

Real-time processing of GPS/INS data for on-the-fly quality control in airborne mobile mapping

Klaus Legat^{1,2}, *Jan Skaloud*¹, *Philipp Schaer*¹

¹ Laboratoire de Topométrie (TOPO), Ecole Polytechnique Fédérale de Lausanne (EPFL), CH

² Institute of Navigation and Satellite Geodesy (INAS), Graz University of Technology, AUT

Biographies

Klaus Legat (legat@tugraz.at), born in 1975, holds a Dipl.-Ing. and Dr.techn. in geodesy from Graz University of Technology where he is engaged at the Institute of Navigation and Satellite Geodesy. In June 2005, he joined the geomatics lab at the EPFL in the scope of an Erwin-Schrödinger Fellowship granted by the Austrian Science Fund (FWF).

Jan Skaloud (jan.skaloud@epfl.ch), born in 1969, is a senior scientist and lecturer at the EPFL. He holds a PhD and MSc in Geomatics Engineering from the University of Calgary and conducts research in GPS/INS/sensor integration, mobile mapping, and kinematic geodesy.

Philipp Schaer (philipp.schaer@epfl.ch), born in 1977, received his MSc in environmental engineering from the EPFL in 2003. Afterwards he joined Swissphoto SA in Zurich, where he mainly worked as a project leader in Lidar survey projects. After two years he returned to the EPFL to start a PhD thesis within the TOPO lab. His thesis deals with the project i-Qual (in-flight quality assessment of airborne Lidar, <http://topo.epfl.ch/laserscanning/research.en.php>) that is done in cooperation with Swissphoto and supported by the Swiss Innovation Promotion Agency (CTI).

Abstract

Scan2Map is an airborne remote-sensing system developed under the lead of the TOPO lab at the EPFL. Typically operated aboard a helicopter, it is dedicated for small mission areas (a few square kilometres at the maximum) and offers a sub-decimetre accuracy of the derived mapping products (orthophotos, digital surface- and terrain models). To achieve this high quality independently from ground control, Scan2Map relies on GPS/INS for determining the parameters of exterior orientation of the image sensors (i.e., a digital camera and a Lidar line scanner).

In order to guarantee the complete coverage of a given area of interest, a detailed flight plan needs to be elaborated and it is essential to follow this plan as closely as possible during the flight. However, certain types of problems may remain hidden during the data recording, e.g., insufficient Lidar coverage due to poor ground reflectance. To avoid suchlike situations, an online monitoring tool is currently under development that allows assessing the quality of the recorded data while in the air. Among other calculations, this tool requires the real-time processing of the GPS/INS data.

The paper describes the principles of airborne mobile mapping by Scan2Map, discusses certain aspects of the real-time GPS/INS processing, and shows some initial results obtained during the development of the quality-assessment tool.

1 Introduction

The task of airborne mobile mapping is the recording of image- and trajectory data for the (offline) generation of geographic information products. Scan2Map is a highly versatile airborne remote-sensing system developed at the EPFL during the last three years, partly in cooperation with the private companies UW+R and VNR Electronique. This system is usually operated aboard a helicopter and combines two different types of imaging sensors (a high-end digital CCD camera and a Lidar scanner) with an integrated GPS/INS navigation system. The latter is employed for determining the parameters of exterior orientation (EO) of the image sensors (i.e., position and attitude) during the flight. The EO is required for establishing the geometrical correspondence between the image contents and the real world (as represented by a suitable reference frame). This process is known as direct georeferencing – in opposition to indirect georeferencing where the EO is derived from known ground-control points (GCPs) signalised within the scene. For further information on direct georeferencing, an up-to-date methodology review is found in Skaloud (2006).

The uniqueness and strength of Scan2Map is its great versatility. Due to its low weight (< 15 kg), small dimensions, and the rugged design of its sensor mounting, it can be placed in small airborne vehicles within a very short period of time (usually less than half an hour). Thereby, the sensor system is not attached rigidly to the carrier vehicle but in a flexible manner. This requires that the flight crew comprises – apart from the pilot – at least two persons (i.e., a sensor operator and a system controller), which may be regarded as a (minor) drawback compared to rigidly attached systems. The flexible mounting, however, allows the sensor operator to pinpoint the image sensors to the area of interest while tracking all translational and rotational motions of the system by the embedded navigation sensors. In other words, the orientation of the sensor assembly is only partly dictated by the attitude of the carrier vehicle which is a significant advantage over conventional systems. Due to their great freedom of motion, helicopters allow flying very close to the terrain. Hence, the spatial resolution and accuracy of the mapping products is excellent (typically in the sub- to one-decimeter domain). Essentially, these products include orthorectified images (orthoimages), digital surface models (DSM), and digital terrain models (DTM).

Purpose of the research

Image sensors typically provide masses of data. As a rule of thumb, generating the final product from the data recorded during a single hour requires ten to fifteen hours or more. Furthermore, the costs per flight mission are not negligible, mainly due to the renting of the helicopter and its pilot. Thus, it is important to ensure that the complete area of interest is covered during the flight which requires a careful planning of the whole mission. Among other aspects, the planning has to take into account the time-dependent availability of the GPS satellites in the area of interest. Thereby, topographic obstacles to the signals are of crucial importance – especially in mountainous environments. Clearly, even the most detailed and accurate flight plan is worthless if one cannot control its realisation during the mission for some reason. To minimise the (financial) risk, the planning is usually done in a tolerant style, meaning that at least the most critical parts of the scene are multiply covered (despite of the general overlap of passive images required in conventional photogrammetric processing). This approach usually works well but is not very economic.

It is evident that the mission could be accelerated and the operator and controller could be much more confident about the results during the flight if a real-time quality monitoring of the imaging process was at their disposal. Apart from the aircraft trajectory, other issues must be taken into account referring to the measurement principles of imaging sensors. First of all, such sensors typically apply some sort of central projection (although probably not in the

strict sense). Therefore, the “footprint” of the recorded image or scan line on the actual terrain is irregular, depending on the variation of the topography. Although this effect may be considered during the planning (there is usually a map or even a coarse DTM available), a deviation from the planned trajectory may lead to a data gap that can go undetected during the flight (Fig. 1). In case of Lidar, another problem may arise due to the limited range of such sensors that varies with the reflectance properties of the terrain. The sensor within Scan2Map is dedicated for close-range applications with a maximal ground separation of around 300-400 m (at 80% reflectance). Under adverse conditions, a small position- or attitude deviation of the scanner from its nominal trajectory might cause a loss of the reflected signal. Further, the angle of incidence plays a decisive roll for the geometrical quality of the data as it influences the reflectance and the propagation of the trajectory errors into the object space.

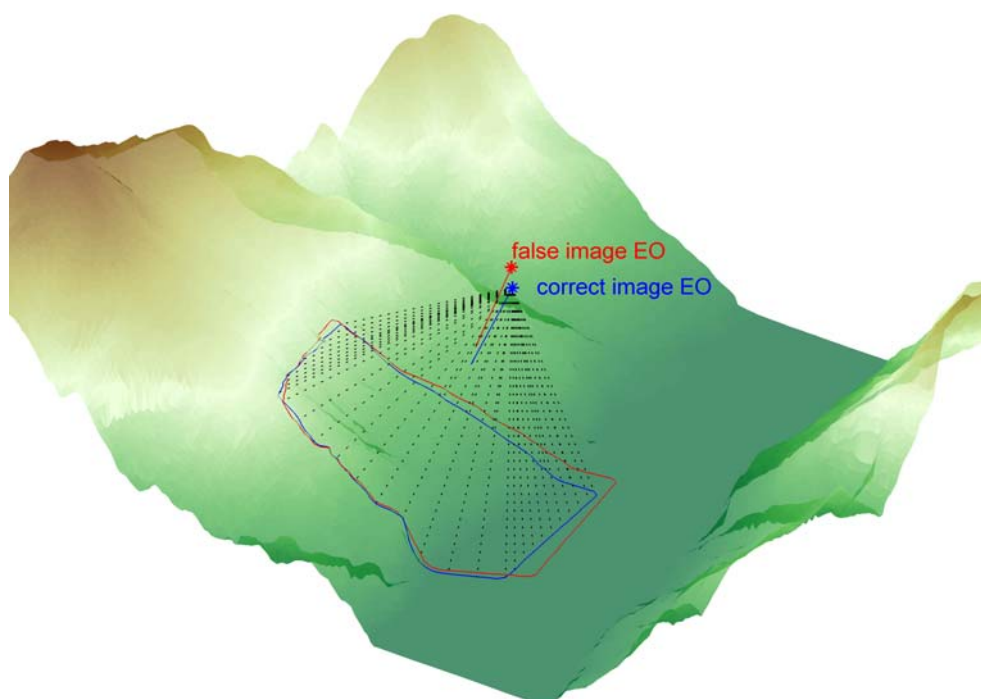


Fig. 1: Exemplary distortion of an image footprint due to erroneous EO parameters (strongly amplified).

The arguments raised above call for a flexible online quality-monitoring tool that allows both supporting the pilot in putting the planned trajectory into practice and enabling the system controller to determine the coverage and accuracy of the recorded data “on the fly”. On the one hand, this requires a real-time processing of the GPS/INS data to obtain the position and attitude of the sensor assembly together with their rates of change; thereby, the quality requirements will certainly be somewhat relaxed as compared to the postprocessing for the final products. On the other hand, if an earlier DTM of the scene is available, the boundary image rays may be artificially intersected with the DTM yielding preliminary image footprints for coverage estimation. In case of the Lidar, the received signal echoes may be analyzed together with the EO and physical characteristics of the laser beam to obtain an impression of the point distribution and quality of the results.

Despite of the mentioned technical aspects, one must not forget some practical issues that are related to the presentation of the online processing results to the flight crew: Both the guidance instructions and the quality figures must be provided in an easily perceivable style to the pilot and controller, respectively. Especially in mountainous terrain, the work load on the pilot is significant and his main responsibility is, after all, the safety of the crew and aircraft.

2 Helicopter-based mobile mapping by Scan2Map

At first, this section briefly discusses some important issues that concern operating a mobile-mapping system aboard a helicopter. Subsequently, the main features of the Scan2Map system are outlined that are important for the navigation matters examined later on.

2.1 System and helicopter dynamics

A helicopter is a relatively harsh environment (with respect to vibrations and accelerations) for high-precision instruments that have not been developed primarily for airborne use. The only sensor being part of Scan2Map that corresponds to aeronautical standards is the IMU (see below). However, this problem is uncritical since the sensor assembly is partly freed from the motion of the aircraft, which allows adjusting its spatial orientation with respect to the terrain. This is achieved by attaching the sensor block elastically to a winch or some other supporting construction installed in the helicopter. Furthermore, the assembly is manually stabilised and pinpointed by the operator according to the needs of the mission. This design absorbs most of the vibrations of the aircraft. The data-recording unit currently involves two rugged notebook computers that are monitored by the controller. The whole processing rack is placed on shock mounts to avoid a damage of its elements.

The freedom of motion of the sensor unit with respect to the helicopter does not permit the use of a GPS antenna directly mounted on the chassis of the aircraft. Instead, the antenna must be rigidly attached to the other sensors for computationally reducing the antenna position to the coordinate centre of the IMU and, subsequently, to the origins of the image sensors. Within Scan2Map, this is achieved by a mast that raises the antenna sufficiently above the chassis to guarantee favourable satellite-reception conditions. The length of the mast is limited by a minimal safety clearance from the rotor (roughly 0.5 m). Hence, the question arises whether the satellite-reception quality is affected by the rotor. Empirical investigations have confirmed that current rotors with nonmetallic blades, once rotating at their nominal operation speed of some 300 to 500 rounds per minute (rpm), have only a relatively small effect on the reception quality: The outage time at the zenith is lower than 4 ms with 20 blade overpasses per second when assuming a three-blade rotor spinning at 400 rpm, a blade width of 30 cm, and a 2 m separation of the GPS antenna from the spin axis. Note that the signal-to-noise ratio (SNR) is somewhat degraded by the rotor: In the RINEX format, e.g., the SNR for the L1 carrier phase drops from an average of 6 to 8 (with 9 being the maximum) to between 5 and 7 (for a specification of RINEX see Gurtner 2006). Complete losses of signal lock and cycle slips occur relatively rarely (mainly during the rotor-acceleration period prior to take-off), and satellites at low elevation angles are usually stronger affected than higher ones.

The motion of the helicopter itself involves six degrees of freedom: Translations and rotations can take place along and about any axis in space, respectively, though some of them in certain limits. The pitch angle, i.e., the inclination about the across axis, usually exhibits the smallest variations (up to $\pm 20\text{-}30^\circ$); although the inclination about the along axis as measured by the roll angle is also restricted, it varies in a considerably larger range (approx. $\pm 50^\circ$). This great flexibility of the helicopter motion can be exploited in a very beneficial way for limiting the error growth within the IMU: Performing sharp turns with high centrifugal acceleration and large roll-angle rates allows for separating the different systematic error effects on the IMU data. These error terms are continuously tracked by a Kalman filter, both offline as well as – in the future – also in real time. A slight limitation to the helicopter motion arises from the lateral placement of the sensor assembly with respect to the chassis. If the helicopter was subject to a steep inclination towards that side where the system is placed, an increased shadowing of the GPS signals would arise. Therefore, these turns are usually avoided even though it is likely that the ambiguities can be correctly re-fixed after brief losses of lock.

Finally, note that the flying speed and height are determined by the requirements of the application and the characteristics of the imaging sensors (see below).

2.2 System architecture

Sensors

The current system architecture involves the following sensors:

- Hasselblad H1 digital CCD camera: This instrument comprises a CCD array of more than 20 million pixels ($\sim 5000 \times 4000$); its focal length is approx. 0.035 m, and the lateral field of view (measured from the optical axis to the borders of the image) is around $30^\circ \times 30^\circ$. At a relative flight height above ground of 350 m (corresponding to an image scale of 1:10 000), the point density on the ground is roughly 120 pt/m².
- Riegl LMS-Q240 line-based laser scanner: The Lidar offers a variable swath angle and up to 10 000 individual range measurements per second that may be flexibly assigned to the number of scan lines (up to 80 per second). The assignment is done in accordance with the desired point spacing on the ground. Using 20 scan lines and a maximal swath angle of $\pm 30^\circ$ at a relative flying height of 350 m yields an average point density of 1 pt/m across track (the scan line is roughly perpendicular to the heading of the helicopter). The along-track point density depends on the flying speed and the number of scan lines. When using 20 scan lines, a speed over ground of 20 m/s yields an along-track density of again 1 pt/m.
- Northrop-Grumman LN-200 tactical-grade IMU: This type of IMU is widely spread in airborne mobile-mapping systems. Its main performance parameters are defined by a gyro-bias repeatability of better than $1^\circ/\text{h}$ and an accelerometer-bias repeatability of better than 0.3 m-g (1 milli-g $\sim 0.01 \text{ m/s}^2$). Thus, this IMU accumulates a horizontal position error of approx. 5 m after 1 min of stand-alone operation. The performance improves substantially when integrating the IMU with GPS inside a Kalman filter that continuously tracks the evolution of the sensor biases. In the integrated mode, the absolute positioning accuracy is dictated by GPS (see below); the best achievable attitude accuracy is at the order of 0.005° (approx. 20 seconds of arc) for the roll and pitch angles, and slightly worse for the yaw angle.
- Javad Legacy GD dual-frequency carrier-phase GPS receiver: In postprocessing mode, the absolute position of the sensor system is derived by a kinematic baseline solution from a nearby base station or from a virtual reference station (VRS). Under favourable conditions, the absolute positioning accuracy is approx. 5-10 cm, presupposed that the carrier-phase ambiguities have been resolved correctly. In case of real-time stand-alone code positioning, the accuracy derives from the GPS standard positioning service (SPS), specified on average to 13 m horizontally and 22 m vertically (Department of Defense 2001). However, the repeatable accuracy is usually much better (i.e., some 2-3 m under good conditions). Although the receiver is capable of processing the correction data supplied by satellite-based augmentation services (SBAS), this feature can only be rarely used in practice as the availability of geostationary satellites in mountainous environments is limited. Last but not least, the receiver provides receiver-autonomous integrity monitoring (RAIM) which is certainly beneficial for the desired application. However, the mountainous surroundings where Scan2Map is usually operated reduce the usability of RAIM as the number of redundant satellites will likely be insufficient for failure detection and isolation/exclusion.

A photo of the system as installed on a helicopter is shown in Fig. 2.



Fig. 2: Scan2Map installed in an Alouette-III.

System integration

The whole measurement system is synchronized to GPS time which is achieved by using the PPS signal and an NMEA data stream (message type ZDA) provided by the Javad receiver. The raw data of the IMU are forwarded to an intermediate processing and synchronization unit specifically developed for the LN-200 by VNR Electronique. From there, the data are transferred via TCP/IP to a notebook for logging. The raw observations of the GPS receiver are stored both internally as well as on a notebook for the postprocessing. The images of the camera are stored in an image bank and the exposure times are marked as external events by the GPS receiver and logged on a PC. Finally, the Lidar measurements are also transmitted via Ethernet and memorised on one of the two notebooks. For the navigation support, another NMEA data stream is currently passed on from the GPS receiver to a moving-map display that is presented to the pilot.

A decisive element for the quality of the final remote-sensing products is the accurate knowledge of the relative orientation parameters between the image sensors and the body frame as realized by the accelerometer triad of the IMU. These parameters comprise the position difference of the centre of projection of the image sensors and the GPS antenna from the IMU coordinate centre – known as lever arm – and a difference in attitude known as boresight (these parameters are different between the camera and the Lidar). While the lever arms can be determined by standard geodetic measurements, the orientation angles are much more critical and the means for their determination differ considerably among the sensors: In case of the camera, a block of images is required and GCPs are beneficial. In case of the Lidar, the calibration techniques are still a field of intensive research. One rigorous approach that considers also some ranging parameters relies on planes of different spatial orientation and inclination measured from flight lines with variable altitudes and course angles (Skaloud and Lichti 2006). In any case, accurate and reliable calibration requires considering the temporal correlations of the GPS/INS data which arise from the filtering/smoothing process and the temporal correlations of the raw sensor data themselves (Skaloud and Schaer 2003).

For further details on the overall system architecture see Skaloud et al. (2005).

3 Real-time processing of GPS/INS data

3.1 Requirements

Quality of the GPS/INS data

Before going into details of the processing, the requirements to be met by the GPS/INS solution for the real-time performance monitoring are briefly examined. This investigation has to address both the quality of the position- and attitude determination.

Position determination

For the sake of simplicity, it would be favourable if the positioning requirements were met by a stand-alone GPS navigation solution based on the SPS (with the possibility to bridge signal outages by inertial navigation). Thus, additional system components could be avoided which would otherwise be required for establishing some sort of differential or relative positioning. (Nevertheless, the medium-term objective is to use RTK at least for small-range missions.)

Concerning the quality monitoring, the effect of errors in the horizontal and vertical coordinates is different. Errors in the horizontal coordinates will translate the image footprints by the same amount. Thereby, the impact is greater for lower relative flight heights due to the smaller image footprints. Assuming a realistic flight height of 200 m, a horizontal position error of 10 m will shift the footprint accordingly. When considering a value of 30° for the lateral field of view of the camera, the side length of the footprint (on a horizontal plane) is approx. 231 m. Thus, the induced error corresponds to less than 5% of the side length. Furthermore, one has to consider that neighbouring sensor positions will probably be subject to similar errors due to the high short-term correlations of the GPS/INS data. In contrast, neighbouring flight lines are usually flown after intervals of several minutes where the temporal correlations of the GPS/INS data are lower. Nevertheless, an accuracy of ± 10 m should be sufficient when taking into account that the pilot probably cannot follow the planned trajectory much closer.

The influence of a height error corresponds to an erroneous scale factor. If the nominal relative flight height of 200 m was in error by up to ± 15 m (the height accuracy of GPS is typically worse by a factor of about 1.5 due to geometrical reasons), the assumed footprint side length of 231 m takes values between 214 m and 248 m, i.e., the maximal deviation is ± 17 m – or ± 8.5 m on both ends. Thus, the resulting errors are similar to those arising from biased horizontal coordinates. Consequently, if only a GPS navigation solution is available, an additional height stabilisation might be helpful (e.g., by using a barometric altimeter after calibration at a known initial height). This also relates to the fact that the vertical channel of an INS is less stable than the horizontal channels (see e.g., Farrel and Barth 1999: Sect. 6.6).

Attitude determination

In the GPS/INS integration, the attitude is almost exclusively determined by the IMU. As was proved by empirical analysis, the attitude accuracy obtained when using SPS position updates in a Kalman filter is around 0.2° for the yaw angle which is twenty times worse than in the smoothed postprocessing; in contrast, the online accuracy of the roll and pitch states is much better. Clearly, the influence of an attitude bias on the image footprints increases with the relative flight height. This effect may be approximated by the product $\Delta h \hat{\psi}$ where Δh is the flight height above the terrain and $\hat{\psi}$ is an attitude error expressed in radians. Thus, when assuming the maximal operation flight height of Scan2Map of 350 m, an attitude error of 0.2° yields a footprint position error of some 1.2 m. Thus, even if the attitude bias was five times greater (i.e., 1°), its influence would still be no larger than that of the position biases.

Processing capacity and specific requirements

A number of real-time implementation issues for Kalman-filtering are discussed, e.g., in Brown and Hwang (1997: Sect. 6.11) and Grewal et al. (2001: Sect. 8.2). Some arguments presented in these references are addressed below:

First of all, one must define for a specific application what “real time” really means; it is unrealistic to assume completely instantaneous measurements. Regardless of the sensor, there is always some latency associated with the output. In case of a GPS receiver, e.g., the delay is some fraction of the update interval. In case of the IMU, the latency of the raw data is negligible but their preprocessing (or rather reorganisation and furnishing with a GPS time stamp) and transmission via TCP/IP requires some time, especially since the latter is a packet-based communication. For the controller to gain an in-depth insight into the measurement progress and preliminary quality features, the delay of the output should not be greater than some three to five seconds – although countermeasures to mend an insufficient coverage will probably only be initiated after finishing a specific flight line (which usually takes a few minutes or more). To fulfil this requirement, the Kalman filter may either be outputting delayed results that are available after the most recent measurement update, or predicted results obtained from a forward projection by applying the dynamical model.

When investigating the required processing power (throughput) of the online computations, the following topics need to be accounted for: the strapdown inertial-navigation processing; the Kalman filter (the workload depends on the dimension of the state vector); the image footprint estimation (being associated with some effort due to geographic data manipulations and spatial intersections when using a pre-existing DTM); and the direct georeferencing of the Lidar measurements (requiring an interpolation of the EO parameters between the GPS/INS solutions). The performance parameters of state-of-the-art notebook PCs allow accommodating all these computations on a single computer, which is certainly also a question of the required level of detail, data handling strategies, and multitasking capabilities.

3.2 Processing strategies

Kalman filter architecture

The online processing is based on a loose-coupling strategy, i.e., the GPS and inertial data are preprocessed before introducing them to the filter. In case of GPS, basically position estimates are used (velocity estimates may be processed in addition). For providing these data, the NMEA format might be used; however, a proprietary format providing more detailed variance-covariance information is better suited. Referring to the IMU, the raw observations include specific-force measurements from the accelerometers and angular-rate measurements from the gyros. These are processed by a strapdown inertial-navigation algorithm that runs at the recording rate of the sensors (400 Hz). In contrast, the Kalman filter is updated at the output frequency of the GPS receiver where 1 Hz will be sufficient (although 10 Hz would be possible).

The dynamical model of the filter is controlled by the coupled error propagation within the strapdown algorithms. The state-vector elements are, thus, the errors in the attitude, velocity, and position estimated for the IMU. These nine elements are complemented by a total of six instrument-bias states making up a 15-state filter. The systematic sensor errors are modelled as Gauss-Markov (GM) models of first order to reflect their temporal correlations.

Note that these tradeoffs yield a suboptimal performance of the filter due to several reasons. On the one hand, the level of detail of the inertial error models is limited. There will be some influences that are not correctly accounted for (e.g., residual scale factors of the sensors). On

the other hand, using GPS update measurements at a data rate of 1 Hz means that the updates are temporally correlated which is not reflected by the chosen filter architecture.

Another important aspect of the GPS/INS integration is a gross-error detection mechanism that should be monitoring the quality of the GPS data before using them for updating the filter states. This is of great relevance especially in mountainous environments where an unfavourable satellite constellation may render a position solution inadequate for updating. The necessity for a gross-error detection corresponds to the fact that a Kalman filter has an infinite impulse response, i.e., the effect of an outlier in one update measurement may “damage” the filter output for extended time periods (Grewal et al. 2001: Sect. 8.3.1).

Initialization

The initialization process mainly refers to determining the initial alignment of the IMU. Generally, the alignment phase can be done both in static and kinematic mode.

Static coarse alignment

The static alignment involves a phase that is known as coarse self-alignment that uses the gravity vector (i.e., the superposition of gravitation and centrifugal force due to the earth rotation) and the earth-rotation vector at the initial position. Comparison with the specific-force and angular-rate vectors supplied by the IMU yields the initial attitude matrix of the platform (see, e.g., Jekeli 2001: Sect. 8.2). Thereby, the roll- and pitch angles follow from the accelerometer data as – in static mode – the only force to be measured should be the (vertical) gravity; in contrast, the yaw angle proceeds from the gyro measurements as the earth-rotation vector has no component in the east direction (this procedure is also known as gyro-compassing). Unfortunately, the instrument biases of the LN-200 are too large to accurately and reliably deduce the initial alignment angles of the IMU in a short period of time. The main difficulty relates to the gyro biases which leave the yaw angle with a large uncertainty. This shortcoming is circumvented by supporting the initial alignment with a digital magnetic compass. Although being subject to magnetic disturbances, the azimuth measurements usually provide a sufficiently good estimate of the initial yaw angle.

Kinematic alignment

The engine and rotor are started only after the coarse alignment. When the rotor spins at its nominal rate, the helicopter is ready for take-off. Once in the air, the pilot initiates a high-speed turn which allows separating the bias terms of the inertial sensors which are not observable in static mode (essentially the gyro biases). Such turns are repeated several times during the flight as the biases do not only vary from turn-on to turn-on but also during operation.

Deduction of the remote-sensing coverage and quality

As already mentioned, the area covered by the image sensors may be simply deduced from the preliminary EO parameters estimated during the flight and using the characteristic geometric features of the sensors. In case of the Lidar, the amount of energy return is taken into account and an error propagation from the air to the ground may be performed, considering the aperture of the laser beam (that determines the size of the “laser dot” on the ground) and the terrain incident angle. This is especially interesting for the future enhancement of the system where the current GPS code solution will be replaced by a precise baseline solution (RTK) in case of small-range missions. In addition, the documentation of GPS signal outages and the position- and attitude-determination quality estimates available from the Kalman filter may be presented in a suitable manner to the controller.

Eventually, it is worth mentioning that certain manoeuvres of the helicopter may considerably deteriorate the GPS/INS quality. This refers especially to quasi-static phases in the air where the craft remains almost at the same spot for some time. Thereby – and probably amplified by side winds –, the platform tends to vibrate at low frequency but with amplitudes of up to several meters which may cause irresolvable situations in the navigation processing. Thus, the online monitoring could raise the controller’s awareness against suchlike situations and allow initiating countermeasures that contribute to a stabilisation of the navigation data. The same is true for less severe problems like when exceeding the nominal flight height or speed.

4 First results and analysis

4.1 Trial characteristics

This section shows the results obtained for a flight trial performed in the Swiss city of Sion (capital of the canton Valais in south-western Switzerland). The presented data refer only to the GPS/INS processing. Note that although only the results of a single flight are shown, these are representative for the envisaged field of application.

Some characteristics of the flight are listed in Tab. 1. In addition, a horizontal position plot and the height profile are shown in Fig. 3 and Fig. 4, respectively. Note that for graphical reasons, the horizontal plot depicts only a 12-minute detail of the whole trajectory comprising two clover-leaf-shaped flight passages. These were performed for calibrating the Lidar by repeatedly passing over a complex of buildings that provided a number of differently oriented and inclined planes. The two periods of almost constant flight height included in the profile plot (located roughly at 700 m and 800 m) also refer to these calibration passages. (For more details on the Lidar calibration algorithms see Skaloud and Lichti (2006)). Note that the height variations during the overpasses demonstrate how difficult it is for the pilot to exactly follow a given 3D flight plan.

Tab. 1: Characteristics of the test flight at Sion.

Feature	Trial characteristics
Objective	Lidar calibration
Surrounding	wide valley, east/west
Flight duration	17 min
Track height (min. / max.)	530 m / 810 m
GPS reception conditions	favourable
GPS type of reference	real station
GPS base separation	max. 4.5 km

The two plots shown in this section were obtained from the GPS baseline solution which was computed post mission using the software package GrafNav. The estimated positioning accuracy (rms) is between 2 and 5 cm for all coordinate directions. The carrier-phase ambiguities were successfully solved during the whole flight. The comparison between the forward and backward processing of the data set that provides some kind of internal quality monitoring yielded deviations of less than 5 cm for the whole on-air period. The number of satellites used in the carrier-phase processing was 6 to 7 on average with short periods of only 4 or 5 spacecraft.




Fig. 3: Horizontal plot of the test flight (detail). The coordinates are relative to the start point.




Fig. 4: Height profile of the test flight (ellipsoidal heights in WGS-84).

4.2 Results and analysis

In the following, a comparison between the high-end postprocessing results and those of the quasi real-time processing is done. Thereby, “quasi real time” means that the processing was done under real-time conditions but in postprocessing. The reason is that the online processing has not as yet been fully implemented (although the required real-time versions of the strapdown- and Kalman-filtering algorithms are readily available) since some modifications to the system-internal communication links need to be done in advance which are currently in progress.

Discrepancies between the results arise from the differences in the GPS processing (kinematic baseline determination based on dual-frequency carrier-phase measurements vs. stand-alone navigation solution based on C/A-code observations), and the differences in the Kalman-filter modelling (25-element state vector vs. 16-element state vector). Note that in case of the latter, the 16-element state vector refers to the version described in Sect. 3.2, where the yaw error (vertical misalignment) was separated into two states for enhanced numerical stability. In case of the 25-element state vector, the above 16 states are complemented by another six residual scale-factor terms modelled as GM processes of first order (one for each inertial sensor) and three additional random constants for the gyro biases. Apart from the extended state vector, the high-end computation comprises also a smoothing phase that improves the results of the forward filtering by a backward sweep. The high-end processing was done using the software package POSProc by Applanix (now part of Trimble).

The differences in the position- and attitude estimates between the two processing variants are of primary interest. The obtained differences for the flight trial are shown in Fig. 5 and Fig. 6. As can clearly be seen from these plots, the results fully correspond to the expectations. The position differences are around ± 3 m for all coordinate directions during most of the time, i.e., the accuracy demands are fully met. Although not shown explicitly, these differences fit very well to the rms accuracy estimates computed by the Kalman filter (which often tend to be too optimistic). Furthermore, the time average of each of the coordinate differences is below 1 m and they tend towards zero, meaning that the quasi real-time results are largely unbiased.

Similar conclusions may be drawn from the attitude results. The initial instability of the yaw difference is due to its high degree of uncertainty during the coarse alignment. However, once the helicopter proceeds at sufficient speed, the yaw error converges very fast. Although the mean value is slightly biased, its residual discrepancy is usually lower than 0.2° .

To sum up, the quasi real-time filtering results meet both the accuracy- and reliability requirements of the given application.




Fig. 5: Position differences between the post-mission and quasi real-time processing variants.




Fig. 6: Attitude differences between the post-mission and quasi real-time processing variants.

5 Summary and outlook

The paper has discussed the principles of on-the-fly quality monitoring of airborne remote sensing data in the context of the Scan2Map system developed at the EPFL. The requirements on the navigation have been elaborated and some initial results obtained from the current development phase have been shown. As the quality of these results is very promising, one can be confident that online quality assessment yields a major contribution to the continuous enhancement of the mapping process in direct georeferencing.

The real-time processing of the GPS/INS data is an important prerequisite for the quality monitoring while on air. Once the current approach based on GPS C/A-code solutions has been successfully implemented in Scan2Map, the next step will be the transition to RTK for small-range applications. The focus of the remote-sensing investigations will increasingly concentrate on the quality of DTMs derived from Lidar measurements, both online as well as post mission.

Acknowledgements

Klaus Legat expresses his gratitude to the Austrian Science Fund (FWF) for supporting his research fellowship at TOPO-EPFL under contract number J2460-N10. Thanks go also to Dr. Julien Vallet for operating the sensor assembly during the flight trial at Sion.

References

- Brown RG, Hwang PYC (1997): Introduction to random signals and applied Kalman filtering, 3rd edition. Wiley, New York.
- Department of Defense (2001): Global Positioning System standard positioning service performance standard. U.S. Assistant for GPS, Positioning, and Navigation, Defense Pentagon, Washington (DC).
- Farrel JA, Barth M (1999): The Global Positioning System and inertial navigation. McGraw-Hill, New York.

- Grewal MS, Weill LR, Andrews AP (2001): Global positioning systems, inertial navigation, and integration. Wiley, New York.
- Gurtner W (2006): RINEX – The receiver independent exchange format version 3.00. Astronomical Institute, University of Berne.
- Jekeli C (2001): Inertial navigation systems with geodetic applications. De Gruyter, Berlin.
- Skaloud J (2006): Reliability in direct georeferencing: An overview of the current approaches and possibilities. EuroSDR workshop EuroCOW on Calibration and Orientation, Castelldefels, Spain.
- Skaloud J, Lichti D (2006): Rigorous approach to boresight self-calibration in airborne laser scanning. Submitted for review to the ISPRS Journal of Photogrammetry and Remote Sensing.
- Skaloud J, Schaer P (2003): Towards a more rigorous boresight calibration. Presented at the ISPRS international workshop on Theory, Technology, and Realities of Inertial/GPS Sensor orientation, Castelldefels, Spain.
- Skaloud J, Vallet J, Keller K, Veyssiere G, Kölbl O (2005): Helimap – Rapid large-scale mapping using handheld Lidar/CCD/GPS/INS sensors on helicopters. Proceedings of the ION GNSS 2005, Long Beach, California, September 13-16.