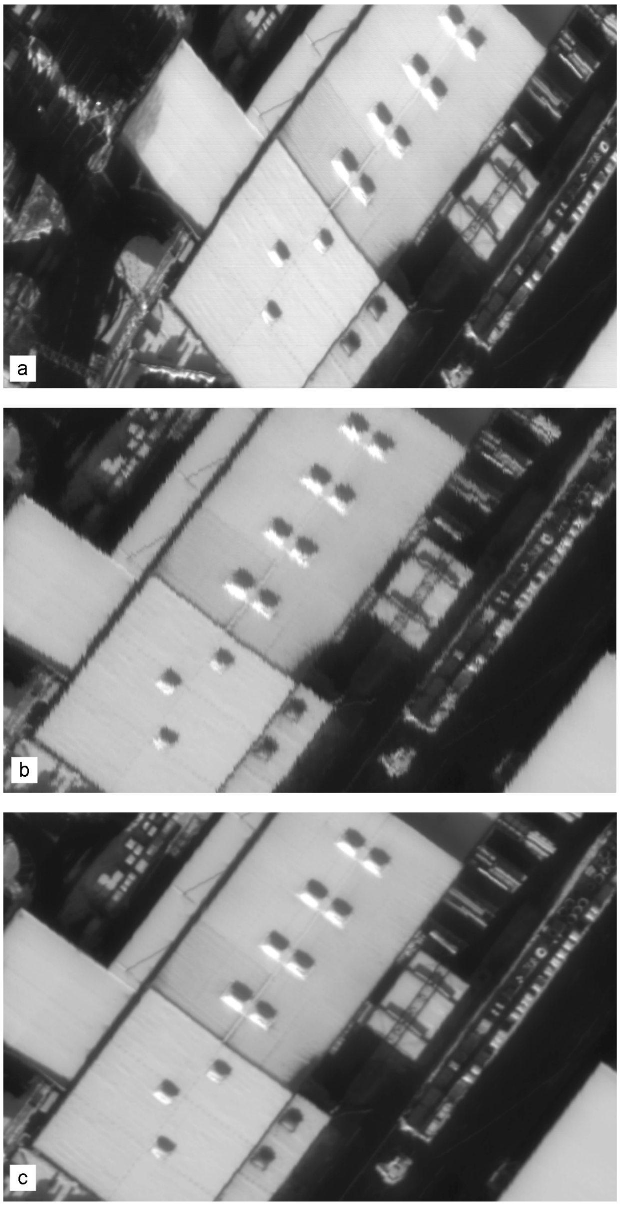

Figure 6: The image of the MFC's central line sensor corrected without (a), with the noisy (b) and with the optically determined (c) orientation

The implementation of our feature tracking approach was able to process 550 frames per second on a laptop with a mobile Intel Core 2 Quad Q9300 CPU for the two cameras simultaneously in the offline processing test. Processing included the loading of the images, the PRNU correction, the determination of the mean image shift and the saving of the resulting shift vector. As exactly these processing steps are necessary for online processing (except that the images are gathered directly from the auxiliary sensors) this is seen as a proof of the online processing capability of the approach.

5 CONCLUSIONS AND OUTLOOK

The results of the empirical test have shown that the optical measurement of fast orientation changes can be performed with standard hardware which can be characterized by low costs, low weight, and low power consumption. Our approach achieved an excellent accuracy under challenging conditions in the airborne case. We could show that it is possible to determine the orientation with a rate of 480 Hz in real time. This allows many new applications for cost-efficient and high-resolution line cameras and laser scanners; for example their installation in small, unmanned aircrafts.

The results give reason to suggest that the presented approach is applicable to remote sensing satellites as well. Very small orientation changes at fractions of the angular resolution of the main sensor could be determined with a very high rate. Additionally, the influence of deviations from the mean terrain height become negligible because to the very large distance from the ground compared to the height deviations.

Especially because of its high measurement rate and precision, the approach can also be used to evaluate exterior orientation measurement systems with regard to the measurement of fast orientation changes. We plan to perform corresponding tests for different INS/GPS systems.

In the empirical test, white noise has been added to the orientation measurement of a high-end INS/GPS system to simulate a small low cost system. This simulation ignores the reduced long-term stability of low cost systems which causes the orientation to drift away faster and in a higher degree. Such drifts can't be compensated by our approach but it allows to perform an initial correction of heavy short-term image distortions in case of a line camera as main imaging sensor. In the pre-corrected images, well-defined tie points can be selected automatically and be used to correct long-term orientation errors via global bundle adjustment, which will be the subject of further investigation.

As water typically doesn't have trackable structures it isn't possible to detect the image shifts for an auxiliary sensor which is entirely capturing images of a water surface. The image shifts of one single sensor aren't enough to determine the orientation. When only two sensors are used it is very likely that this situation occurs. To overcome this problem we propose the use of three or four auxiliary sensors oriented to different directions. Using this configuration it is very unlikely that three of them are directed towards water, except when scanning a large lake or the ocean. As an additional advantage we expect to gain accuracy due to the additional redundancy in optical flow determination.

The presented technique of feature drift compensation could also be used for the correlation based approach. Instead of correcting the shift vectors of single features, the mean shift vectors could be corrected in a similar manner.

ACKNOWLEDGEMENTS

We want to thank Prof. Beate Meffert from the Department of Computer Science of the Humboldt-Universität zu Berlin for her helpful hints. Also our colleagues from the Department of Sensor Concepts and Application of the Institute of Robotics and Mechatronics of the German Aerospace Center we want to thank for preparing and performing the test flight. Especially Jörg Brauchle for the design and construction of the sensor setup and Dr. Sergey Zuev for the processing of the reference measurement.

REFERENCES

Börner, A., Hirschmüller, H., Scheibe, K., Suppa, M. and Wohlfeil, J., 2008. Mfc - a modular line camera for 3d world modelling. Proceedings of Robot Vision, International Workshop RobVis 2008 pp. 319–326.

Bouguet, J.-Y., 2000. Pyramidal implementation of the lucas kanade feature tracker description of the algorithm. Technical report, Intel Corporation Microprocessor Research Labs.

Janschek, K. and Tchernykh, V., 2001. Optical correlator for image motion compensation in the focal plane of a satellite camera. Space Technology 21, Issue 4, pp. 127–132.

Lourakis, M. I. A. and Argyros, A. A., 2004. The design and implementation of a generic sparse bundle adjustment software package based on the levenberg-marquardt algorithm. Technical report, ICS-FORTH, Heraklion, Crete, Greece.

Shi, J. and Tomasi, C., 1994. Good features to track. Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR'94) pp. 593 – 600.

Tchernykh, V., Dyblenko, S., Janschek, K., Seifart, K. and Harnisch, B., 2003. Airborne test results for a smart pushbroom imaging system with optoelectronic image correction. Vol. 5234, pp. 550–559.

Tomasi, C. and Kanade, T., 1991. Shape and motion from image streams: a factorization method - part 3 detection and tracking of point features. Technical Report CMU-CS-91-132, Computer Science Department, Carnegie Mellon University.

Wohfeil, J. and Bucher, T., 2009. A modular, interactive software-concept for radiometric and geometric correction of airborne and spaceborne linescanner images. Remote Sensing f. Environmtl. Monitoring, GIS Apps., and Geology IX Vol. 7478, pp. 1E 1–12.