

# Precise and efficient high-frequency trajectory estimation for LiDAR georeferencing

Florian Pöppel<sup>a,\*</sup>, Andreas Ullrich<sup>b</sup>, Gottfried Mandlbürger<sup>a</sup>, Norbert Pfeifer<sup>a</sup>

<sup>a</sup> Department of Geodesy and Geoinformation, TU Wien, Wiedner Hauptstraße 8/E120, 1040 Wien, Austria

<sup>b</sup> RIEGL Laser Measurement Systems GmbH, Riedenburgstraße 48, 3580 Horn, Austria

## ARTICLE INFO

### Keywords:

Laser scanning  
Sensor orientation  
Inertial navigation  
Vibrations  
Least-squares adjustment

## ABSTRACT

Laser scanners mounted on moving platforms allow for efficient large-scale 3D mapping using light detection and ranging (LiDAR). Because the laser scanner is moving with respect to the earth, its trajectory (position and orientation over time) must be known in order to georeference the scanner measurements to an earth-fixed coordinate system. This is commonly realized through integration with satellite and inertial navigation systems. Sensor fusion algorithms, whether filter-based or adjustment-based, then fuse the navigation data to obtain an estimate of the platform trajectory. Errors in this trajectory cause errors in the 3D point cloud through the georeferencing process. Most processing workflows therefore include a step which optimizes the trajectory based on the LiDAR data itself. In this contribution, we focus on two related aspects of trajectory estimation and LiDAR georeferencing: Firstly, we analyze the impact of high-frequency trajectory dynamics, which cause oscillating errors in the trajectory and negatively impact point cloud precision if the inertial sensors' sampling frequency is too low to properly resolve them. This implies the necessity of recording and processing inertial measurements at a sufficiently high frequency, which can drastically increase computational effort of the sensor fusion algorithm, especially for adjustment-based approaches. Thus, secondly, we propose a method for performing adjustment-based trajectory estimation with high-frequency inertial measurements which efficiently uses downsampled low-frequency inertial measurements within the adjustment while recovering the high-frequency trajectory dynamics from the original measurements. Analysis and processing are performed for two separate datasets acquired with two different platforms, a quadcopter uncrewed aerial vehicle (UAV) and a crewed fixed-wing aircraft. For the former, we demonstrate through analysis of point spread on planar surfaces that a 200 Hz sampling frequency for the inertial measurements is insufficient and leads to reduced point cloud precision. In both cases, the proposed methodology is shown to precisely recover the high-frequency trajectory while drastically reducing memory usage and runtime compared to performing the adjustment with the high-frequency inertial measurements.

## 1. Introduction

Kinematic laser scanning is a widely used surveying technology for topographic mapping, specifically for the efficient acquisition of large-scale 3D point clouds. In kinematic laser scanning, a laser scanner is mounted on a moving platform such as a land vehicle, drone or aircraft. The laser scanner makes use of light detection and ranging (LiDAR) together with a scanning mechanism which varies the deflection angle and thereby the direction of the laser beam. In combination with the motion of the carrier platform, this enables 3D sampling of the environment. The laser scanner thus measures the range to the target and the angle of the laser beam. In a first step, these polar coordinates are converted to Cartesian coordinates, which are still relative to the

scanner itself. Georeferencing is then the process of transforming the measurements from the laser scanner's coordinate system to an earth-referenced coordinate system. For this, the platform's trajectory, its position and orientation over time, must be known. Commonly, a global navigation satellite system (GNSS) receiver and antenna, and an inertial measurement unit (IMU) are mounted on the platform together with the laser scanner. GNSS, IMU, and the laser scanner itself operate with different measurement frequencies (i.e., at different temporal scales). While GNSS provides long-term absolute positioning, usually at a rate of 1–10 Hz, the IMU is responsible for short-term relative orientation and outputs measurements between about 100 and 1000 Hz. High-end survey-grade LiDAR systems output individual measurements at

\* Corresponding author.

E-mail addresses: [florian.poeppel@geo.tuwien.ac.at](mailto:florian.poeppel@geo.tuwien.ac.at) (F. Pöppel), [aullrich@riegl.com](mailto:aullrich@riegl.com) (A. Ullrich), [gottfried.mandlbuerger@geo.tuwien.ac.at](mailto:gottfried.mandlbuerger@geo.tuwien.ac.at) (G. Mandlbürger), [norbert.pfeifer@geo.tuwien.ac.at](mailto:norbert.pfeifer@geo.tuwien.ac.at) (N. Pfeifer).

<https://doi.org/10.1016/j.isprsjprs.2025.03.007>

Received 9 November 2024; Received in revised form 8 March 2025; Accepted 10 March 2025

Available online 26 March 2025

0924-2716/© 2025 The Authors. Published by Elsevier B.V. on behalf of International Society for Photogrammetry and Remote Sensing, Inc. (ISPRS). This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).




Fig. 1. Common ranges for the measurement frequencies of GNSS, IMU and LiDAR as part of a laser scanning system.

very high rates, commonly between 100 kHz and 3 MHz for single-channel devices (Fig. 1). The platform's trajectory is estimated by fusing absolute position information from GNSS with measurements of linear acceleration and angular velocity from the IMU. However, errors in GNSS and IMU measurements cause errors in the trajectory, which in turn cause errors in the georeferenced point cloud. In this context, we categorize trajectory errors as *low-frequency* and *high-frequency*, with low-frequency corresponding roughly to the frequency domain of the GNSS, and referring to anything above as high-frequency. Low-frequency trajectory errors cause slowly varying deformations in the point clouds, which manifest, e.g., as discrepancies between overlapping flight strips. High-frequency trajectory errors on the other hand impact the point cloud on smaller time scales, creating wave-like artifacts in the point clouds and generally increasing point spread and reducing precision.

The industry-standard processing workflow for kinematic laser scanning involves first estimating the trajectory using a Kalman filter, and then correcting low-frequency trajectory errors in a subsequent strip adjustment (e.g., Glira et al. (2016) and Jonassen et al. (2023)). Lately, holistic methods have been proposed based on batch non-linear least-squares (NLS) adjustment, which integrate the laser scanner measurements into the trajectory estimation (Brun et al., 2022; Pöpl et al., 2024). Such methods are the main focus of this article, as the tight coupling between IMU and LiDAR is integral in reducing both low-frequency and high-frequency errors. However, the observations regarding vibrations and IMU sampling rate apply in general to inertial navigation and thus also to other sensor fusion approaches, such as Kalman filters.

### 1.1. Inertial navigation and vibrations

There are a number of factors related to inertial navigation which influence trajectory quality (Farrell, 2008, Ch. 11), whether in standalone inertial navigation or when fused with other sensors. Many deterministic errors, such as constant offsets or output scale errors can be determined in a laboratory calibration and corrected for, taking into account various factors such as deformation of the material casing or sensors themselves due to temperature changes. On the other hand, stochastic errors such as measurement noise or bias drift cannot be pre-calibrated and can only be compensated in-run by integration with other sensor data (Farrell et al., 2022). Statistical models for the stochastic processes are derived from specific calibration procedures and used to improve the in-run estimation of the stochastic error components.

Many errors in inertial sensors and subsequently the inertial navigation system (INS) output are caused or amplified by the presence of vibrations (Groves, 2013, Ch. 4). The most critical type of error is the vibration rectification error (VRE), a systematic measurement error which depends on the frequency of the sensed signal. The underlying physical principles causing this error depend on the sensor's operating principle and its implementation.

Generally, the magnitude of acceleration and angular velocity due to vibrations adds to that of the underlying platform motion. Since accelerometers and gyroscopes can only output measurements in a certain range, values outside the measurement range are clipped. Strong vibrations may therefore cause the measurement range to be exceeded, even if the actual motion of the platform is within the specified range of the sensor. As vibration often exceeds platform motion in terms of

maximum signal amplitude, it may also implicitly cause errors which would otherwise be invisible, e.g., non-linearity errors occurring only at large signal magnitudes.

In kinematic laser scanning, the inertial measurement unit is mounted rigidly with respect to the laser scanner itself and the GNSS antenna<sup>1</sup>. In practice, the different components will not always be completely rigid and vibration will lead to a time-varying relative orientation thus breaking the assumption of rigidity.

As these vibration-related errors are generally hard to calibrate for and cannot be reliably compensated in-run, most IMU/INS manufacturers recommend keeping vibrations as low as possible through mechanical means (e.g., Hexagon (2023), SBG Systems (2024)). Use of appropriate materials and dampening (Periu et al., 2013; Li et al., 2015; Liu et al., 2024) can reduce vibrations and their impact on the payload sensors, but this is not always practically possible due to constraints in size or weight.

From a theoretical standpoint, it is clear that insufficiently high sampling frequency causes a mismatch between the continuous physical quantity and its measured discrete representation. According to the Nyquist theorem, a signal can be represented without aliasing error if and only if the sampling frequency is at least twice the highest frequency of the signal. In other words, sampling with a frequency of  $[N]$  Hz only allows resolving signal components with frequency below  $[\frac{N}{2}]$  Hz, with aliasing effects occurring for signal components of higher frequencies. Most higher-frequency acceleration or angular velocity is caused by vibrations (i.e., periodic oscillating motion,  $\geq 10$  Hz) either of the carrier platform (e.g., UAV rotors) or the scanner system itself (e.g., rotation of scanning mirror). In the worst case, aliasing causes a periodic motion with zero mean (i.e., stationary on average) to look like a time-varying bias error which in turn compounds into position and orientation errors within strap-down inertial navigation. As most IMUs internally sample higher than the output frequency, the output is usually anti-aliased thereby mitigating this error. However, the higher-frequency components are now removed from the signal and thus missing from the estimated trajectory.

The errors induced by aliasing or low-pass filtering, i.e., insufficient sampling rate in the presence of vibrations, are often only viewed in the context of overall navigation performance. For most applications, the high-frequency trajectory components are unimportant as long as the overall navigation performance (i.e., the low-frequency trajectory components) are not negatively affected by the vibration. Thus, vibrations are often viewed as undesirable noise and filtered out from the signal (e.g., Alam and Rohac (2015), Suwandi et al. (2019) and Li et al. (2024)). However, removing the vibrations from the inertial measurements and thus from the trajectory is not desirable for kinematic laser scanning, where any unresolved vibration translates into trajectory errors that in turn cause errors in the georeferenced point cloud.

In Wei and Li (2022), the authors investigate effects of GNSS and IMU sampling rate on the trajectory estimated via GNSS/IMU Kalman filter, and find that a reduced IMU sampling rate has a measurable but small negative impact on the positioning accuracy. However, orientation accuracy is not evaluated, which is arguably more important in the context of kinematic laser scanning. Vibration-induced errors in the orientation or boresight are addressed in Ma and Wu (2012), who simulate and analyze the impact of oscillating orientation errors on the point cloud. Schlager et al. (2022) investigate an automotive LiDAR sensor with an internal IMU and find that its 50 Hz sampling rate is insufficient if higher-frequency vibrations are present, which may occur in normal use-cases such as the vehicle traveling over rough asphalt or gravel. The magnitude of such errors depends on the

<sup>1</sup> In airborne laser scanning, *gyro-stabilization mounts* (GSM) are used to help keep the system leveled. In this case, the GNSS antenna is not rigid w.r.t. the IMU, but the relative orientation is measured by the GSM itself.

sensors used, the mounting of these sensors with respect to the carrier platform, the motion of the carrier platform and the frequency of the vibrations. As such, the values reported in the specific studies are hard to generalize. Nevertheless, different studies have analyzed vibration-related effects and found them to be of significance for kinematic LiDAR applications. However, little literature is available analyzing the effects of vibrations for commonly-used survey-grade laser scanning systems, in particular airborne laser scanning with uncrewed aerial vehicles (UAVs) or crewed aircraft.

In this article, we attempt to close these gaps and analyze platform dynamics of two airborne laser scanning platforms, a remotely piloted quadcopter and a crewed fixed-wing survey aircraft, with the goal of quantifying and mitigating the impact of vibrations on point cloud precision. Our analysis shows that in the presence of strong vibrations, a sampling rate of 200 Hz or less leads to reduced precision of the point cloud. However, a higher sampling rate would imply significantly higher computational effort for the trajectory processing. This is especially the case for tightly-coupled IMU/LiDAR integration approaches which have become popular both in robotics and in photogrammetry. For this reason, we suggest the use of downsampled IMU measurements within the NLS adjustment but to recover the high-frequency trajectory afterwards, ensuring that all vibrations present in the original high-rate inertial measurements are properly resolved but at the same time computational effort remains reasonable.

## 1.2. Downsampling and pre-integration

Downsampling is a common way to reduce computational load in strap-down inertial navigation by reducing the measurement rate. For GNSS/IMU integration based on Kalman filtering, it is common to run the processing at two distinct temporal resolutions (Al-Jailaty and Mansour, 2021): At a high rate, inertial measurements are read from the underlying sensor, suitably aggregated, and passed on to the filter itself<sup>2</sup>. Then, the filter itself runs at a lower rate using these aggregated (in other words, downsampled) measurements. Due to the inherent nonlinearity of the strap-down navigation equations, care has to be taken to avoid introducing errors through this process. Savage (1998a,b) provide equations for coning and sculling corrections, which have been widely used since. More recently, the similar concept of IMU pre-integration (Lupton and Sukkarieh, 2009) has found widespread use in robotics (Le Gentil and Vidal-Calleja, 2021). Here, the non-linear manifold structure of the inertial measurements is taken into account when performing the pre-integration. In the context of simultaneous localization and mapping (SLAM), the IMU measurements are pre-integrated in-between poses, thus reducing the number of parameters and observations in the SLAM problem and therefore drastically increasing computational efficiency. The use of IMU pre-integration for dynamic networks (Rouzaud and Skaloud, 2011), a type of least-squares adjustment, is investigated in Cucci and Skaloud (2019). There, performance benefits and also better numerical stability for the pre-integrated measurements are confirmed through simulation. However, accuracy deteriorates as more measurements are pre-integrated into one: the frequency of the pre-integrated measurements decreases and the IMU biases become harder to estimate accurately. Because IMU pre-integration inherently reduces the measurement rate, the issue of aliasing or signal cut-off, discussed in Section 1.1, is also relevant here. If higher-frequency components are present in the actual platform motion, these cannot be recovered from lower-frequency IMU measurements (whether downsampled or pre-integrated) and cannot be represented by a lower-frequency trajectory.

To address these two issues related to IMU downsampling,

## 1. unresolved high-frequency signal components and

<sup>2</sup> This is often realized inside the sensor hardware, with the high-frequency measurements inaccessible to the user.

## 2. decreased accuracy due to difficulties estimating biases,

we propose a two-stage strategy for incorporating downsampled inertial measurements in an NLS adjustment-based trajectory estimation. First, the trajectory is modeled as a combination of high-frequency and low-frequency parts, where the low-frequency part is estimated as part of the NLS adjustment, and the high-frequency part is reconstructed directly from the high-frequency measurements. Then, to improve the accuracy of the estimated IMU biases and to allow for modeling additional types of IMU errors (especially scale factors), we iterate the downsampling of the IMU measurements as better estimates of the IMU errors become available. In this way, we avoid having to propagate IMU errors and noise through the measurement equations. Specifically, we do not modify any measurement equations, which means the proposed strategy can be realized on top of existing NLS-based GNSS/IMU/LiDAR-integration algorithms. However, it makes use of the original high-frequency measurements and is thus not applicable if the downsampling is already performed within the inertial sensor during data acquisition. While this strategy is simple, it is shown to be effective, as it is both efficient and precise. This is demonstrated on real-world data.

## 1.3. Main contributions

In summary, the main contributions of this article are (1) analysis of platform dynamics for two representative airborne laser scanning platforms based on in-flight inertial measurements together with a heuristic quantification of errors caused by vibrations with frequencies exceeding the IMU measurements' Nyquist rate, and (2) a strategy for efficiently processing high-rate IMU measurements in the context of trajectory estimation via batch non-linear least-squares (NLS) adjustment. The methodology is demonstrated and evaluated in detail for an uncrewed laser scanning (ULS) and a (crewed) airborne laser scanning (ALS) dataset.

Background and theoretical considerations for the analysis of the inertial measurements and platform dynamics are given in Section 2. The methodology for the downsampling of the IMU measurements, and the strategy for integrating them in the NLS trajectory estimation is described in Section 3. Then, real-world data is processed and thoroughly analyzed in Section 4. Lastly, Section 5 discusses the results and gives an outlook on further investigations.

## 2. Trajectory dynamics

Before describing the downsampling and sensor fusion methodology, we provide a motivating example and introduce simple tools for analyzing the dynamics that the platform and sensor assembly are subjected to. In short, we examine the high-rate inertial measurements in the frequency domain in order to determine the presence of vibrations and to predict the magnitude of trajectory errors resulting from aliasing or signal-cut-off.

### 2.1. Motivating example and theoretical background

Consider as an example a sinusoidal signal made up of three components

$$s(t) = 0.25 \sin(50 \cdot 2\pi t) \quad (2.1)$$

$$+ 0.1 \sin(120 \cdot 2\pi t) \quad (2.2)$$

$$+ 0.01 \sin(310 \cdot 2\pi t). \quad (2.3)$$

The ideal signal is sampled at 800 Hz with random white measurement noise (Fig. 2). Additionally, the 800 Hz signal is decimated by a factor of 4, simply by averaging 4 consecutive samples. Lastly, the 800 Hz signal is first passed through an anti-aliasing (low-pass) filter, and then decimated as before. This example aims to demonstrate the




Fig. 2. Example signal with frequency components 50 Hz, 120 Hz, and 310 Hz. The ideal signal is tainted by additive white noise and sampled at 800 Hz, which is then decimated with and without anti-aliasing to 200 Hz.




Fig. 3. Spectrum of the example signal  $s$ , idealized (black), obtained by discrete Fourier transform from the 800 Hz signal (blue) and from the decimated 200 Hz signals with (orange) and without (red) anti-aliasing. The y-axis shows the linear magnitude  $|A(v)|$  for the complex Fourier coefficient  $A(v) \in \mathbb{C}$  at frequency  $v$ .

effect of the downsampling when signal frequencies larger than the Nyquist rate are present. This somewhat mirrors the internal processing of IMUs, albeit only with linear averaging and for one axis, thereby disregarding the topic of coning and sculling (Savage, 1998a,b).

Fig. 3 shows the spectrum of the example signal. Both the theoretical/ideal spectrum and the actual 800 Hz signal spectrum show the three frequency peaks and the white noise. The 200 Hz decimated signal spectrum is cut off at 100 Hz, half the sampling rate. The 120 Hz and 310 Hz signal are aliased and now appear at 80 Hz and 90 Hz, respectively. For the anti-aliased signal, which is first low-pass filtered and then decimated, the incorrect 90 Hz peak disappears completely. Note that the incorrect 80 Hz is still visible but with reduced amplitude; in practice the filter response is never a perfect ‘brick wall’.

As the exact mode of rate-reduction for specific IMUs differs or is not even disclosed to the user, we cannot exactly reproduce the inner workings of any given IMU. However, we expect that some form of anti-aliasing (and, for higher-end IMUs possibly also coning and sculling corrections) are applied, even if the details (e.g., filter characteristics) are unknown. Thus for our analysis, we focus on the case of a downsampled but anti-aliased signal, no matter whether the downsampling is performed in real-time on-board the sensor or in post-processing (similar to pre-integration).

## 2.2. A-prior error estimates

To obtain a heuristic for the expected trajectory error due to unresolved vibrations, consider a signal  $s(t)$  with period  $T$  on the interval

$[0, T]$ , reconstructed from  $v_s = 800$  Hz samples. We write  $s(t)$  as its Fourier sum and split it into low-frequency and high-frequency parts

$$\begin{aligned} s(t) &= \sum_{|v| \leq \frac{v_c}{2}} A(v) e^{2\pi i v t} \\ &= \sum_{|v| \leq \frac{v_c}{2}} A(v) e^{2\pi i v t} + \sum_{\frac{v_c}{2} < |v| \leq \frac{v_c}{2}} A(v) e^{2\pi i v t}, \end{aligned} \quad (2.4)$$

where  $v = \frac{k}{T}$ ,  $k = 0, \pm 1, \pm 2, \dots$  and  $v_c$  is the reduced sampling frequency. If the original signal  $s(t)$  is anti-aliased with an ideal low-pass filter and decimated to  $[v_c]$  Hz, the low-frequency signal’s spectrum is cut off at  $[\frac{v_c}{2}]$  Hz, i.e.,

$$s_{\text{LF}}(t) = \sum_{|v| \leq \frac{v_c}{2}} A(v) e^{2\pi i v t}. \quad (2.5)$$

Consider now the angular velocity measurements  $w(t)$  of a single gyroscope as the signal. Assuming rotation is only present around that axis, the angle  $\theta$  is simply the integral of the angular velocity  $\theta(t) = \int_0^t \omega(\tau) d\tau$ . Thus, the root mean square (RMS) error in the interval  $[0, T]$  of  $\theta_{\text{LF}}$ , incurred by the lower sampling rate, is

$$\text{RMS}(\theta - \theta_{\text{LF}}) \quad (2.6)$$

$$= \sqrt{\frac{1}{T} \int_0^T \left( \int_0^t \omega(\tau) - \omega_{\text{LF}}(\tau) d\tau \right)^2 dt} \quad (2.7)$$

$$= \sqrt{\frac{1}{T} \int_0^T \left( \sum_{|v| > \frac{v_c}{2}} \int_0^t A_{\omega}(v) e^{2\pi i v \tau} d\tau \right)^2 dt} \quad (2.8)$$

$$= \sqrt{\sum_{|v| > \frac{v_c}{2}} \left( \frac{|A_{\omega}(v)|}{2\pi v} \right)^2}, \quad (2.9)$$

which can be computed from the Fourier coefficients  $A_{\omega}(v)$  of  $\omega(t)$ . It can be seen directly from Eq. (2.9) that for the same signal amplitude, the orientation error is proportional to the inverse of the respective frequency.

Note that when the vibrations are concentrated in one frequency, the peak error is related to the RMS error by

$$\max_{t \in [0, T]} |\theta(t) - \theta_{\text{LF}}(t)| = \sqrt{2} \text{RMS}(\theta - \theta_{\text{LF}}). \quad (2.10)$$

A similar metric can be derived for the position, integrating the accelerometer signal  $f(t)$  twice to obtain the position  $p(t) = \iint_0^t f(\tau) d\tau$ . The corresponding RMS error of the position  $p_{\text{LF}}$  derived from downsampled accelerations  $f_{\text{LF}}$  is then given by

$$\text{RMS}(p - p_{\text{LF}}) = \sqrt{\sum_{|v| > \frac{v_c}{2}} \left( \frac{|A_f(v)|}{(2\pi v)^2} \right)^2}. \quad (2.11)$$

In this case, the error for a given amplitude is proportional to the squared inverse of the respective frequency.

Eqs. (2.9) and (2.11) allow us to make a heuristic a-priori prediction of the trajectory errors induced by cutting off the high-frequency signal components (i.e., vibrations) in the case of downsampling with anti-aliasing. This prediction, based only on the simplified 1D case of strap-down inertial navigation, is independent of the actual sensor fusion methodology. In our specific case, it will be experimentally verified by comparing it with the results of NLS-based trajectory estimation.

## 3. Trajectory estimation

The trajectory estimation approach used here, based on non-linear least-square estimation, is described in detail in Pöppl et al. (2024). In the following, a short overview of this methodology is given, followed by a more detailed description of the downsampling itself and the integration of downsampled inertial measurements in the trajectory estimation.



### 3.1. Non-linear least-squares trajectory estimation

At the core, NLS estimation involves finding unknown parameters  $\mathbf{x}$  by minimizing the difference between the functional model  $\mathbf{y}_i := f_i(\mathbf{x})$  and observations  $\tilde{\mathbf{y}}_i$  of the true values  $\mathbf{y}_i$

$$\begin{aligned}\tilde{\mathbf{y}}_i &= \mathbf{y}_i + \epsilon_i \\ &= f_i(\mathbf{x}) + \epsilon_i, \\ &\text{with } \epsilon_i \sim N(\mathbf{0}, \Sigma_i), \quad 1 \leq i \leq n,\end{aligned}\quad (3.1)$$

where the measurements  $\tilde{\mathbf{y}}_i$  are tainted by the additive zero-mean Gaussian noise  $\epsilon_i$ .

The NLS estimate  $\mathbf{x}^*$  of the parameters  $\mathbf{x}$  is then obtained by minimizing the sum of squared residuals

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} \sum_{i=0}^n \sum_{j=0}^{m_i} \rho(r_{i,j}^2), \quad (3.2)$$

$$\text{with } r_i := \Sigma_i^{-\frac{1}{2}} (\tilde{\mathbf{y}}_i - f_i(\mathbf{x})) \in \mathbb{R}^{m_i}, \quad (3.3)$$

where  $\rho : [0, \infty) \rightarrow [0, \infty)$  is a suitable loss function. In our case, this is either the standard linear loss  $\rho(r) = r$  or the robust Huber loss (Huber, 1964).

The residuals include four main types:

- Priors, which encode a-priori information about the stochastic properties of some parameters. IMU bias and scale factors are modeled stochastically as Gauss–Markov processes and functionally as time-varying (using linear splines).
- IMU gyroscope and accelerometer observations, derived from the standard strap-down inertial navigation equations.
- GNSS position observations describing the discrepancy between the modeled position of the GNSS antenna and the measured GNSS position as provided by external GNSS processing. Due to the presence of outliers, the Huber loss is used for the GNSS residuals.
- LiDAR observations describing the discrepancies between multiple scans of the same object. In our case, planes are extracted from the point clouds. The observed planes are then matched and merged into one modeled plane, for which three parameters (offset and two slopes) are explicitly estimated. LiDAR boresight misalignment can also be modeled and estimated if necessary. As the matching of corresponding planes may fail causing gross errors, the Huber loss is also used for the LiDAR residuals.

All GNSS, IMU and LiDAR observations depend on position and orientation at specific points in time. The trajectory itself is modeled using splines, with spline nodes set for each IMU sample time. The IMU measurement equations are discussed further in Section 3.2. The other measurement equations and the LiDAR plane extraction are the same as in Pöpl et al. (2024), only the representation of the trajectory differs. Here, Hermite splines are used instead of B-Splines for modeling the trajectory (see Section 3.4). In contrast to B-Splines, the Hermite spline coefficients are equivalent to the spline function value at the corresponding node, thus simplifying the construction of interpolatory splines for given function values.

Since the least-squares optimization is non-linear, suitable starting values for the parameters  $\mathbf{x}$ , especially the trajectory, are required. For this reason, the optimization is first performed incrementally in a sliding-window mode, with a window of length 60 s. The full processing workflow then includes the following steps (Fig. 4):

1. Sliding-window NLS optimization with GNSS and IMU observations to obtain an initial trajectory.
2. Batch optimization with GNSS and IMU observations, starting from the sliding-window trajectory.
3. Preliminary georeferencing of the LiDAR point cloud with the GNSS/IMU trajectory.

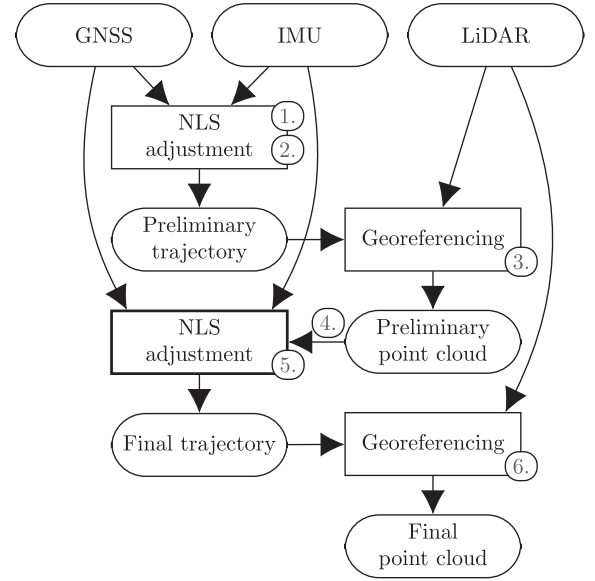


Fig. 4. Trajectory estimation workflow from Pöpl et al. (2024).

4. Extraction and matching of planar features from the preliminary LiDAR point cloud.
5. Batch optimization with GNSS, IMU and LiDAR observations.
6. Final georeferencing of the LiDAR point cloud with the GNSS/IMU/LiDAR trajectory.

### 3.2. Inertial measurement equations

The inertial sensors, accelerometer and gyroscope, measure respectively specific force and angular velocity of the body frame (i.e., the sensor assembly) with respect to an inertial frame. The trajectory itself is modeled with respect to an earth-fixed frame, and the point cloud is then derived from scanner measurements and the trajectory, and therefore also referenced to the earth-fixed frame. Specifically, the three Cartesian coordinate systems relevant for the inertial measurements are

1. the  $e$ -system, the standard earth-centered earth-fixed (ECEF) coordinate system,
2. the  $i$ -system, the standard earth-centered inertial (ECI) coordinate system, which rotates with respect to the  $e$ -system around their common  $z$ -axis, and
3. the  $b$ -system, with origin in the IMU center and axes derived from the carrier platform (front, right, down).

We keep with the notation of Pöpl et al. (2024), which itself follows Groves (2013):  $\mathbf{x}_{\beta\alpha}^{\gamma}$  refers to a property of frame  $\alpha$  w.r.t. frame  $\beta$  resolved in frame  $\gamma$ , and  $\mathbf{R}_{\alpha}^{\beta}$  denotes the rotation from  $\alpha$ -frame to the  $\beta$ -frame. For example,  $\mathbf{x}_{eb}^e$  and  $\dot{\mathbf{x}}_{eb}^e$  denote position and velocity of the  $b$ -system with respect to and resolved in the  $e$ -system, whereas  $\mathbf{R}_e^b$  is the rotation from  $e$ -system to  $b$ -system.

The IMU outputs measurements at times  $t_k \in [0, T]$ ,  $1 \leq k \leq n$ . We interpret these measurements as incremental, i.e.,  $\Delta \mathbf{v}$  and  $\Delta \theta$  measurements, which correspond to the increment in velocity and orientation between two sample times  $t_{k-1}$  and  $t_k = t_{k-1} + \Delta t_k$

$$\Delta \theta_{ib}^b(t_{k-1}, t_k) = \log(\mathbf{R}_i^b(t_{k-1}) \mathbf{R}_b^i(t_k)) \quad (3.4)$$

$$= \log\left(\mathbf{R}_e^b(t_{k-1}) \exp(\omega_{ie}^e \Delta t_k) \mathbf{R}_b^e(t_k)\right), \quad (3.5)$$

$$\Delta \mathbf{v}_{ib}^b(t_{k-1}, t_k) = \mathbf{R}_i^b(t_{k-1}) \left( \dot{\mathbf{x}}_{ib}^i(t_k) - \dot{\mathbf{x}}_{ib}^i(t_{k-1}) - \Delta t_k \gamma_{ib}^i(t_{k-1}) \right) \quad (3.6)$$




Fig. 5. Illustration of the high-rate and low-rate sample intervals and corresponding increments.

$$= \mathbf{R}_e^b(t_{k-1}) \left( \dot{\mathbf{x}}_{eb}^e(t_k) - \dot{\mathbf{x}}_{eb}^e(t_{k-1}) + \Delta t_k \left( [\omega_{ie}^e]^2 \mathbf{x}_{eb}^e(t_{k-1}) + 2 [\omega_{ie}^e]_{\wedge} \dot{\mathbf{x}}_{eb}^e(t_{k-1}) - \gamma_{ib}^e(t_{k-1}) \right) \right), \quad (3.7)$$

where  $\omega_{ie}^e$  is the earth's angular velocity and  $\gamma_{ib}^e$  is the acceleration due to gravitational force<sup>3</sup>. The functions log and exp are the SO(3) Lie-group logarithm and exponential, which map between a 3D axis-angle vector and the corresponding rotation matrix, with  $\mathbb{R}^{3 \times 3}$  interpreted as subset of SO(3). The hat operator  $[\cdot]_{\wedge}$  maps a vector to the skew symmetric cross-product matrix, so that  $[\mathbf{x}]_{\wedge} \mathbf{y} = \mathbf{x} \times \mathbf{y}$ .

The inertial measurement equations are simpler when specified with respect to the  $i$ -system (Eqs. (3.4) and (3.6)), but are reformulated here with respect to the  $e$ -system (Eqs. (3.5) and (3.7)) as this is what both trajectory and point cloud are referenced to.

### 3.3. Downsampling of the inertial measurements

The strategy with which inertial measurements are downsampled (in other words, *aggregated*) can be seen as a simple averaging, but averaging in a way that respects the Lie-group structure of SO(3), and takes into account the dependence of acceleration on the current orientation. This is essentially the same as used in the on-manifold pre-integration strategy introduced in Forster et al. (2015), except that we do not aim to propagate IMU noise and biases through the integration. Also, we do not *pre-integrate* in-between specific keyframes, but rather always aggregate  $d$  subsequent samples, thereby effectively reducing the sampling frequency from  $\nu_s$  to  $\nu_c = \frac{\nu_s}{d}$  and the number of samples from  $n$  to  $\lfloor \frac{n}{d} \rfloor$ .

For the original 'high-frequency' measurements, the lower-case index  $k$  is used, while the downsampled 'lower-frequency' measurements are indexed with upper-case  $K$

$$k \in \{1, \dots, n\}, K \in \{1, \dots, N\}, N := \lfloor \frac{n}{d} \rfloor. \quad (3.8)$$

To simplify notation, the coordinate systems are omitted as they are always the same, and the index now refers to the sample interval

$$\Delta \theta_k := \Delta \theta_{ib}^b(t_{k-1}, t_k), \quad \Delta \mathbf{v}_k := \Delta \mathbf{v}_{ib}^b(t_{k-1}, t_k), \quad \text{and} \quad (3.9)$$

$$\Delta \theta^K := \Delta \theta_{ib}^b(t_{d(K-1)}, t_{dK}), \quad \Delta \mathbf{v}^K := \Delta \mathbf{v}_{ib}^b(t_{d(K-1)}, t_{dK}). \quad (3.10)$$

These  $\Delta \theta$  and  $\Delta \mathbf{v}$  values will later take the role of the measurements  $\mathbf{y}$  in the NLS adjustment (Eqs. (3.1) and (3.2)). The aggregated low-frequency incremental measurements for the sample times  $t^{K-1} := t_{d(K-1)}$  and  $t^K := t_{dK} = t^{K-1} + \Delta t^K$  (see Fig. 5) are derived

from the high-frequency incremental measurements by substituting the low-frequency sample intervals  $[t^{K-1}, t^K]$  into Eqs. (3.4) and (3.6), yielding

$$\Delta \theta^K \stackrel{(3.4)}{=} \log \left( \mathbf{R}_i^b(t_{d(K-1)}) \mathbf{R}_b^i(t_{dK}) \right) \quad (3.11)$$

$$= \log \prod_{j=1}^d \mathbf{R}_i^b(t_{d(K-1)+j-1}) \mathbf{R}_b^i(t_{d(K-1)+j}) \quad (3.12)$$

$$\stackrel{(3.4)}{=} \log \prod_{j=1}^d \exp \left( \Delta \theta_{d(K-1)+j} \right), \quad (3.13)$$

$$\Delta \mathbf{v}^K \stackrel{(3.6)}{=} \mathbf{R}_i^b(t_{d(K-1)}) \left( \dot{\mathbf{x}}_{ib}^i(t_{dK}) - \right. \quad (3.14)$$

$$\left. \dot{\mathbf{x}}_{ib}^i(t_{d(K-1)}) - \Delta t^K \gamma_{ib}^i(t_{d(K-1)}) \right) \approx \mathbf{R}_i^b(t_{d(K-1)}) \sum_{j=1}^d \left( \dot{\mathbf{x}}_{ib}^i(t_{d(K-1)+j}) - \right. \quad (3.15)$$

$$\left. \dot{\mathbf{x}}_{ib}^i(t_{d(K-1)+j-1}) - \Delta t_{d(K-1)+j} \gamma_{ib}^i(t_{d(K-1)+j-1}) \right) \stackrel{(3.6)}{=} \sum_{j=1}^d \prod_{\ell=1}^{j-1} \exp \left( \Delta \theta_{d(K-1)+\ell} \right) \Delta \mathbf{v}_{d(K-1)+j} \quad (3.16)$$

$$\mathbf{R}_i^b(t_{d(K-1)}) \mathbf{R}_b^i(t_{d(K-1)+j-1})$$

The approximation in Eq. (3.16) is due to using the gravitational acceleration  $\gamma_{ib}^i$  at times  $t_{d(K-1)+j-1}, t_{d(K-1)+j}, \dots$  instead of  $\gamma_{ib}^i(t_{d(K-1)})$ . Since the gravitational acceleration varies slowly, this is not an issue in practice.

### 3.4. Trajectory upsampling

The aggregated low-frequency measurements  $\Delta \theta^K$  and  $\Delta \mathbf{v}^K$  correspond precisely to the same underlying measurement equations as the high-frequency measurements  $\Delta \theta_k$  and  $\Delta \mathbf{v}_k$  in Eqs. (3.4) and (3.6), except for larger time intervals. Thus, if we were interested only in determining the position at those large time intervals and for the moment disregarded IMU errors, there would be no inherent loss of information. However, as argued in Section 2, estimating the trajectory only at a lower frequency means disregarding motion present at higher frequencies (i.e., vibrations). In the context of LiDAR georeferencing, the resulting errors in the trajectory propagate to the LiDAR point cloud. Because the LiDAR measurements are made at a rate far higher than the IMU sampling rate, it is in any case necessary to interpolate the trajectory for georeferencing. The argument regarding frequency cut-off can also be seen from the viewpoint of interpolation, as the interpolation error grows with the distance between the data points. Here, we address this issue by recovering the high-frequency trajectory components from a low-frequency trajectory using the original high-frequency IMU measurements. This may be interpreted as an interpolation method which is aided by the IMU measurements themselves in order to improve accuracy within the larger sample intervals.

Assuming that the trajectory estimation (Section 3.1) was performed with downsampled IMU measurements, estimates of position and orientation at each low-frequency sample time  $t_{dK} = t^K$

$$\mathbf{x}_{dK}, \mathbf{R}_{dK}, \quad K = 1, \dots, N, \quad (3.17)$$

are available in addition to the original IMU measurements at the high-frequency sample times

$$\Delta \theta_k, \Delta \mathbf{v}_k, \quad k = 1, \dots, n. \quad (3.18)$$

The aim is to recover the high-frequency position and orientation for each sample time  $t_k$

$$\mathbf{x}_k, \mathbf{R}_k, \quad k = 1, \dots, n. \quad (3.19)$$

<sup>3</sup> Gravitational acceleration is derived using the ellipsoidal gravity model from Groves (2013, Ch.2).

To simplify notation, the coordinate systems are not explicitly specified, but both position and orientation refer to the body coordinate system (*b*-system) with respect to the earth-fixed coordinate system (*e*-system).

For the low-frequency sample times  $t_{dK}, K = 1, \dots, N$ , the position and orientation are already known. For these sample times, we define the following auxiliary variables

$$\begin{aligned} \mathbf{x}_{dK}^+ &:= \mathbf{x}_{dK}^- := \mathbf{x}_{dK}, \\ \mathbf{R}_{dK}^+ &:= \mathbf{R}_{dK}^- := \mathbf{R}_{dK}, \end{aligned} \quad (3.20)$$

which will allow defining recursive equations for the forward- and backward-integrated position and orientation values within each  $[t_{d(K-1)}, t_{dK}]$  interval. For every  $K = 1, \dots, N-1$  and  $j = 1, \dots, d-1$  the corresponding high-frequency index is  $k = dK + j$ . The high-frequency trajectory for each  $k$  is reconstructed from a convex combination of the forward-integrated and the backward-integrated high-frequency inertial measurements.

For the orientation, this leads to

$$\begin{aligned} \mathbf{R}_k^+ &:= \exp(-\Delta\theta_k) \mathbf{R}_{k-1}^+ \exp(\Delta t_k \omega_{ie}^e), \\ \mathbf{R}_k^- &:= \exp(\Delta\theta_{k+1}) \mathbf{R}_{k+1}^- \exp(-\Delta t_{k+1} \omega_{ie}^e) \end{aligned} \quad (3.21)$$

$$\mathbf{R}_k := \mathbf{R}_k^+ \exp\left(\frac{j}{d} \log((\mathbf{R}_k^+)^T \mathbf{R}_k^-)\right), \quad (3.22)$$

where  $k = dK + j$ ,  $1 \leq K < N$  and  $0 < j < d$ .

Note that this fills the gaps between each  $t_{dK}$  and  $t_{d(K+1)}$ , but does not change the orientation at the  $[t_{dK}, t_{d(K+1)}]$  interval endpoints.

For the position, the procedure is similar but involves integrating twice to obtain first the velocity and then the position, forwards and backwards. The rotation computed previously is used to reference each measurement to the correct time. This yields

$$\begin{aligned} \mathbf{a}_k^+ &:= \gamma_{k-1} - [\omega_{ie}^e]^2 \mathbf{x}_{k-1}^+ - 2[\omega_{ie}^e]_{\wedge} \mathbf{v}_{k-1}^+, \\ \mathbf{v}_k^+ &:= \mathbf{v}_{k-1}^+ + \mathbf{R}_{k-1}^T \Delta \mathbf{v}_k + \Delta t_k \mathbf{a}_k^+, \\ \mathbf{x}_k^+ &:= \mathbf{x}_{k-1}^+ + \frac{\Delta t_k}{2} (\mathbf{v}_k^+ + \mathbf{v}_{k-1}^+), \end{aligned} \quad (3.23)$$

for the forward pass and

$$\begin{aligned} \mathbf{a}_k^- &:= \gamma_{k+1} - [\omega_{ie}^e]^2 \mathbf{x}_{k+1}^- - 2[\omega_{ie}^e]_{\wedge} \mathbf{v}_{k+1}^-, \\ \mathbf{v}_k^- &:= \mathbf{v}_{k+1}^- - \mathbf{R}_k^T \Delta \mathbf{v}_{k+1} - \Delta t_{k+1} \mathbf{a}_k^-, \\ \mathbf{x}_k^- &:= \mathbf{x}_{k+1}^- - \frac{\Delta t_{k+1}}{2} (\mathbf{v}_k^- + \mathbf{v}_{k+1}^-), \end{aligned} \quad (3.24)$$

for the backwards pass, where  $\gamma_k = \gamma_{ib}^e(t_k)$ . Finally, the high-frequency positions are again determined by convex combination of the forward and backward integrals

$$\mathbf{x}_k := \mathbf{x}_k^+ + \frac{j}{d} (\mathbf{x}_k^- - \mathbf{x}_k^+), \quad (3.25)$$

where  $k = dK + j$ ,  $1 \leq K < N$  and  $0 < j < d$ .

### 3.5. Implicit modeling of the high-frequency trajectory

The above procedure allows to perform trajectory estimation with downsampled IMU measurements, yet recover the high-frequency trajectory for the purpose of georeferencing. For the NLS trajectory estimation, the GNSS and LiDAR measurements must be referenced to the high-frequency trajectory. To achieve this without introducing additional parameters into the NLS adjustment, we split the trajectory into a low-frequency and a high-frequency part (see also Fig. 6)

$$\begin{aligned} \mathbf{R}(t) &:= \mathbf{R}_{LF}(t) \cdot \mathbf{R}_{HF}(t), \\ \mathbf{x}(t) &:= \mathbf{x}_{LF}(t) + \mathbf{x}_{HF}(t), \end{aligned} \quad (3.26)$$

where  $\forall K = 1 \dots N$ ,

$$\begin{aligned} \mathbf{R}(t_{dK}) &= \mathbf{R}_{LF}(t_{dK}), \\ \mathbf{x}(t_{dK}) &= \mathbf{x}_{LF}(t_{dK}), \end{aligned} \quad (3.27)$$




Fig. 6. Illustration of the decomposition into low-frequency and high-frequency parts. Note that the high-frequency part is zero at all low-frequency sample times  $t^K, K = 1, \dots, N$ .

and only estimate the former within the adjustment.

Splines are used to represent the continuous-time position and orientation. Specifically, we use cubic quaternion Hermite splines (Kim et al., 1995) for orientation and quintic Euclidean Hermite splines for position, which are once and twice continuously differentiable, respectively. The splines are defined by their nodal values

$$\begin{aligned} \mathbf{R}_{LF}(t_{dK}) &= \mathbf{R}_{dK}, \\ \mathbf{x}_{LF}(t_{dK}) &= \mathbf{x}_{dK}, \text{ and} \end{aligned} \quad (3.28)$$

$$\begin{aligned} \mathbf{R}_{HF}(t_k) &= \mathbf{R}_{LF}^T(t_k) \mathbf{R}_k, \\ \mathbf{x}_{HF}(t_k) &= \mathbf{x}_k - \mathbf{x}_{LF}(t_k). \end{aligned} \quad (3.29)$$

Each spline's nodes are placed at the respective sample times  $t_k$  and  $t_{dK}$ . Hermite splines are used because their spline coefficients are simply their nodal values, whereas recovering a quadratic or cubic B-spline's coefficients from nodal values requires solving large equation systems with at least  $n$  parameters, which is exactly what we aim to avoid in the first place.

Within the NLS adjustment, only the  $N$  sets of coefficients of the low-frequency splines  $\mathbf{x}_{LF}$  and  $\mathbf{R}_{LF}$  are optimized. Note that regardless of the downsampling factor  $d$ , both parameters and observations are reduced accordingly and the redundancy of the adjustment remains the same.

### 3.6. Iterative estimation of IMU errors

In the previous sections, no attention was paid to errors in the IMU measurements. The IMU error model (Pöpl et al., 2024) is standard: the measured values  $\Delta \tilde{\mathbf{v}}, \Delta \tilde{\theta}$  are tainted by additive white noise  $\epsilon_*$ , biases  $\mathbf{b}_*$  and scale factors  $\mathbf{S}_* = \text{diag}(s_{x*}, s_{y*}, s_{z*})$ . The measurement model for the velocity and orientation increments for an interval of length  $\Delta t$  is

$$\begin{aligned} \Delta \tilde{\mathbf{v}} \frac{1}{\Delta t} &= (\mathbf{I} + \mathbf{S}_v) \Delta \mathbf{v} \frac{1}{\Delta t} + \mathbf{b}_v(t) + \epsilon_v, \\ \Delta \tilde{\theta} \frac{1}{\Delta t} &= (\mathbf{I} + \mathbf{S}_\theta) \Delta \theta \frac{1}{\Delta t} + \mathbf{b}_\theta(t) + \epsilon_\theta. \end{aligned} \quad (3.30)$$

When aggregating the IMU measurements according to Eqs. (3.11) and (3.14), the presence of bias, scale factor and random errors will cause errors in the aggregated measurements and in the upsampled trajectory. Such errors cannot be precisely modeled in the same way after aggregation due to the non-linearities involved. One option to mitigate this problem would be to move the computation of the aggregated measurements into the adjustment, but that is computationally expensive. Often, the contribution of the biases is considered in the adjustment through linearization (Forster et al., 2015; Cucci and Skaland, 2019). Both approaches require modifying the observation equations. We take a different approach, and re-compute both the upsampled trajectory and the downsampled measurements iteratively. This process is easily parallelizable and in the context of batch NLS adjustment,

where all parameters are estimated simultaneously, the computational effort required to solve the linear equation system is much larger. The stochastic model for the IMU noise  $\epsilon_*$  remains unchanged (see Pöppl et al. (2024)); the error variance is simply derived from the a-priori determined noise densities (Farrell et al., 2022), i.e., scaled to account for the length of the time interval  $\Delta t$ .

The NLS adjustment problem is inherently iterative, and is performed until convergence or for a total number of  $M$  iterations. However, every  $\sqrt{M}$  iterations or earlier in case of convergence, it is stopped. At this point, an estimate of the IMU biases and scale factors is available. This is used to correct the high-frequency IMU measurements and then both the downsampled low-frequency measurements and the upsampled high-frequency trajectory are re-generated from the corrected measurements. Note that the IMU bias and scale factors need to be re-applied on top of the new downsampled measurements, as otherwise the stochastic model for these errors would be invalid. The NLS iterations are then restarted again. This process can be seen as a sort of fixed point iteration which, ideally, converges to the actual bias and scale factor values. Practically, the iteration is stopped at the maximum of  $M$  total NLS iterations or if the change in the re-generated downsampled IMU measurements is below a certain threshold.

The full estimation procedure is thus made up of two nested iterations: The outer iteration, termed *D/U iteration*, re-generates the downsampled IMU measurements and the upsampled trajectory. The inner iterations are standard *NLS iterations*, performed within each D/U iteration in order to obtain the NLS estimate Eq. (3.2) given the currently available downsampled IMU measurements and upsampled trajectory. The individual steps are:

1. Starting from an initial trajectory estimate, compute the corresponding high-frequency trajectory using the original IMU measurements  $\widehat{\Delta\theta}_k, \widehat{\Delta v}_k$  and compute the downsampled measurements  $\widehat{\Delta\theta}^K, \widehat{\Delta v}^K$ .
2. Run the NLS optimization with downsampled IMU observations, GNSS position observations and (optionally) LiDAR plane observations to obtain error estimates  $b_*, s_*$  for the downsampled IMU measurements, i.e.,

$$\begin{aligned} \widehat{\Delta\theta}^K \frac{1}{\Delta t} &= \text{diag}(1 + s_\theta) \Delta\theta^K \frac{1}{\Delta t} + b_\theta + \epsilon_\theta, \\ \widehat{\Delta v}^K \frac{1}{\Delta t} &= \text{diag}(1 + s_v) \Delta v^K \frac{1}{\Delta t} + b_v + \epsilon_v. \end{aligned} \quad (3.31)$$

The NLS iterations are stopped after  $\lfloor \sqrt{M} \rfloor$  iterations or in case of convergence, where convergence is assumed if the RMS change of optimized trajectory between iterations is below a given threshold.

3. Correct the original high-frequency IMU measurements with the error estimates

$$\begin{aligned} \widehat{\Delta\theta}_k &\leftarrow \text{diag}(1 + s_\theta)^{-1} (\widehat{\Delta\theta}_k - b_\theta), \\ \widehat{\Delta v}_k &\leftarrow \text{diag}(1 + s_v)^{-1} (\widehat{\Delta v}_k - b_v). \end{aligned} \quad (3.32)$$

4. Using the corrected high-frequency IMU measurements, re-compute the upsampled trajectory and the downsampled IMU measurements  $\widehat{\Delta\theta}_k^K, \widehat{\Delta v}_k^K$  and re-apply the estimated errors to the downsampled measurements

$$\begin{aligned} \widehat{\Delta\theta}^K &\leftarrow \text{diag}(1 + s_\theta) \widehat{\Delta\theta}_k^K + b_\theta, \\ \widehat{\Delta v}^K &\leftarrow \text{diag}(1 + s_v) \widehat{\Delta v}_k^K + b_v. \end{aligned} \quad (3.33)$$

5. Run the NLS optimization with downsampled IMU observations, GNSS position observations and (optionally) LiDAR plane observations.

In the D/U iterations, steps 2–4 are repeated until the maximum absolute change in the re-generated downsampled measurements is below the given tolerance or the maximum number of NLS iterations  $M$  is reached. Practically,  $M = 16$  is used in which case both termination conditions are reached before the maximum number of iterations. As at least one U/D-iteration is always performed, this scheme generally increases the required number of NLS iterations and thus incurs a small runtime overhead (see Section 4.1.4).



Fig. 7. DJI M350 with a RIEGL miniVUX-3UAV laser scanner and integrated RiLOC-E navigation system.

## 4. Results

Remote sensing payloads are usually mounted with mechanical dampening components in order to reduce the impact of vibrations on the measurements, whether caused by the carrier platform or the scanning system itself. However, it is not physically possible to fully dampen all vibrations, especially for smaller UAV platforms where payload weight and size is limited. In airborne laser scanning with crewed aircraft, the laser scanning system is often mounted in a gyro stabilization mount which ensures consistent orientation and low dynamics of the payload. Such mounting is not always feasible, depending on the payload and aircraft used.

We analyze inertial data from two representative airborne platforms: a small quadcopter UAV with a maximum payload of 2.7 kg, and a crewed fixed-wing survey aircraft with the payload mounted in a pod affixed to the aircraft wing strut. All data processing is carried out on a Linux workstation with an AMD Ryzen 7 2700X CPU, 64 GiB RAM and an NVMe drive. The non-linear least-squares adjustment is implemented using the CERES solver (Agarwal et al., 2022).

### 4.1. Case study: UAV LiDAR

The first dataset is acquired with a quadcopter UAV (DJI Matrice M350) featuring as payload a RIEGL miniVUX-3UAV laser scanner integrated with a RIEGL RiLOC-E navigation system (Fig. 7). The latter comprises a low-SWaP<sup>4</sup> consumer-grade MEMS IMU, in this case configured with an output rate of 850 Hz, and a dual-frequency GNSS receiver with an output rate of 7 Hz. Traditional high-performance GNSS/IMU solutions used for UAV-based surveying provide orientation with an expected accuracy of approximately 30 mdeg. In comparison, the IMU used here is an entry-level option meant for tight-coupling with the laser data (Pöppl et al., 2023, 2024), and the expected accuracy of purely GNSS/IMU-derived orientation is up to an order of magnitude less accurate. However, evaluation of absolute georeferencing accuracy is not the goal of this work; rather, the aim is to investigate the impact of unresolved vibrations on the trajectory itself as well as the achieved precision (single-strip and multi-strip) of the georeferenced point cloud when using downsampled IMU measurements as described in Section 3.3.

The test dataset (Fig. 8) is an intentionally small data acquisition consisting of only six flight strips, flown in pairs with 100% overlap, at a flying height of 50 m above ground level (AGL) and with flight speed of 6  $\frac{m}{s}$ . The data acquisition is kept short (approx. 10 min) to keep computational effort and storage needs reasonable for evaluation with a large number of different processing runs and therefore point

<sup>4</sup> Size, weight and power (SWaP).





Fig. 8. ULS point cloud colored by reflectance (dark/bright = low/high reflectance, hillshade) and flight trajectory, superimposed onto map (© OpenStreetMap contributors). The X/Y-axes are those of a local east/north/up (ENU) Cartesian coordinate system. The point clouds from the two flight strips in the dashed white rectangle are later evaluated in detail.

clouds. The study area contains an industrial complex, which is chosen specifically as it comprises extensive flat surfaces (asphalted ground and buildings with mostly flat roofs, as well as cuboid shipping pallets).

#### 4.1.1. Analysis of in-flight inertial measurements

To get an a-priori estimate of the errors which might be incurred by aliasing or cut-off effects when downsampling, a spectral analysis of the 850 Hz inertial sensor data is performed. Fig. 9 shows the spectrum for the full flight period.

Apart from the actual platform motion, there are further peaks due to vibration at around 45 Hz, 90 Hz and 110–120 Hz. Because this is aggregated for the whole flight, there will be periods with peak vibrations exceeding the magnitudes shown here, as well as periods with less vibrations. Generally, vibrations increase during maneuvers (e.g., turns) and are slightly lower within the flight lines where a consistent speed and heading is maintained. The exact source of the vibrations is not known, but possibly related to either the scan mechanism or UAV rotor rotation. Recovering the 120 Hz signal would require a sampling rate of at least 240 Hz. A computable heuristic for the error caused by insufficient sampling rate is provided by Eq. (2.9), and shown in Fig. 10 for the orientation. The position is less relevant here, for two reasons: The high-frequency position error is relatively small (less than 0.15 mm, see also Fig. 11), and it results directly in a translation error in the point cloud of similar magnitude. On the other hand, an orientation error of 18 mdeg (corresponds approximately to the roll error for a downsampling factor of 4), would cause a point displacement of 3.1 cm at 100 m range<sup>5</sup>. It can also be seen that the sequential increases in predicted error occur at downsampling factors corresponding to the peaks in the signal spectrum. The increases at  $d = 4$  and  $d = 8$  correspond to sampling rates of 212.5 Hz and 106.25 Hz. Note that 200 Hz, slightly below  $d = 4$ , is a sampling frequency often used by inertial navigation systems for surveying applications. Thus, for this specific carrier platform and payload, the standard choice of sampling rate already leads to appreciable signal




Fig. 9. Spectrum of inertial measurements from the DJI M350 quadcopter UAV with the RIEGL miniVUX-3UAV as payload. The graph's x-axis shows the frequency in Hz, and the y-axis the corresponding accelerometer or gyroscope signal amplitude.




Fig. 10. Predicted root mean square error (RMSE) in orientation when decimating the IMU measurements with a given factor. IMU data is used for the full flight of the DJI Matrice M350 and RIEGL miniVUX-3UAV.

cut-off. The (predicted) orientation errors increase markedly as higher downsampling factors are used - which would be desirable in order to reduce runtime and memory usage of optimization-based GNSS/IMU-integration algorithms. To investigate the actual impact of these errors on point cloud precision, the processed trajectory and georeferenced point clouds are now analyzed.

#### 4.1.2. Differences in the estimated trajectories

In order to quantify the effects of downsampling on the estimated trajectory, we analyze the differences between the baseline trajectory (standard processing, i.e., with IMU observations at the full output rate) and trajectories processed with downsampled IMU measurements. We do this for standard downsampling as well as for the suggested down-/upsampling strategy, and compare both the trajectories themselves

<sup>5</sup> 100 m is the distance from scanner to the left or right strip boundary at 50 m AGL.

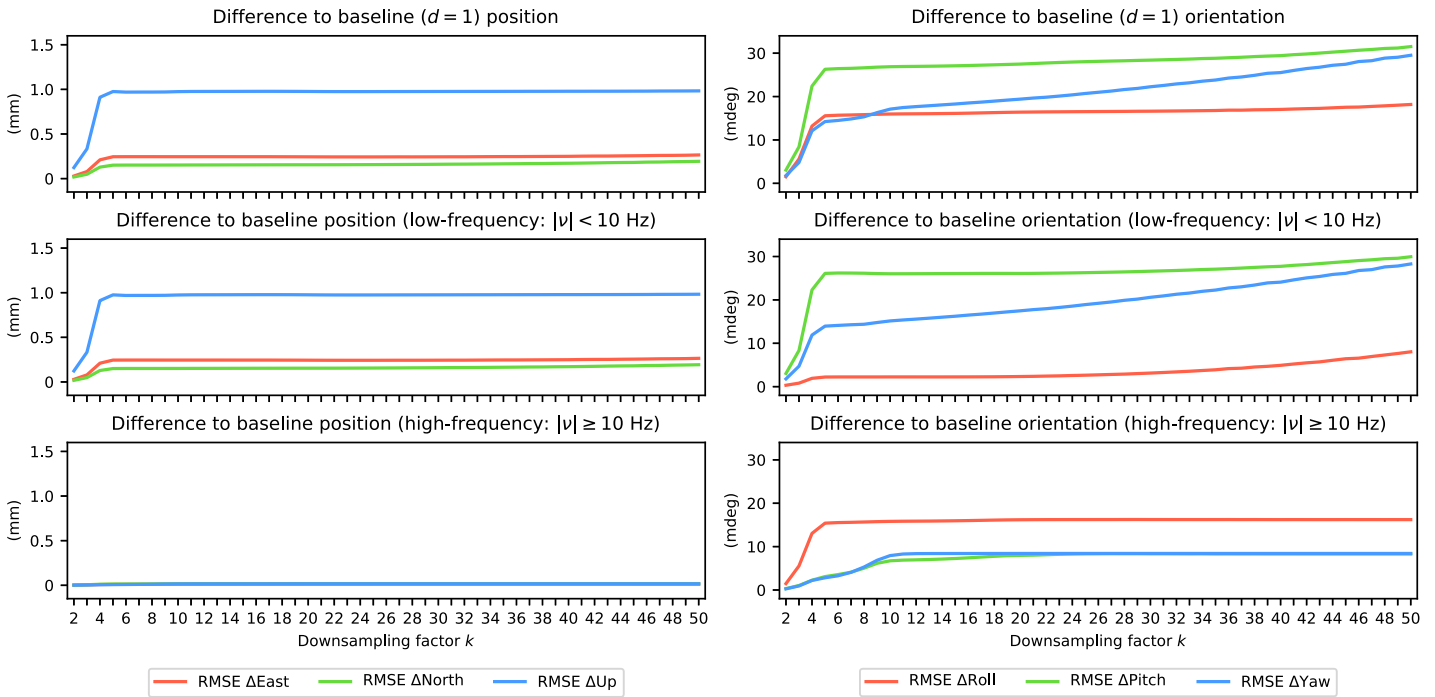


Fig. 11. Difference between baseline GNSS/IMU trajectory and trajectory processed with downsampled IMU measurements without reconstruction of high-frequency components (V1), for the full flight period.

as well as the point cloud georeferenced with these trajectories. This analysis is carried out once for GNSS/IMU integration and once for GNSS/IMU/LiDAR-integration.

The trajectory estimation based on GNSS/IMU integration is performed as described in Section 3. For the LiDAR-specific measurement equations and feature matching/extraction method see Pöpl et al. (2024). The processing is done once for  $d = 1$ , i.e., standard processing with no downsampling, and as two distinct variants for  $d > 1$ :

- **V1:** Anti-aliasing via low-pass filter followed by downsampling (see Section 3.3). This variant is meant to mimic the output of an IMU set to a lower output rate, which performs the anti-aliasing and downsampling in real-time and outputs the corresponding lower-frequency integrated  $\Delta v$  and  $\Delta \theta$  values.
- **V2:** Downsampling as above, but without anti-aliasing and with reconstruction of the higher frequency trajectory components (see Section 3.4) and implicit modeling of the higher-frequency components in the adjustment (see Section 3.5). The downsampling of the IMU measurements, upsampling of the high-frequency trajectory and the NLS optimization are iterated until the RMS change in the downsampled IMU measurements (due to convergence of the IMU error estimates) is below  $0.01 \frac{m}{s^2}$  and  $0.025 \frac{mdeg}{s}$  (see Section 3.6).

Note that in both cases, the number of estimated parameters and measurements in the adjustment are the same, additional effort lies only in the repeated downsampling of the IMU measurements and in possibly increased number of NLS iterations.

This processing is performed once at full IMU frequency ( $d = 1$ ) to obtain a baseline trajectory. This will be the reference for evaluating the impact of the downsampling on the trajectory. Then for each  $d \in \{2, \dots, 50\}$ , two separate trajectories are processed as described above. The overall difference between the respective  $d > 1$  trajectory and the baseline trajectory is split into two parts, a high-frequency ( $v > 10$  Hz) part and a low-frequency ( $v \leq 10$  Hz) part. This is done to highlight the effects of downsampling, which manifest in two ways: oscillating errors caused by unresolved high-frequency trajectory dynamics and low-frequency errors caused by inconsistent IMU error estimates (cf. Section 3.6).

Fig. 11 shows the differences in position and orientation between baseline trajectory ( $d = 1$ ) and  $d > 1$  trajectories derived from  $d$ -times downsampled (V1) IMU measurements. These emulate trajectories obtained from IMUs with  $d$ -times reduced sampling rate where measurements are appropriately aggregated internally but only output at lower rate. As expected, the downsampling has relatively little effect on the position, which exhibits a maximum RMS difference (w.r.t. the baseline position) below 1 mm in the low-frequency [0–10] Hz band, and below 0.15 mm in the high-frequency [ $>10$ ] Hz band. The orientation error is more significant: in the low-frequency band, the downsampling indirectly causes errors up to 30 mdeg due to differing IMU bias and scale factor estimates. In the high-frequency band, the anti-aliasing and downsampling cuts off the respective frequency components above the reduced Nyquist rate causing an RMS difference of up to 16 mdeg. The errors after trajectory estimation, i.e., after integration with GNSS, generally agree with the heuristics derived from the frequency spectrum (Fig. 10) but are slightly lower. The increase is not linear but increases sharply at  $d = 4$  (corresponding to a Nyquist frequency of 106.25 Hz, cf. Fig. 9). When downsampling at exactly the frequency of strong vibration, some residual aliasing occurs. In the context of trajectory estimation, this causes an additional problem: a signal which oscillates with possible high amplitude now looks constant, and IMU bias and scale factor can no longer be distinguished. This makes it even more difficult to estimate IMU biases for the downsampled IMU measurements, especially at these frequencies. This is true both for position (respectively, the accelerometer) and orientation (respectively, the gyroscope). However, low-frequency errors can be compensated in a strip adjustment or LiDAR-integrated trajectory estimation, while the high-frequency errors generally cannot.

The techniques introduced in Sections 3.4–3.6 aim to mitigate both the low-frequency errors (due to inconsistent IMU error estimates) and the high-frequency errors (due to unresolved vibrations). Fig. 12 shows that for the GNSS/IMU adjustment, now performed with trajectory upsampling and IMU error iteration (V2), the previously seen discrepancies between baseline trajectory and  $d > 1$  trajectories reduce drastically. There are residual low-frequency errors at the downsampling factors corresponding to frequencies with strong vibrations and

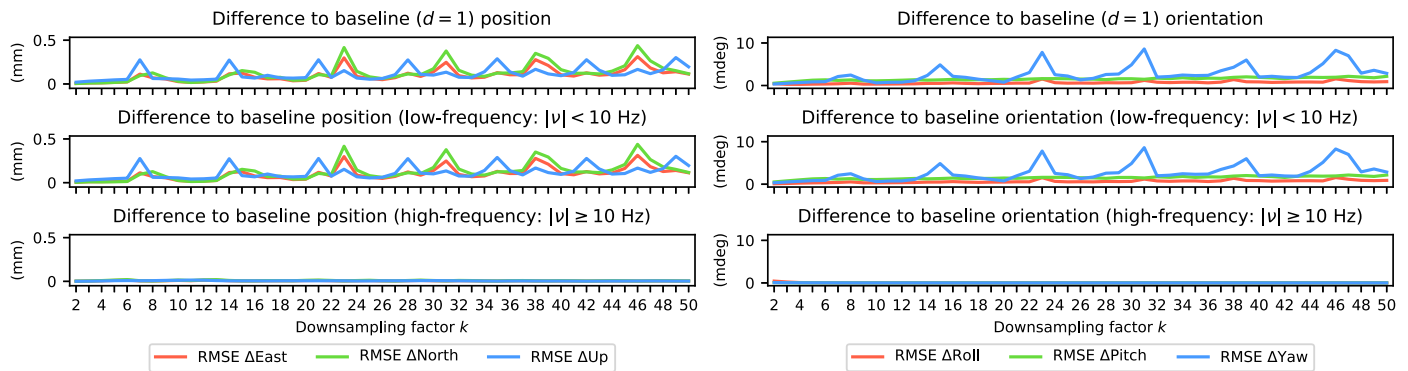


Fig. 12. Difference between baseline GNSS/IMU trajectory and trajectory processed with GNSS and downsampled IMU measurements with reconstruction of high-frequency components (V2), for the full flight period.

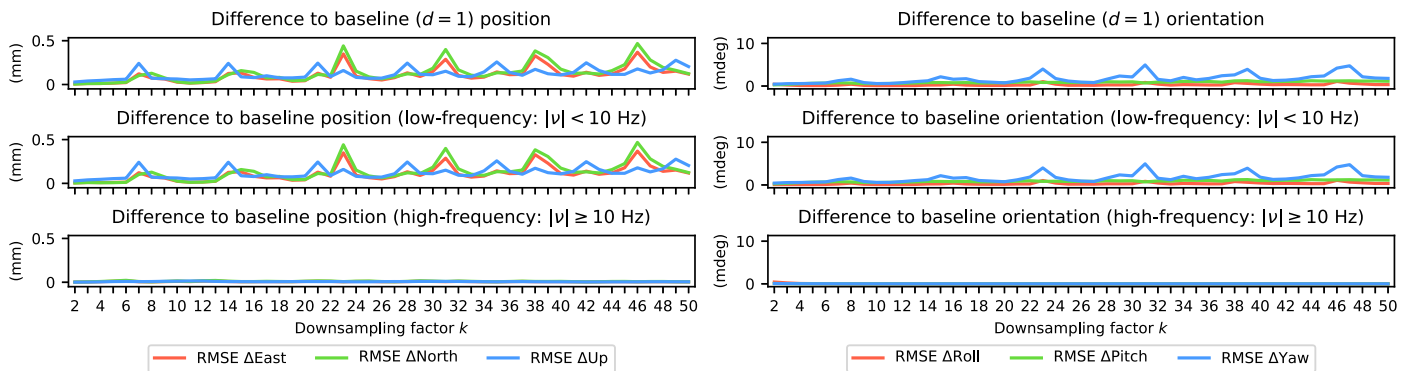


Fig. 13. Difference between baseline GNSS/IMU/LiDAR trajectory and trajectory processed with GNSS, LiDAR and downsampled IMU measurements with reconstruction of high-frequency components (V2), for the full flight period.

multiples thereof, both for position and orientation. However, these are at most 0.5 mm and 10 mdeg respectively, and therefore well below both the errors in Fig. 11 and the expected accuracy of the GNSS/IMU solution.

The remaining low-frequency differences, occurring due to difficulties in estimating IMU errors when using highly downsampled measurements in the presence of vibrations at the downsampling rate, are then further reduced when integrating LiDAR-derived plane observations in the adjustment (Fig. 13). As the aim is specifically to investigate the effects of downsampling itself, the matching/subsampling (Pöpl et al., 2024) of the LiDAR plane features is done only once for  $d = 1$  and then re-used for the other processing runs in order to eliminate any differences due to the feature matching process. The suggested strategy of repeated downsampling and IMU error iteration (Section 3.6) successfully mitigates low-frequency trajectory errors and completely eliminates high-frequency trajectory errors. The RMS difference in the estimated trajectory (w.r.t.  $d = 1$ ) is below 0.5 mm in position and 5 mdeg in orientation, even for downsampling factors of up to  $d = 50$  (corresponding to an IMU sampling rate of 17 Hz).

#### 4.1.3. Precision of the georeferenced point cloud

To evaluate the single-strip point cloud precision for the GNSS/IMU-derived point cloud, a single flight strip is analyzed and the single-strip precision of each  $d > 1$  point cloud compared to that of the  $d = 1$  baseline point cloud. For each point cloud, a precision grid is derived by computing a PCA within each grid cell. Then, for each grid cell, the precision  $\sigma_1$  of the baseline  $d = 1$  point cloud is subtracted from the downsampling-derived  $d > 1$  point cloud's precision  $\sigma_d$  to obtain the relative precision change  $\Delta\sigma_d := \sigma_d - \sigma_1$ .

This analysis of precision (i.e., point spread) on planar surfaces assumes that these surfaces are sufficiently planar, implying that most of the point spread is due to measurement errors. To ensure this is the case

in the study area, a reference point cloud (Fig. 14(a)) of the core area of interest was acquired via terrestrial laser scanning. Fig. 14(b) shows the corresponding precision grid derived from the TLS point cloud, with standard deviations consistently below 3 mm. In comparison, the ULS precision grid (for the baseline  $d = 1$  point cloud, Fig. 15(a)) shows standard deviations of approximately 1 cm on planar surfaces (ground, roofs) with slightly higher point spread at the strip edges in some areas. It can therefore be assumed that the major part of the  $[ \approx 1 ]$  cm standard deviation is not due to model errors (i.e. non-planar geometry), but due to measurement errors. These measurement errors may be ranging errors, scanning errors or trajectory errors, or errors due to high angle of incidence and consequently elongated laser footprint causing view-dependent discrepancies on non-smooth surfaces. While these cannot be distinguished for the baseline  $d = 1$  solution, the additional point spread incurred from downsampling can be quantified through the relative precision  $\Delta\sigma_d$ , obtained by subtracting the baseline precision grid from the corresponding  $d > 1$  precision grid. Since we are interested only in sufficiently planar surfaces, grid cells with a standard deviation above 10 cm are masked.

The relative precision grids for downsampling (V1) with  $d = 4$  and  $d = 8$  are shown in Figs. 16(a) and 16(b). At  $d = 4$ , a small increase in point spread is visible, predominantly at the strip edges due to high-frequency roll errors. This corresponds to a sampling frequency of 212.5 Hz, which is already insufficient to properly resolve the vibrations present, especially in the forward-axis corresponding to the UAV roll angle (cf. Fig. 9). In this case, the standard deviation at the strip edges increases by up to 0.5 cm compared to the baseline point cloud. The magnitude of this effect increases with the downsampling factor  $d$ . Fig. 18(a) shows the aggregated (RMS and 99-th percentile) relative standard deviation for all  $d \in \{2, \dots, 50\}$ . Note that the relative standard deviation increases sharply up to  $d = 8$  (106.25 Hz), and only slowly afterwards. At this frequency, most of the vibrations are no longer present



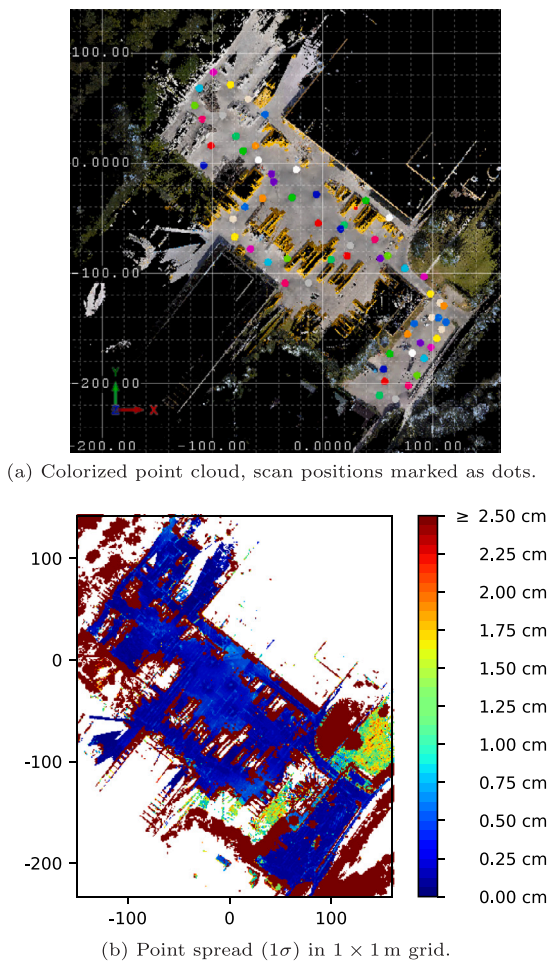


Fig. 14. Terrestrial laser scanning (TLS) data acquisition consisting of a total of 62 scan positions, performed simultaneously with the ULS data acquisition.

in the downsampled signal, and further downsampling incurs mainly low-frequency errors which are not visible in the single-strip precision at a grid size of  $1 \times 1$  m. Figs. 16(c) and 16(d) show the corresponding relative precision grids for the down-/upsampling variant (V2) with  $d = 4$  and  $d = 8$ , which do not exhibit a decrease in precision. In this case, the standard deviation increases only for the downsampling factors  $d$  corresponding to the main vibration frequencies, in which case the increase is still below 1 mm (see Fig. 18(a)).

The multi-strip precision is evaluated in the same way for the GNSS/IMU/LiDAR-derived point cloud. Figs. 17(a) and 17(b) now show the relative precision grids for downsampling (V1) with  $d = 4$  and  $d = 8$  of the downsampling-derived point clouds compared to the baseline point cloud (Fig. 15(b)). Now, the precision grid includes points from both flight strips, thus showing both high-frequency errors (increasing point spread within each strip) and low-frequency errors (increasing overall point spread). An increased standard deviation by  $[\approx 1]$  cm can be seen for  $d = 4$ , and above 1.5 cm for  $d \geq 8$ . The relative difference tops out at  $d = 8$  (Fig. 18(b)), at which point the high-frequency errors no longer increase and the low-frequency errors are corrected in the LiDAR-integrated adjustment. The up-/downsampling variant (V2) does not exhibit notably increased point spread (decreased precision) for any  $d$  (Figs. 17(c), 17(d) and 18(b)). In comparison to the GNSS/IMU adjustment, the errors at downsampling factors corresponding to frequencies with significant vibration are now also mitigated. Consequently, while point spread is shown to increase with standard

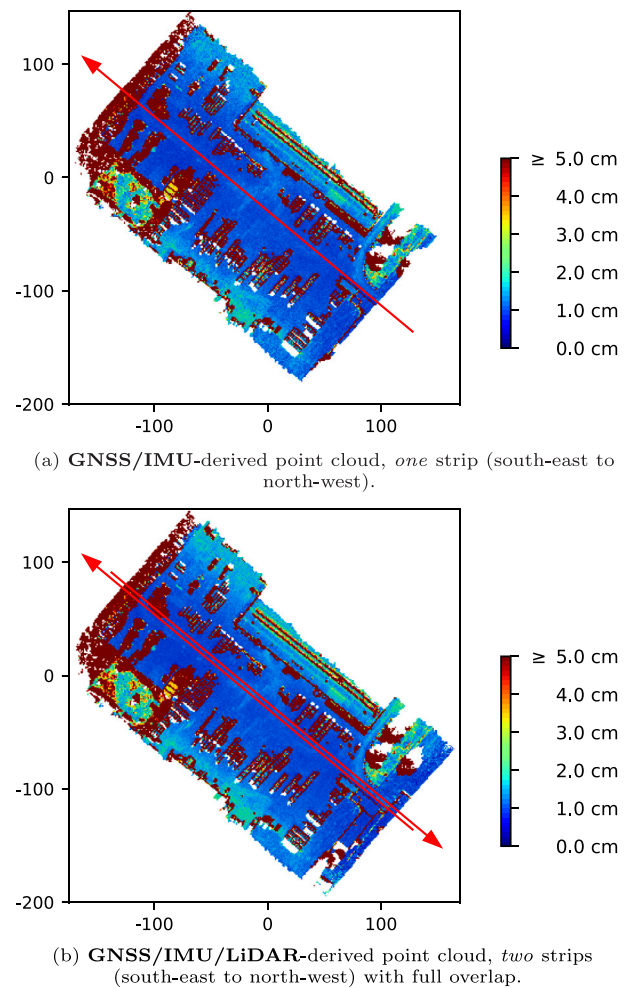


Fig. 15. Standard deviation ( $1\sigma$ ) of ULS flight strip(s) in a  $1 \times 1$  m grid. The point cloud is georeferenced with the GNSS/IMU and GNSS/IMU/LiDAR  $d = 1$  trajectories. In each grid cell, a principal component analysis of all points is performed, and the grid cell is colored by the resulting standard deviation in normal direction. The lower figure shows two flight strips, i.e., the multi-strip standard deviation computed from both strips' points.

downsampling (V1, see Section 3.3), there is no indication of degradation in point cloud quality when downsampling with the proposed strategy (V2, described in Sections 3.3–3.6).

#### 4.1.4. Runtime and memory usage

The key benefit of using downsampled measurements is a significant reduction in the number of observations and parameters in the non-linear optimization. This directly reduces the size of the linear equation system which must be solved at each NLS iteration step. Consequently, the runtime and memory usage of the adjustment is expected to reduce in proportion with downsampling factor  $d$ , down to a lower limit. Fig. 19 provides runtime, memory usage and number of iterations for all processing runs. For both variants, memory usage for  $d = 1$  is approximately 12 GiB, and decreases proportionally to  $\frac{1}{d}$  with a lower limit of approximately 1.3 GiB for the GNSS/IMU processing and 2.3 GiB for the GNSS/IMU/LiDAR processing. Without D/U iteration, the runtime is approximately  $[(\frac{120}{d} + 3)]$  s for the GNSS/IMU adjustment and  $[(\frac{140}{d} + 6)]$  s for the GNSS/IMU/LiDAR adjustment. With D/U iteration, the factor with which runtime decreases as  $d$  increases is slightly less than  $\frac{1}{d}$ , due to the higher total number of NLS iterations required. The runtime is also reported as normalized with respect to the number of downsampled IMU measurements. This



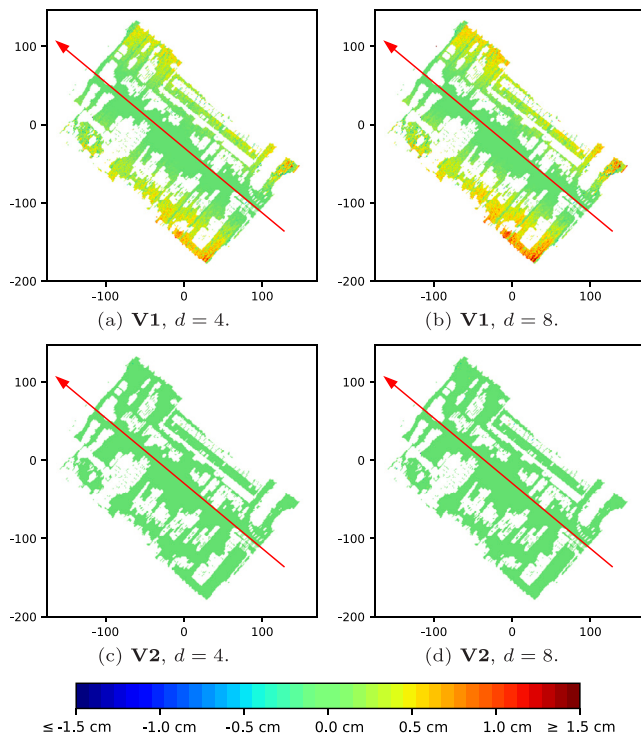


Fig. 16. Single-strip relative precision  $\Delta\sigma_d$  of downsampling-derived  $d > 1$  point clouds (south-east to north-west, one strip) compared to the baseline  $d = 1$  point cloud (Fig. 15(a)), georeferenced with the GNSS/IMU trajectory.






Fig. 17. Multi-strip relative precision  $\Delta\sigma_d$  of downsampling-derived  $d > 1$  point clouds (south-east to north-west, two strips, full overlap) compared to the baseline  $d = 1$  point cloud (Fig. 15(b)), georeferenced with the GNSS/IMU/LiDAR trajectory. This shows two flight strips, and thus the multi-strip relative standard deviation computed from both strips' points.



(a) Relative precision values for the point clouds (south-west to north-east, one strip) georeferenced with the GNSS/IMU trajectory.



(b) Relative precision values for the point clouds (south-east to north-west, two strips with full overlap) georeferenced with the GNSS/IMU/LiDAR trajectory.

Fig. 18. Relative precision values for  $d \in \{2, \dots, 50\}$ . Given is the 99-th percentile of the relative standard deviations  $P_{99}(\Delta\sigma_d)$ , and the root mean square relative standard deviation  $RMS(\Delta\sigma_d)$ . The red line corresponds to the standard downsampling (V1) and the blue line to the suggested down-/upsampling approach (V2).

normalized runtime is almost constant for the GNSS/IMU adjustment without downsampling, but increases as  $d$  increases when LiDAR-derived observations are included and when performing D/U iterations. For the GNSS/IMU/LiDAR adjustment, the lower bound for the runtime depends on the number of LiDAR-derived observations and parameters. Here, 7000 plane observations (with 3 observation each) corresponding to 1700 modeled planes (with 3 parameters each) are used. In this case, there is no benefit with respect to runtime for  $d > 16$ .

In this analysis, runtimes do not include time required for georeferencing (performed twice, see Fig. 4) and LiDAR feature extraction & matching, as these steps are independent of IMU downsampling. Similarly, the runtime values do not include coarse trajectory initialization, required for the NLS optimization. An initial trajectory may be obtained with arbitrary downsampling factor, or even completely outside the adjustment with e.g., a Kalman filter. Here it is realized with a sliding-window optimization with matching downsampling factor  $d$ . The total runtime for  $d = 1$  is 660s and includes all processing steps: initialization, GNSS/IMU adjustment, preliminary georeferencing, feature extraction and matching, GNSS/IMU/LiDAR adjustment, final georeferencing. The two NLS adjustments take up almost half of the runtime. The time required for initialization decreases with the downsampling factor in the same manner as the NLS runtime (cf. Fig. 19), while all other processing steps have the same runtime, independent of downsampling factor.

#### 4.2. Case study: Airborne LiDAR

The performance of the proposed methodology applied to airborne laser scanning (ALS) with crewed fixed-wing aircraft is assessed using an ALS dataset acquired over the town of Horn, Lower Austria. Here, a RIEGL VQ-580-II laser scanner is used, mounted on a Cessna aircraft in a RIEGL VQX-1 wing-pod. The laser scanner is integrated with a triple-frequency GNSS receiver/antenna and a navigation-grade IMU. This IMU is configured for an output rate of 600 Hz, and the GNSS receiver output rate is 10 Hz. In contrast to the ULS dataset, a high-grade IMU is used here which comprises pendulous accelerometers and




Fig. 19. Runtime (overall, and divided by number of downsamped IMU measurements), memory usage, and iteration counts of the different processing runs.

fiber-optic gyroscopes. This is comparable to survey-grade navigation systems used in airborne mapping with an expected orientation accuracy of 15 mdeg. Again, the focus is on the impact on the trajectory itself as well as the achieved precision (single-strip and multi-strip) of the georeferenced point cloud, specifically when using downsamped IMU measurements together with the upsampling and iteration strategy described in Section 3.

The LiDAR data acquisition consists of 9 flight strips, flown at two different altitudes (550 m and 715 m AGL) with a ground velocity of 50–60  $\frac{m}{s}$  (Fig. 20). Additionally, GNSS and IMU data is available and processed for the full trajectory, which includes approximately 50 km round-trip to and from the airfield with a total flight time of 45 min. The absolute accuracy is evaluated by comparison to 62 independently surveyed height check points. The final ALS point cloud is however not fully independent, as the z-component of the GNSS lever arm was determined from comparison with the ground data; it was not precisely measured prior to flight and could not be determined with sufficient accuracy from the in-flight data alone.

#### 4.2.1. Analysis of in-flight inertial measurements

The signal spectrum (Fig. 21) shows that the overall level of vibrations is an order of magnitude lower than for the ULS dataset due to the type of mounting and less stringent restrictions on size and weight, which allow for better vibration dampening. Peaks are present at frequencies of approximately 110–120 Hz, 150–160 Hz and 220–240 Hz. The expected high-frequency error caused by anti-aliasing and downsampling (effectively cutting-off the respective frequency components) can be estimated according to Eq. (2.9). In contrast to the ULS dataset, the errors increase slower at first, but peak at a similar level of 30 mdeg for high downsampling factors (respectively, low sampling rates), as shown in Fig. 22. At  $d = 3$  the expected RMSE is below 0.5 mdeg for roll, pitch and yaw, indicating that in this case, a 200 Hz sampling rate would be sufficient.

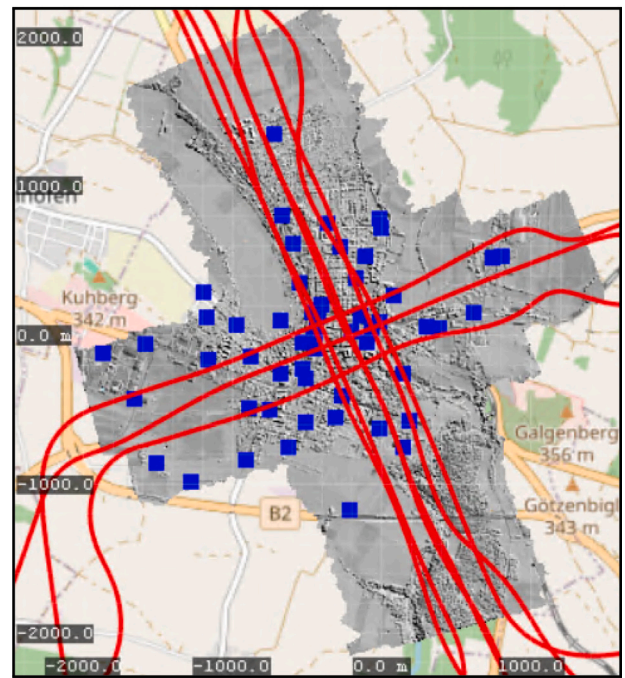


Fig. 20. ALS point cloud by reflectance (dark/bright = low/high reflectance, hillshade) and flight trajectory, superimposed onto map (© OpenStreetMap contributors). The X/Y-axes are those of a local east/north/up (ENU) Cartesian coordinate system. The 62 height check points are marked with blue rectangles.

#### 4.2.2. Differences in the estimated trajectory

The effects of downsampling on the estimated trajectory are again evaluated by comparing a baseline trajectory, processed with no downsampling (i.e.,  $d = 1$ ) to trajectories processed with downsamped IMU measurements. The common case of 200 Hz ( $d = 3$ ) and subsequent halving of the output rate (i.e.,  $d \in \{6, 12, 24, 48\}$ ) are investigated in more detail. For each  $d$ , the adjustment is performed once with only downsampling (V1), and once with down-/upsampling and error iteration (V2), and respectively with only GNSS/IMU observations and again for GNSS/IMU/LiDAR observations.

The RMS differences of the resulting  $d > 1$  trajectories to the baseline  $d = 1$  trajectory are given in Table 1a–1d. The differences are in all cases much lower than for the ULS dataset. This is due to less vibrations and the higher grade of the IMU, especially lower biases, scale factors, and less temporal variation thereof. Specifically, there is only about 0.15 mm and 0.5 mdeg RMS difference between the trajectory derived from the full 600 Hz inertial measurements and the downsamped (V1) 200 Hz measurements (Table 1a). This is expected when comparing to Fig. 22, which indicates no notable high-frequency orientation errors will occur at  $d = 3$ . For higher downsampling factors, the errors do increase again with an RMS orientation difference of over 20 mdeg for  $k = 48$ .

The difference in orientation between  $d = 1$  and  $d > 1$  is reduced by more than a factor of 10 when using upsampling and error iteration (V2, Table 1b). For the adjustment including the LiDAR planar observations (Table 1c), the RMS orientation differences for downsampling-only (V1) processing are on the same level as without LiDAR observations. In comparison, these errors are again reduced by a factor of at least 10 when performing upsampling and error iteration (V2, Table 1d). For this final trajectory, the RMS difference between  $d = 1$  and  $d \in \{3, 6, 12\}$  is below 0.1 mm and 0.1 mdeg, with marginally higher position differences for  $d = 24$  and  $d = 48$ .

#### 4.2.3. Precision and accuracy of the georeferenced point cloud

The point cloud precision of the proposed methodology, now applied to airborne data, is evaluated in the same way as in Section 4.1.




Fig. 21. Spectrum of inertial measurements from a Cessna fixed-wing aircraft with the laser scanner mounted inside a wing-pod, for the duration of the LiDAR data acquisition.




Fig. 22. Predicted root mean square error (RMSE) in orientation if high-frequency measurements are disregarded. Plot shows predicted RMSE when downsampling with a given factor, for data from the VQ-580-II mounted on the wing-strut of a Cessna aircraft.

In the interest of brevity, we analyze only the final product, i.e., the point cloud georeferenced with the trajectory estimated from GNSS, IMU, and LiDAR observations. Fig. 23 shows the  $1\sigma$  precision grid for the  $d = 1$  point cloud. On relatively flat surfaces (e.g., roads, rooftops), a standard deviation of 1 cm is reached. Again, relative precision grids are computed by comparison of the  $d > 1$  precision grid to the baseline ( $d = 1$ ). In planar areas (ground, building roofs, etc.), an increase in the standard deviation can be attributed to trajectory errors; in vegetated or otherwise non-smooth areas, this is not the case and thus grid cells with a standard deviation above 10 cm are again masked. The precision grids for the downsampling-only (V1) processing runs are shown in Figs. 24(a)–24(d). As expected from the trajectory analysis (Fig. 22), no decrease in precision is discernible for  $d = 3$ , but standard deviation

Table 1

Root mean square (RMS) differences of  $d > 1$  trajectories with respect to the baseline  $d = 1$  trajectory.

(a) Downsampling (V1), GNSS/IMU.					
$d$	3	6	12	24	48
X (mm)	0.05	0.05	0.07	0.18	0.66
Y (mm)	0.04	0.05	0.06	0.15	0.56
Z (mm)	0.13	0.13	0.14	0.25	0.82
Roll (mdeg)	0.25	1.00	2.31	7.16	20.85
Pitch (mdeg)	0.32	0.42	1.59	5.34	16.35
Yaw (mdeg)	0.40	0.51	1.51	2.94	7.03
(b) Down-/upsampling (V2), GNSS/IMU.					
$d$	3	6	12	24	48
X (mm)	0.04	0.06	0.08	0.14	0.20
Y (mm)	0.03	0.05	0.06	0.12	0.18
Z (mm)	0.03	0.07	0.10	0.34	0.41
Roll (mdeg)	0.02	0.02	0.03	0.08	0.12
Pitch (mdeg)	0.02	0.03	0.03	0.07	0.11
Yaw (mdeg)	0.01	0.03	0.05	0.07	0.14
(c) Downsampling (V1), GNSS/IMU/LiDAR.					
$d$	3	6	12	24	48
X (mm)	0.08	0.18	0.64	0.81	3.36
Y (mm)	0.10	0.28	0.50	0.73	2.17
Z (mm)	0.14	0.17	0.29	0.66	3.19
Roll (mdeg)	0.24	1.00	2.31	7.17	20.89
Pitch (mdeg)	0.32	0.43	1.59	5.35	16.43
Yaw (mdeg)	0.41	0.54	1.53	3.00	7.60
(d) Down-/upsampling (V2), GNSS/IMU/LiDAR.					
$d$	3	6	12	24	48
X (mm)	0.02	0.06	0.07	0.15	0.20
Y (mm)	0.03	0.05	0.07	0.15	0.19
Z (mm)	0.03	0.07	0.09	0.33	0.40
Roll (mdeg)	0.01	0.02	0.02	0.04	0.06
Pitch (mdeg)	0.01	0.01	0.02	0.04	0.07
Yaw (mdeg)	0.01	0.03	0.03	0.06	0.10

Table 2

Root mean square (RMS) height distance (in cm) of point clouds processed with downsampling-only (V1) and down-/upsampling (V2) to the 62 height check points.

$d$	1	3	6	12	24	48
Downsampling only		2.9	2.9	2.9	3.7	6.2
Down-/upsampling	2.9	2.9	2.9	2.9	2.9	2.9

does increase for  $d \in \{6, 12, 24, 48\}$ . Aggregated metrics are shown in Fig. 25. The RMS and  $P_{99}$  values of the relative standard deviation increase already for  $d = 6$  and again for  $d = 12$ ,  $d = 24$  and  $d = 48$ , with the RMS standard deviation for  $d = 48$  being 3 cm higher than for  $d = 1$ . This is mitigated by the proposed upsampling and error iteration strategy (V2), which results in equally precise point clouds for all  $d$ .

To ensure there is no negative effect on the georeferencing accuracy, Table 2 reports the RMS distances of each point cloud to the height check points. By this measure, the proposed down-/upsampling strategy (V2) retains the baseline accuracy for all  $d$ . For the downsampling-only (V1) point cloud, the RMS with respect to the height check points is at a similar level for  $d \in \{3, 6, 12\}$  but does increase for  $d = 24$  and  $d = 48$ .

#### 4.2.4. Runtime and memory usage

In this case, the total runtime including all processing steps for  $d = 1$  is 2300 s (=38.3 min) of which 720 s (=12.0 min) are spent on initialization (i.e., sliding-window NLS adjustment) and 820 s (=13.6 min) on the two full NLS adjustments. Even though the IMU sampling rate (and therefore the frequency of trajectory parameters) is slightly lower than for the ULS dataset, the flight is longer, resulting in an overall much larger number of parameters and observations. Consequently, the runtime and memory usage is higher (Fig. 26). Runtime and memory usage again decrease with the downsampling factor, with a lower limit



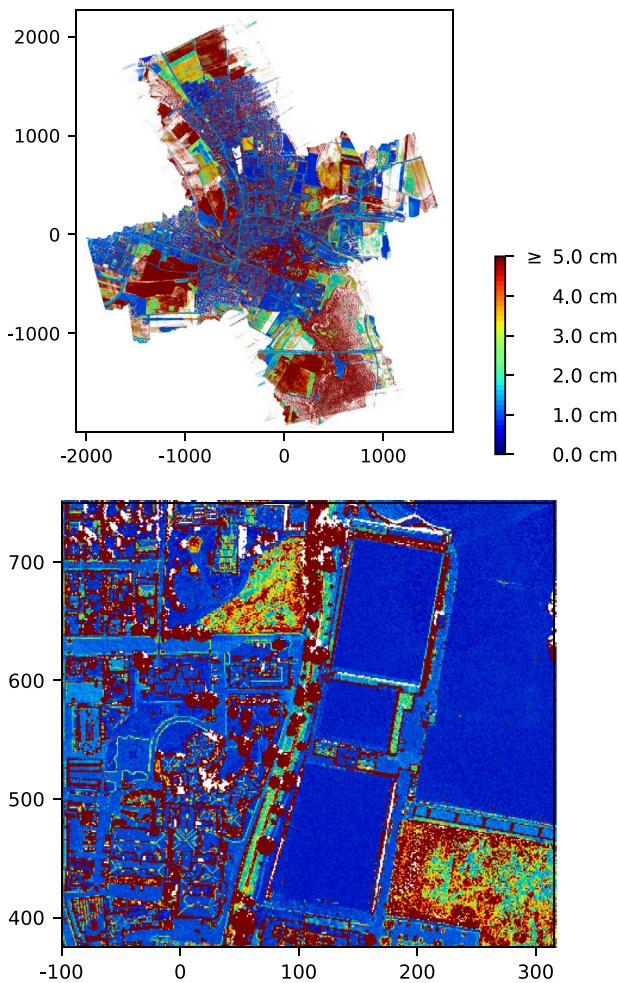


Fig. 23. Standard deviation ( $1\sigma$ ) in a  $1 \times 1$  m grid. The point cloud is georeferenced with the GNSS/IMU/LiDAR  $d = 1$  trajectory. In each grid cell, a principal component analysis of all points is performed, and the grid cell is colored by the resulting standard deviation in normal direction. The standard deviation grid is shown for the full flight block (top) and magnified for a smaller area of interest in the center of the block (bottom). The latter contains large flat surfaces (building roofs, parking spaces, sport courts) with a standard deviation below 1 cm.

of approximately 1.2 min and 4.5 GiB. The achievable reduction in runtime and memory usage is limited by the number of LiDAR-derived measurements and the corresponding plane parameters. In this case, both runtime and memory usage are reduced up to a factor of 6, with no discernible impact on point cloud accuracy and precision.

## 5. Discussion & conclusion

In this contribution, we have presented methodology for quantifying and mitigating the effects of platform vibration on LiDAR point cloud precision in the context of GNSS/IMU/LiDAR-based trajectory estimation and georeferencing. Specifically, we propose using downsampled IMU measurements together with a simple procedure for recovering high-frequency trajectory components from a low-frequency trajectory and high-frequency inertial measurements. This procedure is applied to adjustment-based GNSS/IMU and GNSS/IMU/LiDAR integration, where it allows usage of downsampled inertial measurements without sacrificing trajectory accuracy, thereby significantly reducing computational runtime and memory usage of such methods. Results for this approach are presented and evaluated for two airborne data acquisitions with different but typical carrier platforms, a small quadcopter UAV and a crewed fixed-wing aircraft.

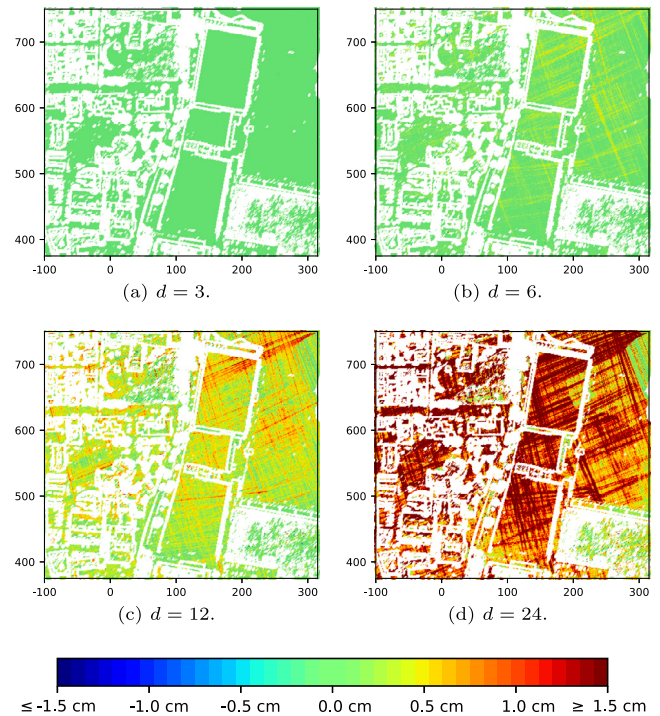


Fig. 24. Multi-strip relative precision  $\Delta\sigma_d$  of downsampling-derived  $d > 1$  (V1) point clouds compared to the baseline  $d = 1$  point clouds, georeferenced with the GNSS/IMU/LiDAR trajectory.




Fig. 25. Relative precision values for  $d \in \{3, 6, 12, 24, 48\}$ , point cloud georeferenced with GNSS/IMU/LiDAR trajectory. Given is the 99-th percentile of the relative standard deviations  $P_{99}(\Delta\sigma_d)$ , and the root mean square relative standard deviation  $RMS(\Delta\sigma_d)$ . The red lines correspond to the standard downsampling (V1) and the blue lines to the suggested down-/upsampling approach (V2).

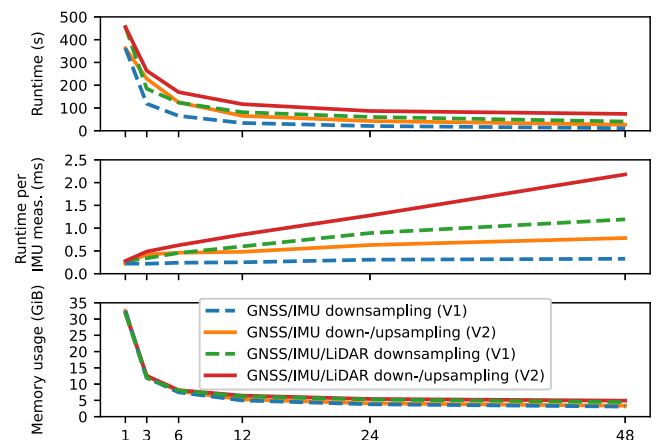


Fig. 26. Runtime (overall, and divided by number of downsampled IMU measurements) and memory usage of the different processing runs.



By analyzing in-flight inertial data as well as the estimated trajectories and the georeferenced point clouds, we demonstrate that for the UAV platform, an IMU sampling rate of 200 Hz can be a bottleneck for the precision of the point cloud. The UAV system used here comprises only an entry-level MEMS IMU and features strong vibrations, which is not atypical for such platforms. On the other hand, the navigation-grade IMU used for the ALS data acquisition exhibits much lower biases, scale factor errors and noise; at the same time, less vibrations are present. In this case, no reduction in point cloud precision is apparent for the 200 Hz sampling rate, although sampling rates at or below 100 Hz do lead to degraded point cloud quality.

It is also demonstrated that the proposed methodology reduces the frequency and number of IMU measurements and trajectory parameters involved in the NLS adjustment, while having negligible impact on the estimated trajectory and in consequence no discernible negative impact on the precision or accuracy of the final point cloud. Adjustment-based trajectory estimation methods, where parameters are estimated simultaneously, are often limited by runtime and memory usage. For both datasets used here, the time spent for NLS adjustment is about half of the total runtime, which includes GNSS/IMU adjustment, preliminary georeferencing, plane extraction and matching, GNSS/IMU/LiDAR adjustment and final georeferencing. Memory usage is highest for the GNSS/IMU/LiDAR adjustment. With the proposed methodology, runtime and memory usage are reduced by 85% to approximately 1/6 of the baseline while preserving precision and accuracy of the georeferenced point cloud, thus allowing for the processing of larger datasets or higher measurement rates. Nevertheless, gains in runtime and memory usage of the adjustment-based trajectory estimation are limited by the LiDAR observations. Runtime and memory usage of the GNSS/IMU/LiDAR adjustment depend on the number of measurements and parameters, which is influenced by various factors such as GNSS, IMU, and LiDAR sensor measurement rates as well as the number of plane features and corresponding objects. The latter depends on point cloud overlap between and within strips. Overlaps between different strips, especially with large temporal separation, reduce the sparsity of the equation system and thereby cause the adjustment's runtime and memory usage to scale superlinearly with the acquisition time. Larger datasets also tend to require proportionally more LiDAR correspondences to ensure consistency between flight strips. While it is possible to split up large datasets into separate tiles and process them individually, this is again a multi-step approach where some tiling has to be applied first, and consistency between tiles must be ensured subsequently. In contrast, the holistic and contiguous processing employed here is attractive due to increased redundancy in the estimation and the possibility of explicitly modeling and exploiting the continuity in trajectory and time-varying IMU errors. Thus, to allow such adjustment-based trajectory estimation methodologies to efficiently process even very large-scale data acquisitions (i.e., up to 10 hours of flight time) contiguously, further work is required in this regard. Processing of large-scale multi-flight data acquisitions, possibly with multiple different platforms and at different times, provides further potential for improving the quality of georeferencing. However, efficient processing of such datasets will require scalability not only temporally (achieved e.g., by IMU down-sampling as demonstrated here) but also spatially (possibly through domain decomposition, or by aggregating LiDAR measurements).

In summary, joint trajectory and strip adjustment with multiple acquisitions, platforms, and sensors is a promising avenue towards increased accuracy, precision, and reliability of LiDAR georeferencing - if processing can be done efficiently.

#### CRedit authorship contribution statement

**Florian Pöppl:** Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Data curation, Conceptualization, Writing – review & editing. **Andreas Ullrich:** Writing – review & editing, Software, Resources, Data curation. **Gottfried Mandlbauer:** Writing – review & editing, Supervision, Resources, Data curation. **Norbert Pfeifer:** Writing – review & editing, Supervision, Resources, Project administration, Funding acquisition.

#### Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: This work was carried out as part of the project ZAP-ALS (883660), a joint project between the Technische Universität Wien and the company RIEGL, which is funded by the Austrian Research Promotion Agency (FFG). Dr. Andreas Ullrich is the Technical Director of RIEGL Laser Measurement Systems GmbH. Florian Pöppl consults for RIEGL on selected technical issues related to satellite and inertial navigation but outside the scope of this work.

#### Acknowledgments

This work was carried out as part of the project ZAP-ALS (883660), a joint project between the *Technische Universität Wien* and the company *RIEGL*, which is funded by the Austrian Research Promotion Agency (FFG)<sup>6</sup>. The authors also acknowledge TU Wien Bibliothek for financial support through its Open Access Funding Program.

#### References

- Agarwal, S., Mierle, K., The Ceres Solver Team, 2022. Ceres solver.
- Al-Jailaty, H., Mansour, M.M., 2021. Efficient attitude estimators: A tutorial and survey. *J. Signal Process. Syst.* <http://dx.doi.org/10.1007/s11265-020-01620-4>.
- Alam, M., Rohac, J., 2015. Adaptive data filtering of inertial sensors with variable bandwidth. *Sensors* 15 (2), 3282–3298. <http://dx.doi.org/10.3390/s150203282>.
- Brun, A., Cucci, D.A., Skalous, J., 2022. LiDAR point-to-point correspondences for rigorous registration of kinematic scanning in dynamic networks. *ISPRS J. Photogramm. Remote Sens.* 189, 185–200. <http://dx.doi.org/10.1016/j.isprsjprs.2022.04.027>.
- Cucci, D.A., Skalous, J., 2019. On raw inertial measurements in dynamic networks. *ISPRS Ann. Photogramm. Remote. Sens. Spat. Inf. Sci.* IV-2/W5, 549–557. <http://dx.doi.org/10.5194/isprs-annals-IV-2-W5-549-2019>.
- Farrell, J., 2008. Aided navigation: GPS with high rate sensors. *Electronic Engineering*, McGraw-Hill, New York.
- Farrell, J.A., Silva, F.O., Rahman, F., Wendel, J., 2022. Inertial measurement unit error modeling tutorial: Inertial navigation system state estimation with real-time sensor calibration. *IEEE Control. Syst.* 42 (6), 40–66. <http://dx.doi.org/10.1109/MCS.2022.3209059>.
- Forster, C., Carlone, L., Dellaert, F., Scaramuzza, D., 2015. IMU preintegration on manifold for efficient visual-inertial maximum-a-posteriori estimation. In: *Robotics: Science and Systems XI*. Robotics: Science and Systems Foundation, <http://dx.doi.org/10.15607/RSS.2015.XI.006>.
- Glira, P., Pfeifer, N., Mandlbauer, G., 2016. Rigorous strip adjustment of UAV-based laserscanning data including time-dependent correction of trajectory errors. *Photogramm. Eng. Remote Sens.* 82 (12), 945–954. <http://dx.doi.org/10.14358/PERS.82.12.945>.
- Groves, P.D., 2013. Principles of GNSS, inertial, and multisensor integrated navigation systems, second ed. *GNSS Technology and Application Series*, Artech House, Boston.
- Hexagon, 2023. Inertial Navigation Systems and Vibration. Technical Report.
- Huber, P.J., 1964. Robust estimation of a location parameter. *Ann. Math. Stat.* 35 (1), 73–101. <http://dx.doi.org/10.1214/aoms/1177703732>.
- Jonassen, V.O., Kjörsvik, N.S., Gjevestad, J.G.O., 2023. Scalable hybrid adjustment of images and LiDAR point clouds. *ISPRS J. Photogramm. Remote Sens.* 202, 652–662. <http://dx.doi.org/10.1016/j.isprsjprs.2023.07.007>.
- Kim, M.-J., Kim, M.-S., Shin, S.Y., 1995. A general construction scheme for unit quaternion curves with simple high order derivatives. In: *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques - SIGGRAPH '95*. ACM Press, pp. 369–376. <http://dx.doi.org/10.1145/218380.218486>.
- Le Gentil, C., Vidal-Calleja, T., 2021. Continuous integration over SO(3) for IMU preintegration. In: *Robotics: Science and Systems XVII*. Robotics: Science and Systems Foundation, p. 9. <http://dx.doi.org/10.15607/RSS.2021.XVII.078>.
- Li, Z., Yan, Y., Jing, Y., Zhao, S.G., 2015. The design and testing of a LiDAR platform for a UAV for heritage mapping. *Int. Arch. Photogramm. Remote. Sens. Spat. Inf. Sci.* XL-1/W4, 17–24. <http://dx.doi.org/10.5194/isprarchives-XL-1-W4-17-2015>.
- Li, J., Yuan, S., Cao, M., Nguyen, T.-M., Cao, K., Xie, L., 2024. HCTO: Optimality-aware LiDAR inertial odometry with hybrid continuous time optimization for compact wearable mapping system. *ISPRS J. Photogramm. Remote Sens.* 211, 228–243. <http://dx.doi.org/10.1016/j.isprsjprs.2024.04.004>.

<sup>6</sup> Österreichische Forschungsförderungsgesellschaft (FFG), Wien, Austria, [www.ffg.at](http://www.ffg.at).

- Liu, J., Liu, Z., Chen, W., Lv, J., Jiang, Z., Pang, J., Du, L., 2024. Response spectrum-based analysis of airborne radar random vibration and multi-point control improvement. *Sci. Rep.* 14 (1), 7565. <http://dx.doi.org/10.1038/s41598-024-56783-y>.
- Lupton, T., Sukkarieh, S., 2009. Efficient integration of inertial observations into visual SLAM without initialization. In: 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE, St. Louis, MO, USA, pp. 1547–1552. <http://dx.doi.org/10.1109/IROS.2009.5354267>.
- Ma, H., Wu, J., 2012. Analysis of positioning errors caused by platform vibration of airborne LiDAR system. In: 2012 8th IEEE International Symposium on Instrumentation and Control Technology (ISICT) Proceedings. IEEE, London, United Kingdom, pp. 257–261. <http://dx.doi.org/10.1109/ISICT.2012.6291650>.
- Periu, C., Mohsenimanesh, A., Laguë, C., McLaughlin, N., 2013. Isolation of vibrations transmitted to a LIDAR sensor mounted on an agricultural vehicle to improve obstacle detection. *Can. Biosyst. Eng.* 55 (1), 2.33–2.42. <http://dx.doi.org/10.7451/CBE.2013.55.2.33>.
- Pöpl, F., Teufelsbauer, H., Ullrich, A., Pfeifer, N., 2023. Mobile laser scanning with low-cost navigation sensors: Compensating for low-grade IMU with dual-GNSS and tightly-coupled LiDAR. In: The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences. Vol. XLVIII-1/W1-2023, pp. 403–410. <http://dx.doi.org/10.5194/isprs-archives-XLVIII-1-W1-2023-403-2023>.
- Pöpl, F., Ullrich, A., Mandlbürger, G., Pfeifer, N., 2024. A flexible trajectory estimation methodology for kinematic laser scanning. *ISPRS J. Photogramm. Remote Sens.* 215, 62–79. <http://dx.doi.org/10.1016/j.isprsjprs.2024.06.014>.
- Rouzaud, D., Skalous, J., 2011. Rigorous integration of inertial navigation with optical sensors by dynamic networks. *Navigation* 58 (2), 141–152. <http://dx.doi.org/10.1002/j.2161-4296.2011.tb01797.x>.
- Savage, P.G., 1998a. Strapdown inertial navigation integration algorithm design part 1: attitude algorithms. *J. Guid. Control Dyn.* 21 (1), 19–28. <http://dx.doi.org/10.2514/2.4228>.
- Savage, P.G., 1998b. Strapdown inertial navigation integration algorithm design part 2: velocity and position algorithms. *J. Guid. Control Dyn.* 21 (2), 208–221. <http://dx.doi.org/10.2514/2.4242>.
- SBG Systems, 2024. How to handle vibrations? <https://www.sbg-systems.com/support/technology/how-to-handle-vibrations/>.
- Schlager, B., Goelles, T., Behmer, M., Muckenhuber, S., Payer, J., Watzenig, D., 2022. Automotive Lidar and vibration: Resonance, inertial measurement unit, and effects on the point cloud. *IEEE Open J. Intell. Transp. Syst.* 3, 426–434. <http://dx.doi.org/10.1109/OJITS.2022.3176471>.
- Suwandi, B., Kitasuka, T., Aritsugi, M., 2019. Vehicle vibration error compensation on IMU-accelerometer sensor using adaptive filter and low-pass filter approaches. *J. Inf. Process.* 27, 33–40. <http://dx.doi.org/10.2197/ipsjip.27.33>.
- Wei, Y., Li, Y., 2022. Impact of sensor data sampling rate in GNSS/INS integrated navigation with various sensor grades. *Int. Arch. Photogramm. Remote. Sens. Spat. Inf. Sci.* XLVI-3/W1-2022, 205–211. <http://dx.doi.org/10.5194/isprs-archives-XLVI-3-W1-2022-205-2022>.