ACCURATE ORIENTATION FOR AIRBORNE MAPPING SYSTEMS

Jan Skaloud and Klaus-Peter Schwarz

Department of Geomatics Engineering
The University of Calgary
Calgary, Alberta, Canada

ISPRS Commission II, Working Group 1

KEY WORDS: Integration, Integrated GPS/INS-System, Direct Exterior Orientation, Photogrammetry, Filtering Techniques

ABSTRACT

Stringent requirements on the accuracy of attitude determination are currently a major challenge for strapdown INS/GPS integration which is at the core of self-contained airborne remote sensing and mapping systems. This paper reviews the error models for INS/GPS integration and focuses in detail on designing filtering methods for improving attitude accuracy in the bandwidth in which an inertial system does not benefit from frequent GPS position/velocity updates.

Several filtering methods are designed based on the spectral analysis of the raw inertial signal in a dynamic environment. These include a spectral technique for dither spike removal and a class of low-pass Finite Impulse Response (FIR) filters operating in forward/backward manner for achieving zero phase distortion. The orientation performance of the whole system with different filters is evaluated by comparing it to the 'true' attitude information provided by a photogrammetric block adjustment. Results show clearly that the choice of an appropriate filter is decisive for attitude accuracy. Overall, the INS/GPS integration combined with the most suitable pre-filtering method agrees with the external orientation reference to 0.005° (19") RMS over the whole test period, while the flight-line consistency is typically 0.003° (10") RMS. The best filter in the comparison has an RMS seven times smaller than the Butterworth filter which is frequently applied in the industrial designs of INS.

1. INTRODUCTION

Increasing demand for up-to-date information in spatially referenced Geographic Information Systems (GIS) requires development of fast, reliable and accurate acquisition systems. With advances in satellite and inertial (INS/GPS) georeferencing techniques and the ready availability of digital imaging sensors (CCD), a considerable portion of GIS information can be acquired from moving platforms operating on land, water or in the air. Currently, land-vehicle based acquisition systems are capable of delivering coordinates in object space with a typical absolute accuracy of 1-2 meters, see for instance Baraniak (1994). More recently, an accuracy of 10-30 cm has been achieved with the VISAT system (El-Sheimy, 1996). However, the carriers which are in highest demand by industry for semiautomatic mapping are aircraft and helicopter. To design an airborne survey system with an accuracy of few decimeters, the method of direct georeferencing needs further improvements (Schwarz et al., 1993). The general concept of georeferencing by INS/GPS is shown schematically in Figure 1.

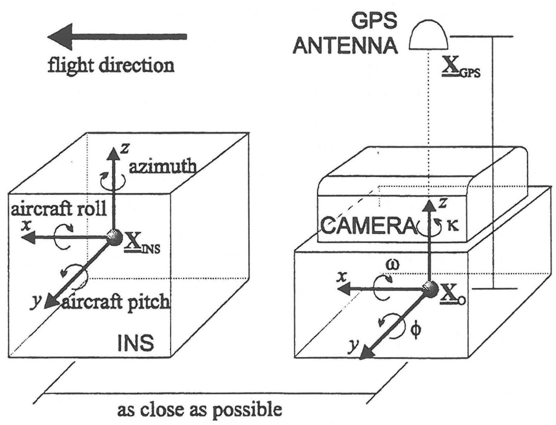

Figure 1. Concept of georeferencing: Determination of image orientation (ω, Φ, κ) and position (X_o, Y_o, Z_o) by INS/GPS.

Its main idea is quite simple: for each captured image determine orientation (ω, Φ, κ) and position (X_o, Y_o, Z_o) of the principal point at the moment of exposure, so that it can be directly used in a chosen mapping frame. An INS/GPS system can provide this information with a quality which depends on the navigation accuracy (X_{INS}, Y_{INS}, Z_{INS}, *roll, pitch, azimuth*) and on the accuracy of calibration parameters relating camera and INS body frames. The accuracies required to fully utilize the resolution of current airborne imaging and gravity sensors are shown in Table 1.

type of sensors	position [m]	orientation
aerial camera, m < 1:2000	0.05~0.1	15"~30"
aerial camera, m > 1:5000	0.75~1	50"~60"
CCD camera or scanner (correlated with pixel size)	0.25~1.0	1'~3'
interferometric SAR	1 ~ 2	10"~40"
airborne gravimetry resolution 0.5 ~ 300 km	acceleration 1 mGal	10"~20"

Table 1. Georeferencing accuracy requirements for airborne survey systems.

The processing chain contributing to the overall accuracy of an acquisition system is affected by the accuracy of the measured image data, INS/GPS position and attitude, system calibration, optical properties of the cameras and the effect of image geometry.

Although the structure of the error model for direct georeferencing is quite general, the significance of each error source is different from one application to the next. For an airborne carrier, the objects of interest are, in general, further from the imaging sensors than for land-vehicle carrier. Thus, they are more sensitive to rotation than to trcnslation. For instance, an attitude errors of 1-3 arc minutes in the VISAT system introduces a positioning error of 1-2.5 cm at a distance of 30 meters from the cameras. The same attitude error at a flying height of 1240 meters would misplace an object on the ground by 1 meter. Thus, the same attitude error contributes only about 10% of the total error budget to the first application, while it may be as much as 60-80% in the second application.

The attitude performance of an integrated INS/GPS is a complex process depending on variety of parameters: quality and type of inertial sensors, sensor placement and configuration, operational aspects and estimation algorithms. Each of these effects needs to be carefully investigated when trying to achieve further improvements at minimal cost. In this paper, the estimation part of the processing chain is scrutinized with the aim to recover the orientation information with the highest accuracy possible.

In Chapter 2, the general concept of INS/GPS integration is introduced and the estimation model is formulated. After reviewing the error characteristics of the primary attitude sensor - a strapdown inertial navigation system (INS) - the error budget of the differential GPS (DGPS) is given. Then, the general concept of integration is presented together with a discussion of long-term sensor error modeling. A discussion on the limits of DGPS for improving inertial attitude determination over shorter time periods concludes this section.

Based on a frequency analysis of the inertial signal in the airborne environment, an improvement of short-time attitude estimation is sought by designing fixed gain filters for data preprocessing. The design of such filters is the central subject of Chapters 3 and 4. Chapter 3 describes a spectral method, originally introduced by Czompo (1990), for removing dithering noise from the gyro and accelerometers outputs. As will be shown later, this method can be very effective in laboratory conditions; however, its contribution to the improvement of attitude accuracy in the airborne environment is rather limited. Chapter 4 describes a class of low-pass filters which can successfully attenuate all higher frequency signals while preserving important signal characteristics in the passband. After stressing the sensitivity of the inertial signal to any phase distortion introduced by filtering, a solution to this problem is found by constructing linear phase finite impulse response filters operating in forward/backward manner. Then, several optimal and sub-optimal methods for constructing these filters are described.

In Chapter 5, the presented filtering schemes are tested with flight data. An independent source of orientation reference is provided by large scale photogrammetry and accurate ground control. Since the detailed description of the test can be found in Skaloud et al. (1996), the emphasis of the analysis is directed towards the evaluation of INS/GPS attitude accuracy achieved with different filters. The significance of the correct filter is illustrated in Figure 2. It compares the Butterworth filter, frequently applied in commercial systems, to a FIR forward/backward filter. The orientation accuracy improves from 0.038° (137") in the Butterworth filter to 0.005° (19") in the FIR filter. In photogrammetric terms, the resulting object space error for a flying height of 2000 m would reduce from 1.33 m to 0.18 m. The paper concludes with a summary of the contribution and with an indication of related problems when using INS/GPS derived orientation for other on-board devices used in the survey.

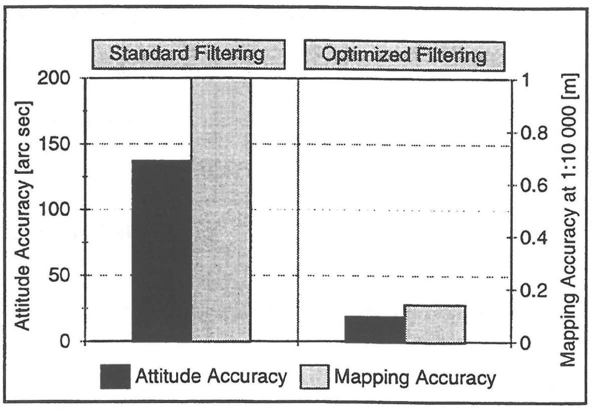

Figure 2. Accuracy gain through optimized filtering.

2. SENSORS AND THEIR ERROR CHARACTERISTICS

This section summarizes the basic properties of strapdown inertial navigation systems, kinematic DGPS with carrier phase observations and the advantages of integrating them. It also discusses some limitations and potential observability problems in complex models.

2.1 Strapdown INS with Ring-Laser Gyros

Inertial navigation systems provide relative position and orientation at high data rates, typically at 50 Hz or higher. If operated in unaided mode, their accuracy degrades quickly due to time dependent systematic errors (Britting 1971). High-quality strapdown inertial systems are usually equipped with ring-laser gyros. The errors of their angular rate output are characterized by the following equation (Savage 1978):

$$d\omega_{ib}^{b} = N_{\omega}\omega_{ib}^{b} + S_{\omega}\omega_{ib}^{b} + d_{\omega} + \mu_{\omega} + \delta_{\omega} \qquad (1)$$

where $\omega_{ib}{}^{b}$ is the vector of angular velocities between the body frame and the inertial frame and $d\omega_{ib}{}^{b}$ its error, N_{ω} is a skew-symmetric matrix describing the non-orthogonality of the axes defining the INS body frame, S_{ω} is a diagonal matrix of scale factor errors, d_{ω} denotes gyro drifts, δ_{ω} is noise due to dither and quantizer effects and μ_{ω} includes other type of gyro errors which are difficult to express analytically, e.g. random gyro drifts including correlated errors, random walk and white noise. The attitude errors of a strapdown system can then be described as a combination of initial alignment errors, including deflections of the vertical, and integrated gyro errors. The well known analysis of misalignment errors shows a periodic, bounded influence on attitude determination while any model for gyro drift leads to an unbounded growth of the attitude errors (Britting, 1971). Schwarz and Wei (1995) showed that uncorrelated errors may also cause attitude divergence. The unbounded growth of this type of errors can be reduced by integration with GPS position and velocities.

2.2 Differential GPS With Carrier Phase Observations

GPS has been recognized in airborne remote sensing, aerotriangulation and airborne gravimetry as a positioning system of superior accuracy, see for instance Schwarz et al. (1993 and

1997). Recent research indicates that post-mission kinematic positioning at sub-meter accuracy is possible with narrow correlator C/A code receivers after double-differencing. These results can be further improved when measurements from GPS are combined with measurements from GLONASS satellites. The accuracy of carrier phase differential positioning using quality receivers is at the level of 10 cm with remote-monitor receiver separation below 30 km, correctly resolved ambiguities, and a consistently good satellite configuration. Over longer baselines, the positioning errors for dual frequency receivers strongly depend on the choice of the tropospheric model. Systematic errors of up to 30 cm are possible over extended periods (Shi and Cannon, 1994).

The GPS velocity estimates are usually obtained from raw Doppler measurements which are also called phase rate measurements. The receiver can theoretically obtain these measurements by counting the number of carrier phase cycles over a small period divided by that time period. Practically, these are averaged in the receiver over time to produce smoother estimates. The resulting raw Doppler measurement for a particular epoch is therefore based on the averaging time interval. Doppler measurements can also be directly derived from carrier phase data. The resulting 'carrier-phase-derived Doppler' is generally smoother than the raw Doppler output from the GPS, since the averaging of the phase measurements is now done at the data rate used. Thus, the GPS velocity based on carrier-phase-derived Doppler would be of higher accuracy, specifically in a low dynamic environment, as typically encountered in airborne mapping. Although empirical values are difficult to obtain because no velocity reference of superior accuracy is available, typical accuracies for a good satellite constellation are at the level of a few cm/s.

Differential GPS with carrier phase observations can also be used for attitude determination by employing a configuration of several antennas. However, the accuracy of the derived orientation parameters is rather limited by antenna separation and receiver noise. The reported values in airborne applications are 6'~30' (Cohen and Parkinson, 1992) and cannot compete with those derived from strapdown INS.

2.3 INS/GPS Integration and Estimation of Long-Term Orientation Errors

Updating an inertial system with navigation information of better quality prevents the unbounded growth of attitude errors. Currently, only GPS can meet such requirements in airborne applications, in terms of precision, range and efficiency.

The aiding of an inertial system with position and velocity data provided by navigation satellites is traditionally implemented via Kalman filtering. The derivation of the dynamic model used by the Kalman filter usually starts with the construction of a full-scale 'truth-error' model, whose order is then reduced based on insight gained into the physics of the problem and a covariance analysis. Typically, the dynamic model is based upon an error model for the three position errors, three velocity errors, and three attitude errors in an INS, driven by sources of uncertainty, and is augmented by some dominant sensor errors, such as accelerometer biases and gyro drifts. If properly done, such a minimal state vector reduces computational burden and accounts for the fact that many of the remaining INS error sources cannot be separated and are best expressed by lumped parameters in the state vector. Hence, models containing 15 to 21 state variables are often appropriate for a high quality strapdown INS.

In the system tested, the correlations in the angular velocity and accelerometer output for each channel are lumped into 'gyro drift' and 'accelerometer bias' state space variables, respectively, which are modeled by a first order Gauss-Markov process with a correlation length of several hours. Although more sophisticated models can be used, they are usually no necessary in low-dynamics scenarios typically encountered in photogrammetric applications. To illustrate this more specifically, the equation for velocity errors in the local-level frame will be discussed in detail. It is of the form

$$\begin{aligned} \delta\dot{v}_e &= f_z\varepsilon_n - f_n\varepsilon_z + b_e \\ \delta\dot{v}_n &= -f_z\varepsilon_e + f_e\varepsilon_z + b_n \\ \delta\dot{v}_z &= f_n\varepsilon_e - f_e\varepsilon_n + b_z \end{aligned} \qquad (2)$$

where the subscripts e, n, z denote east, north and up components, f is the specific force measurement, ε is the misalignment error and b is the accelerometer bias. Equation 2 indicates that the velocity error in a particular channel is generated by misalignment errors coupled with specific force measurements in the other two channels. Since f_z is always large due to gravity, ε_e and ε_n can be observed continuously. In contrast, f_e and f_n have nonzero values only when the vehicle is accelerating in the horizontal plane. Thus, the accuracy of determining ε_z mainly depends on the extent of these maneuvers. On the other hand, the tilt error determination is mainly limited by its correlation with the horizontal accelerometer bias (viz Equation 2). Therefore, in general, better accuracy can be expected in roll and pitch determination because of their strong correlation with vertical gravity. The error in azimuth is only indirectly observable through the horizontal velocity error. Since it takes longer for such an error to develop, azimuth accuracy for a constant velocity flight is usually much poorer that roll and pitch accuracy for a similar flight.

To summarize the above discussion, we can say that DGPS aiding of a strapdown inertial system leads to improvements in attitude estimation within a bandwidth which is a function of the accuracy of the updates, quality of the inertial sensors and the aircraft dynamics. It is also apparent, that the stability of the gyro output is a crucial factor determining the attitude accuracy, since the attitude deterioration due to random-type errors is not detectable by DGPS over a short time period. The subsequent chapters focus on designing filters which will improve short-term attitude estimation by reducing high frequency noise in the gyro and accelerometer outputs caused by aircraft vibration and dithering motion.

3. DITHER SPIKE REMOVAL

Most of the currently manufactured ring-laser gyros undergo dithering motion. This motion is usually referred to as 'gyro dither' and causes vibrations of the whole sensor block. Although dither stabilizes ring-laser gyro output in the long run, it also adds high frequency noise of relatively large amplitude to its output. Manufactures usually implement 'dither-stripping' methods into their data sampling. These methods are designed to remove the effect of dithering from the data. The dither is usually applied in a frequency range of 400-900 Hz. Since most data acquisition systems operate at a much lower sampling rates, the remaining dither frequencies are aliased into the low frequency spectrum limited by the Nyquist frequency. Figure 3 shows the amplitude spectrum of the Z-gyro output, sampled at a rate of 64 Hz, in laboratory and flight conditions. Since the laboratory environment is free of vibration, the peaks in Figure 3(a) at about 9, 20 and 30

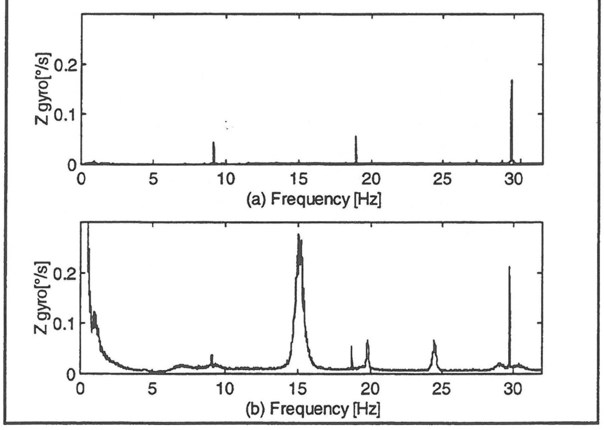

Figure 3. Amplitude spectrum of raw INS signal in (a)-lab conditions, (b)-flight conditions.

Hz correspond to the 3 aliased dither frequencies. They have a maximum amplitude of 0.2° (~720"). Note that these are the amplitudes after the 'dither-striping' methods have been applied. These peaks are also present in the airborne environment, although other high frequency signals, due to vibrations, can be seen (Figure 3b).

Czompo (1990) suggested a special frequency filtering method which detects aliased dither spikes in the inertial raw data and reduces their amplitude to the surrounding noise level. Such a gentle intervention into the spectrum does not change the phase and affects the mean only slightly. It changes the probability distribution of the static data from bimodal to unimodal. In other words, after the filtering, the properties of the sensor noise are closer to the white noise model assumed in the Kalman Filter. The filtering procedure can be described by the three following steps:

Time to frequency domain conversion: Each channel of raw inertial data is divided into time slices of a certain length. The length should be chosen such that a detailed spectrum can be obtained and the Fast Fourier Transform (FFT) can be applied. In our case, subsets of 8192 points were used, which corresponds to 128 seconds of data considering the 64 Hz sampling rate. Then, the data are transformed by FFT to the amplitude and phase spectra for each sensor. The phase spectrum remains unchanged while the dither spikes in the amplitude spectrum will be reduced in the following step.

Dither pattern removal: Due to the aliasing, the dither appears in the amplitude spectrum not as one but as several spikes at different frequencies. The program detects these spikes in predefined windows whose location is known from the static data analysis. The use of non-overlapping windows is convenient since the frequencies where spikes appear can slightly vary for each data slice, but such variation is small enough to allow construction of well separated frequency windows. Practically, the reduction was performed by first computing the average amplitude (m) and standard deviation (σ) from all but 10% of the highest amplitudes, second, the amplitudes higher than ($m + 3\sigma$) were reduced to m.

Frequency to time domain conversion: After the dither pattern was removed from the amplitude spectra, this spectrum and the original unchanged phase spectrum are converted back to the complex spectrum. Then, the inverse FFT is applied to obtain the data in the time domain. The new data with removed dither frequencies are stored and the whole procedure is repeated for the following data subsets.

Czompo (1990) tested such filtering on the INS and INS/GPS derived positions of a land-vehicle. The reference was provided by means of pre-surveyed control points along its trajectory. Section 5 analyses this method for attitude determination in the airborne environment with respect to a 'true' reference provided externally by photogrammetry.

4. LOW PASS FILTERING

This section describes the design of low pass filters which are most suitable for the preprocessing of inertial raw data. From a broad class of low pass filters, sub-classes of filters will be selected which achieve the purpose of getting better short term attitude accuracy while not affecting the long-term performance. Restating this objective in 'filtering language', we can say that the designed filter should strongly attenuate all frequency components which are not due to vehicle motion while not introducing any changes to the frequency band of interest.

The classification of digital filters is usually divided into two major subclasses: recursive, Infinite Impulse Response (IIR) filters, and non-recursive, Finite Impulse Response (FIR) filters. The main difference is that the coefficients of FIR filters operate only on the input data while the IIR filters are also feeding back their own output. Although the description of both filter classes is beyond the scope of this paper we will briefly show that only FIR filters satisfy our design constraints. The input-output relation of an FIR filter can be expressed by the following equation:

$$y_k = \sum_{n=0}^{L} b_n x_{k-n} \tag{3}$$

A filter design that satisfies the above requirements needs to optimize the determination of the $[b_n]$ coefficients with respect to the frequency band of interest.

4.1 Low Pass Filter Design Constraints

Two constraints will be imposed upon the design methods. The first is that of realizability, or causality. This condition implies that the impulse response, which is by definition the system response to a unit sample input at time $k=0$, is equal to zero for $k<0$. Note that this condition is satisfied by the form of Equation 3.

The second restriction is that the filter output can have only a linearly distorted phase. The precise determination of rotation angles is very sensitive to any phase distortion below the cut- off frequency. Since it is impossible to design a filter which would have a zero-phase response and at the same time not violate the causality condition, the generated phase distortion will be such that it can be eliminated later on. Since an FIR filter is nothing but a linear operation on the data, the desired zero phase response can be achieved by applying the same filter again, but in reverse direction. Hence, the output of the forward/backward filter leaves the phase of the original data unchanged.

Linear-phase design is not possible in the case of IIR filters. The linear phase response in terms of transfer function implies that $H(z) = H(z^{-1})$ (Oppenheim and Schafer 1989). In a causal IIR filter design, perfectly linear phase response is not achievable, since the resulting filter would have poles outside the unit circle and would therefore be unstable. For more details, see the extensive literature on this topic. An example of the negative effect of an IIR filter for attitude determination will be shown later in Section 5.

Since all FIR filters are stable by definition, zeros outside the unit circle are no concern in this case and the condition $H(z) = H(z^{-1})$ can be satisfied by requiring zeros to exist in mirror-image pairs inside and outside the unit circle, e.g. z_1, $1/z_1$ (Rabiner and Gold, 1975). Expressing this condition in the time domain requires the FIR filter coefficients to be symmetric such that

$$b_n = b_{L-n}, \quad 0 \leq \text{integer}\left(\frac{L}{2}\right) \tag{4}$$

If L is even, the total number of coefficients is by definition $L+1$. In this case, there is a central sample $b_{L/2}$ about which the coefficients are symmetric. The condition imposed by Equation 4 results in a fixed delay of $L/2$ samples and the corresponding phase response is given by $\theta(\omega) = -\omega(L/2)$, where ω is the normalized frequency in radians. This phase response varies linearly with respect to the frequency.

4.2 FIR Design Via The Ideal Low Pass Filter

This section describes the principle of FIR coefficient derivation by formulating an 'ideal' filter response in the frequency domain and then obtaining its time-domain counterpart by a Fourier transform. Restating again the objective, a filter will be designed with a maximum magnitude response of one in the passband and zero elsewhere, and with a phase response which is a linear function of frequency. This can be expressed in terms of the Fourier transform as

$$H_d(e^{j\omega}) = \begin{cases} e^{-j\omega L/2} & 0 < |\omega| < \omega_c \\ 0 & \omega_c < |\omega| < \pi \end{cases} \tag{5}$$

where H_d is the Fourier transform of the filter coefficients h_d, $\omega_c = 2\pi f_c T$ defines the normalized filter cutoff frequency in radians and corresponds to the delay required to satisfy the causality constraint. The ideal impulse response $h_d(n)$ may be derived by evaluating the inverse Fourier transform of Equation 5 which yields:

$$h_d(n) = \frac{\sin[\,\omega_c(n - L/2)\,]}{\pi(n - L/2)} \tag{6}$$

Since the function in Equation 6 is symmetrical about L/2, the condition on symmetry is satisfied and linear phase response is achieved.

As defined in Equation 6, the impulse response is an infinitely long sequence which has to be truncated in actual filter implementation as $h_d(n) = 0$ for $n > L$. Unfortunately, such abrupt truncation results in an oscillatory effect known as the Gibbs phenomenon, whose magnitude is relatively constant regardless of the filter length (Oppenheim and Schafer 1989). In addition to the choice of cut-off frequency ω_c, the designer has also to choose the filter length (L). Increasing the length of an FIR filter designed via Equations 5 and 6 has four effects, not all of them positive. It

a. reduces the width of the transition bandwidth,
b. increases the frequency of the oscillatory response (Gibbs phenomena),
c. increases the filter delay,
d. increases the computational burden.

From the point of view of post-mission processing the effects (c) and (d) are not important since all 'future' data are available and real-time computation is usually not required. The effect (a) is positive for our purpose because less of the undesired signal leaks into the pass band. Thus, the only problem remaining is the Gibbs phenomenon. Since this undesirable response results from the abrupt truncation of the infinite series, different approaches can be taken to alleviate this problem. A few of them are described in the following sections.

4.3 Time Domain Window Functions

The simplest approach to dampen the oscillatory effect is to perform a somewhat smoother truncation of the infinite series of filter coefficients. Such a smoothing can be achieved by convolving the filter coefficients with a window function. This section discusses the properties of window functions commonly used in FIR filter design. The filters resulting from these functions will be tested in Section 5 on inertial data.

The abrupt truncation of an infinite sequence can be expressed as a product of two sequences

$$h(n) = h_d(n)\, w_R(n) \tag{7}$$

where $h_d(n)$ are the original filter coefficients derived by Equation 5 and $w_R(n)$ is the rectangular window function

$$w_R(n) = \begin{cases} 1 & 0 \leq n \leq L \\ 0 & \text{elsewhere} \end{cases} \tag{8}$$

As has already been mentioned, an improvement in the magnitude response can be expected by using a non-rectangular window with tapered ends. To retain the linear phase characteristics such a window must also be symmetric about its midpoint. The amplitude spectrum of a desirable window has two key characteristics: A narrow main lobe which results in a smaller transition band and side lobes that rapidly decrease in energy for increasing ω, since this results in less oscillatory effect in both the pass-band and the stop-band of the filter. Four of the most commonly used window functions: Bartlett, Hamming, Hanning and Blackman will be later tested and compared on inertial data. The formulas for their derivation can be found in the signal processing literature, e.g. Oppenheim and Schafer (1989), and are therefore not shown here.

4.4 Optimal FIR Low Pass Filter Design

Although the design of a low pass FIR filter by means of a window function is a very straightforward and powerful approach, this method is sub-optimal in the sense that it is possible to design a lower-order FIR filter that has equally good frequency-response characteristics. To design an optimal FIR filter requires the definition of an approximation criterion and the derivation of an algorithm which would satisfy such a requirement. The definition of a particular approximation criterion is twofold: the choice of an error measure to be minimized (e.g. maximum absolute error, sum of square errors, etc.) and the choice a weighting function on the approximation error that enables the designer to choose the relative size of the error in different bands.

It is obvious that the choice of the approximation criteria varies according to specific application requirements. In the case of raw inertial data, one is concerned about any changes to the passband. Attenuation variation in the stop-band is not so critical since this part of the spectrum contains a relatively small part of the total signal power.

The two most popular optimal design techniques will be further tested on actual data, because each has a different error measure. The first is based on minimizing the maximum absolute weighted error between the desired and the approximated filter response and is formulated as a Chebyshev approximation problem. A number of rather complex techniques have been devised for obtaining the solution of this problem (Oppenheim and Schafer, 1989). An iteration technique known as the Remez exchange algorithm is used in the following. It has been described in Parks and McClellan, 1972. The other optimal filter tested in the following can be derived by minimizing an error objective function that reflects the weighted mean-square difference between the ideal amplitude response and the amplitude response of the filter. The symmetrical impulse response ($w(n)=w(L-1-n)$) is then obtained in the least-square sense from a system of linear equations (Oppenheim and Schafer 1989).

5. ANALYSIS OF RESULTS

This section evaluates the orientation performance of the INS/GPS integration approach under actual flight conditions, using the filtering schemes described in Sections 3 and 4. References values were obtained by means of large-scale photogrammetry. The test methodology is briefly summarized in the following section. A more detailed description of the test can be found in Skaloud et al. (1996).

5.1 Test Flight Scenario And Reference Trajectory Accuracy

By hard-mounting a camera on a rigid platform together with the INS, a fixed orientation angle difference between the INS and the camera can be derived. The processing method used to derive the camera orientation is a photogrammetric bundle adjustment. It is a conventional photogrammetric technique which takes advantage of the geometric strength of the overlapping photographic images. By fitting the interlocking image array to ground control points, the orientation parameters are georeferenced and can then be compared to INS/GPS derived orientation parameters.

A well-defined photogrammetric test field close to Cologne, Germany was used to assess the orientation accuracy of an INS/GPS integrated system. Nine photogrammetric strips, three of them repetitive, were flown over the test area in early July 1995 (Figure 4). The length of the strips differs from approximately 1 to 4 km. A subset of 77 center located images was chosen to form a photogrammetric block with 80% forward and 60% side overlap. The average flying height of about 900 m and the 15 cm camera focal length resulted in a photo scale of 1:6 000. The accuracy of the perspective centers of the photographs were estimated by traditional block adjustment using all 47 ground control points. The predicted standard deviations (σ) of the orientation angles range from 0.001°(3") to 0.004°(15"). Their mean values are about 0.002° (7") in roll and pitch and 0.001° (3") in azimuth.

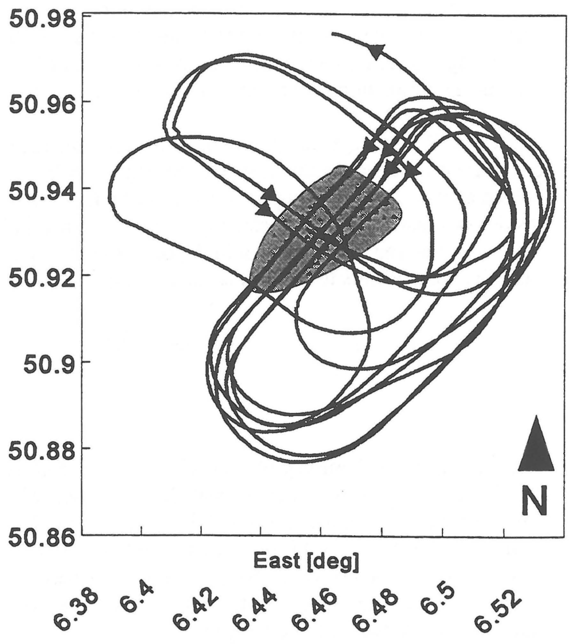

Figure 4. Test flight scenario.

The airborne data collection system contained pairs of Trimble 4000 SSE and Ashtech Z12 dual frequency receivers, a Litton LTN-90-100 strapdown system with gyro drift rates of about 0.03°/h, a portable computer and a Zeiss RMK A aerial camera with a precise shutter pulse output. The time synchronization was realized via a data collection computer receiving raw INS output and GPS data (Ashtech Z12) together with a 1 pulse per second (PPS) signal provided by the receiver. The shutter pulse of the aerial camera was recorded by the receiver (Trimble 4000 SSE) in GPS time.

Since the GPS antenna and the INS system are physically displaced from the perspective center of the imaging sensor, a constant displacement vector $\boldsymbol{dr} = (dx, dy, dz)^T$ and a constant misorientation $\boldsymbol{dR}_p^{\,b} = f(\delta_1, \delta_2, \delta_3)$ exists between the sensors. The components of the translation vector were measured by using conventional surveying techniques before the flight mission. This information is used as a lever arm correction in the INS/GPS integration. Since the body axis of the INS cannot be physically observed, the misorientation matrix $\boldsymbol{dR}_p^{\,b}$ was determined 'in-flight' by means of the first attitude reference. This value was then subtracted from the INS attitude and the remaining differences were considered as the orientation errors of the INS/GPS system.

A key assumption for the attitude comparison is that no changes in relative position and orientation between the imaging device, INS and GPS antenna will occur. Although all sensors were hard-mounted to the same platform, the design of the aerial camera holder most likely prevented perfect camera stabilization in the horizontal (pitch and roll) channels. Therefore, only azimuth will be compared here. Pitch and roll estimated via INS/GPS integration should theoretically outperform azimuth determination; see Section 2.3.

5.2 Accuracy of Azimuth Determination

The performance of the filtering methods presented in Sections 2.3, 3 and 4 will be evaluated by the following statistical parameters computed from the disagreement of the azimuthal reference on 60 points: Mean (m = Σx/n), Standard Deviation (σ

= $[\Sigma(x-m)^2/(n-1)]^{1/2}$), Root Mean Square (RMS = $[\sigma^2+m^2]^{1/2}$) and Maximal Deviation from the mean (MaxDev). The graphical representation of the RMS values is shown in Figure 5. The legend defining the abbreviations for the filter types in this Figure can be found in Table 2. Table 2 also shows other statistical characteristics.

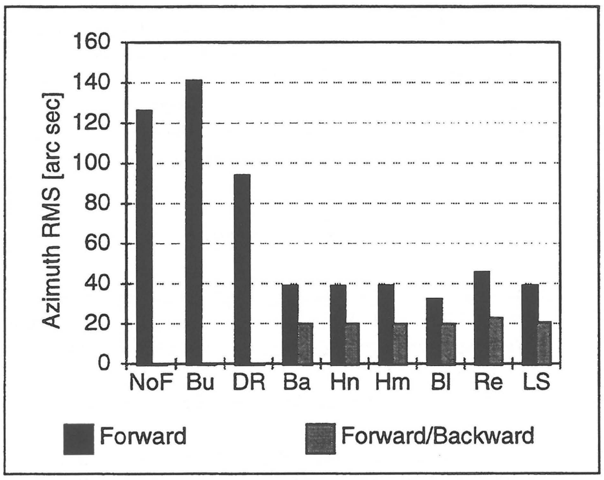

Figure 5. Comparison of filtering methods for azimuth determination by INS/GPS in the airborne environment.

	Abrev.	STD [°]	Mean [°]	MaxDev [°]
No Prefiltering	NoF	0.025	0.025	-0.044
Butterworth	Bu	0.038	0.010	0.166
Dither Removal	DR	0.021	0.014	0.031
Bartlett	Ba	0.005	0.001	0.013
Hanning	Hn	0.005	0.001	0.013
Hamming	Hm	0.005	0.001	0.012
Blackman	Bl	0.005	0.001	0.012
Remez	Re	0.006	0.002	-0.010
LS	LS	0.005	0.001	0.013

Table 2. Comparison of filtering methods for azimuth determination by INS/GPS in the airborne environment. The shown statistical values are computed from comparison to photogrammetry derived attitude using 75 points over a 50 minute period. The filter cut-off frequency is 8 Hz in all cases and the filter order is N=120 for FIR filters and N=9 for the recursive Butterworth filter.

The first bar in Figure 5 and the first row in Table 2 characterize the accuracy of the azimuth determination by INS/GPS integration without INS data prefiltering. They are the result of comparison on 75 control points over 50 minutes of flight time. The detailed comparison for each reference point and the σ-values per flight line are depicted in Figure 6a. Although the relative accuracy for each flight line is within 30", the comparison between them shows systematic changes, the most obvious one being between the third and fourth flight line. This indicates that either a short-term gyro instability or high frequency noise (aircraft vibrations, dither) are integrated into the attitude solution. As to long-term effects, the overall results are stable within a ± 0.050° (3') bound. It can therefore be concluded that aiding the inertial system with DGPS velocity and position dampens the otherwise unbounded growth of the INS attitude errors. When considering the requirements of the most demanding airborne application (viz Table 1), the residual azimuth errors are still too large. As can be seen from Figure 5 and Table 2, some of the inertial data prefiltering methods will help to bring the RMS down.

The comparison of the different filters shows a rather consistent performance for six of them. The best results were achieved when using linear phase non-recursive low pass forward/backward filter designed with a Blackman window as a prefiltering method. The RMS value computed from the comparison to the reference for this filter is 0.005° (19"), which is overall five times better than INS/GPS integration without inertial data prefiltering and seven times better than using a Butterworth filter which is often applied in commercial designs. Moreover, practically zero-mean distribution of the residuals indicates that no distortion has been introduced to the passband by the prefiltering process. The detailed plot of the residuals in Figure 6b indicates no outliers.

Referring again to the statistical comparison in Table 2, it can be said that the dither removal method did only marginally improve the attitude solution. The explanation for the small improvement is on the one hand the low signal power of the dither spikes despite their high amplitude and the fact that the use of the FFT does not result in a very accurate estimate of the amplitude. In other words, this method will always produce an amplitude spectrum with a standard deviation of 100% of the estimated value, independent of the data length (Kay and Marple, 1981). This means, that the location of the dither spikes may not always be correct and the average amplitude may not be accurately established. A better accuracy of the amplitude spectra can be achieved by either constructing the spectrum with finer resolution than needed (by using a longer data set) and then average certain number of consecutive frequencies, or by combining several spectral estimates from a partitioned data set. Unfortunately, using such alternative methods would create difficulties in the spectrum-to-time domain transformation.

All tested linear phase low pass filters improved the azimuth estimation very significantly. Moreover, the approach of achieving zero phase shift by forward/backward filtering introduces further improvements with respect to the forward-only filtering and thus demonstrates the sensitivity of the inertial data to any phase distortion. This fact is further illustrated by showing the influence of a non-linear phase IIR Butterworth filter on the attitude

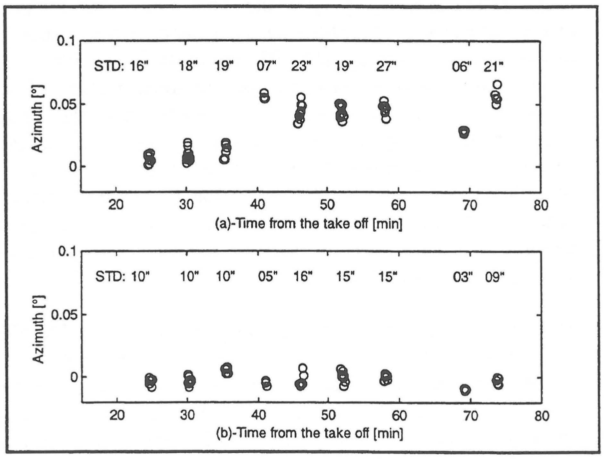

Figure 6. Differences in INS/GPS azimuth determination with respect to 75 control points, (a)-No data prefiltering, (b)-Prefiltering with the forward/backward low pass filter designed with Blackman window.

determination. This type of filter is often used for INS data filtering prior to the mechanization. As can bee seen from Figure 5 and Table 2, the use of a Butterworth low-pass filter actually affects the azimuth estimation negatively.

Comparing the different windowing functions, it can be concluded that all methods performed equally well. This is due to the fact that the inertial signal power at the 8 Hz cut-off frequency is rather low and, therefore, a broader transition band is not as critical as achieving good attenuation in the stop band. Although some frequencies in the stop band have significant power, the attenuation of 30-40 dB of the Bartlett window seems to be sufficient in this case. However, if the use of a different aircraft causes stronger vibrations, the application of window functions with better attenuation would be more appropriate (e.g. Blackman 60-80 dB). If aliased frequencies of high amplitude would be closer to the cut-off frequency, then using a higher-order filter can make up for a larger main lobe width to sharpen the transition band, while maintaining approximately the same ripple attenuation. Empirical testing showed that a filter order of approximately 120 yields the best results.

Within the class of optimally designed low-pass filters, the least-squares approach seems to slightly outperform the Remez-exchange algorithm. However, this difference is small and could be possibly caused by the type of 'RMS' evaluation which obviously favors the square error minimization. If designed with heavy weights in the pass band, the 'optimal' low pass filters work comparably well with respect to the window design approach.

6. CONCLUSIONS

The research reported in this paper showed that prefiltering of inertial data is a vital step for achieving high accuracy with an integrated INS/GPS system for airborne mapping and remote sensing. It removes up to 80% of the high frequency errors which cannot be modeled by frequent GPS updates. Results also show that the choice of the filtering method affects performance significantly.

Since inertial data are very sensitive to changes introduced by the filter to the pass-band, it is important to design a filter with minimum phase distortion. Nevertheless, any low-pass filter introduces a phase shift in the data. However, if this shift is linear, it can be removed by applying the same filter again but in the reverse direction, because filtering is a linear operation. Since it is impossible to design a stable IIR low-pass filter with exactly linear-phase response, only the class of linear-phase FIR was investigated. The optimal and sub-optimal design of these filters is presented. Testing using empirical data showed a comparable performance with an improvement of the orientation accuracy of about five times. Testing also showed that the spectral method of 'dither spike removal' does not affect attitude determination in the airborne environment where aircraft vibration cause the prevailing short-term noise in the inertial signal.

These results lead to the conclusion that the method of INS/GPS integration with inertial data prefiltering will satisfy the orientation requirements for even the most demanding airborne applications if high frequency errors are removed by pre-filtering and medium and low frequencies by GPS-aided Kalman filtering with appropriately chosen sensor error models. In addition, the calibration procedure must resolve the misalignment between the georeferencing sensor and the imaging sensor to the same level of accuracy and the mounting of the two components must be such that no differential change in rotation can occur during the duration of the mission.

ACKNOWLEDGMENTS

Funding for this research was obtained through an NSERC Operating Grant of the second author. Flight time and expertise was provided by Rheinbraun AG-Department of Photogrammetry, Cologne. This support is gratefully acknowledged. Thanks are due to the Stuttgart Institute for Photogrammetry, and especially Mr. Michael Cramer, for the time and effort spent on organizing and executing the flight test and to Mr. Werner Schneider for performing high quality photogrammetric measurements.

REFERENCES

Baraniak, D.W., 1994. Creating High Accuracy Digital Maps Using GPS and Stereo Digital Imaging Technologies, Wisconsin Professional Engineer, May/June 1994, Vol. 35, No. 4, pp. 12-14.

Britting, K.R., 1971. Inertial Navigation System Analysis, Wiley-Interscience, New York.

Cohen, C. and Parkinson B., 1992. Aircraft Applications of GPS Based Attitude Determination, *Proc. of ION GPS-92*, Fifth International Technical Meeting, pp. 775- 282, Albuquerque.

Czompo, J., 1990. Use of Spectra Methods in Strapdown ISS Data Processing, *Proc. of KIS-90*, Banff, Canada, Sep. 10-14.

El-Sheimy, N., 1996. A Mobile Multi Sensor System for GIS Applications in Urban Centers, *Proc. of ISPRS'96*, Commission II, WG 1, Vienna, Austria, July 9-19.

Kay S.M. and Marple, S.L., 1981. Spectrum Analysis - A Modern Perspective, *Proc. of IEEE*, Vol. 69, No. 11, Nov.

Oppenheim, A. V., and Schafer, R.W., 1989. *Discrete-Time Signal Processing*, Englewood Cliffs, N.J., Prentice-Hall, Inc.

Parks, T.W. and McClellan, J.H., 1972. A program for the Design of Linear Phase Finite Impulse Response Filters, *IEEE Trans. Audio Electronics*, Vol. AU-20, No. 3, pp. 195-199, Aug.

Rabiner, L.R. and Gold B., 1975. *Theory and Application of Digital Signal Processing*, Prentice-Hall.

Savage, P.G., 1978. Strapdown Sensors, in AGARD Lecture Series No. 95, NATO, Neuilly sour Seine.

Shi, J., and Cannon M.E., 1994. Precise airborne DGPS positioning with a multi-receiver configuration: Data processing and accuracy evaluation. *Proc. KIS-94*, Banff, Canada, August 30 - September 2.

Skaloud, J., Cramer M., and Schwarz, K.P., 1996. Exterior Orientation By Direct Measurement of Camera Position and Attitude, *XVII CONGRESS of ISPRS - Spatial Information from Images*, Vienna, Austria, July 9-19.

Schwarz, K.P., Chapman, M.A., Cannon, M.E, and Gong, P., 1993. An Integrated INS/GPS Approach to the Georeferencing of Remotely Sensed Data, *PE&RS*, Vol. 59, No. 11, pp. 1667-1674.

Schwarz, K.P. and Wei M., 1995. "Modeling INS/GPS for Attitude and Gravity Applications", 3rd International Workshop for High Precision Navigation, Stuttgart, April 3-6.

Schwarz, K.P. and Glennie, C., 1997. "Improving Accuracy and Reliability of Airborne Gravimetry by Multiple Sensor Configurations", *Scientific Assembly of IAG-97*, Rio de Janeiro, Brazil, September 3-9.