

ON RAW INERTIAL MEASUREMENTS IN DYNAMIC NETWORKS

D. A. Cucci^{1,*}, J. Skaloud¹

¹ Geodetic Engineering Laboratory, EPFL, Bâtiment GC - Station 18 - Lausanne, Switzerland - (davide.cucci, jan.skaloud)@epfl.ch

Commission I, WG I/9

KEY WORDS: Dynamic Networks, IMU-preintegration, Inertial/GNSS navigation

ABSTRACT:

Dynamic Networks have been introduced in the literature to solve multi-sensor fusion problems for navigation and mapping. They have been shown to outperform conventional methods in challenging scenarios, such as corridor mapping or self-calibration. In this work we investigate the problem of how raw inertial readings can be fused with GNSS position observations in Dynamic Networks (DN) with the goal of i) limiting the number of unknowns in the estimation problem and ii) improving the conditioning of the normal equations arising in least-squares adjustments in the absence of spatial constraints (e.g., image observations). For that we propose a modified version of the well known IMU-preintegration method, accounting for a non-constant gravity model, the Earth rotation and the apparent Coriolis force, and we compare it with the conventional DN formulation in a emulated scenario. This consists of a fixed-wing UAV flying four times over a 2 Km long corridor.

INTRODUCTION

In navigation, we would like to determine the position and the orientation (pose, hereafter) of a moving rigid object, represented by a body frame, with respect to a reference one. Nowadays, the most popular solution in outdoor scenarios is based on the fusion of inertial measurements, coming from one or multiple triads of accelerometers and gyroscopes, with Global Navigation Satellite Systems (GNSS) position and (possibly) velocity observations. Inertial/GNSS fusion has been extensively studied in the last sixty years and well established solutions exist based on recursive Bayesian estimation and extended Kalman filters (Titterton and Weston, 1997).

In many applications the solution of the navigation problem is only an intermediate step. For instance, in direct integrated sensor orientation, the pose estimates obtained with inertial/GNSS fusion are used to assign geographic locations to the observations of another sensor, e.g., a LiDAR, while another adjustment follows to determine other parameters related to system calibration (Skaloud and Lichti, 2006) or relative alignment between flight lines. As another example, modern Assisted Aerial Triangulation methods (AAT) (Blázquez and Colomina, 2012) rely on prior knowledge of the camera poses to eliminate or reduce the number of Ground Control Points (GCPs) needed for globally accurate geo-referencing or to de-correlate camera internal and external parameters.

In general, inertial/GNSS fusion is handled as an initial step to obtain a preliminary navigation solution, which is then used, as it is or further adjusted, in subsequent processing steps. For instance, in AAT, the poses determined by inertial/GNSS fusion are used, together with their assumed uncertainty, as prior constraints on the camera exterior parameters in the bundle adjustment. While this approach has proven successful in most cases, it is sub-optimal, and fails in challenging scenarios, for instance when neither the inertial/GNSS observations, or the camera ones, taken alone, are sufficient to obtain a globally accurate solution, as in corridor mapping or when the GNSS signals reception is imperfect. A reason for this is that, in

the classical AAT pipeline the following is neglected: i) the non-linear dependency between the inertial/GNSS poses and the systematic errors, such as inertial sensors biases or those due to imperfect initial alignment, and ii) the time correlation of body frame poses. One consequence is that in geometrically weak AAT it is difficult to correct for a general inertial/GNSS trajectory, e.g., because of poor initial alignment or incorrect inertial biases determination, as discussed in (Cucci et al., 2017b)

Such limitations motivated the development of methods for the joint adjustment of GNSS and inertial observations together with other sensors of interest, e.g., cameras and LiDARs. One such method goes under the name of Dynamic Networks (DNs), in the geodesy literature (Colomina and Blázquez, 2004, Glennie et al., 2011), or pose-graph optimization, in robotics, e.g., see, among the others, (Indelman et al., 2013, Cucci and Matteucci, 2014). In DN, all the unknown poses are determined by means of a single least-squares adjustment step (differently from recursive estimation schemes) minimizing the residuals associated to all the available measurements from all the available sensors, being those GNSS receivers, inertial sensors, cameras, LiDARs, etc. The main advantages of DN are: i) all the different measurements are processed simultaneously, ii) no marginalization of variables is ever performed, as in recursive estimation, and the full non-linear observation models are always employed, iii) in contrast with Kalman filtering/smoothing, spatial constraints can be easily expressed as functions of multiple poses at different time instants. Such advantages come at the price of severely increased computational complexity, especially when dealing with high rate sensors, such as modern Inertial Measurement Units (IMUs).

In this work we discuss how raw inertial observations can be handled, together with GNSS ones, in Dynamic Networks, or pose-graph optimization. At first, in Section 2. and 3., we will review the state of the art method, in which each single inertial observation it is used to form an observation model constraining multiple successive poses. This method is rather straightforward and it has been successfully employed, e.g., see (Cucci et al., 2017b), to show that the DN formulation is superior to the cascade processing of inertial/GNSS, first, and then image observations in AAT. However, the state of the art

*Corresponding author

method suffer of two issues that become critical as the rate of inertial observation increases. These are the number of unknowns in the estimation problem, and the ill-conditioning of the normal equations involved in the solution of the least-squares adjustment.

In Section 4. we present an alternative technique called *IMU-preintegration*, first proposed in (Lupton and Sukkarieh, 2009) and now well established in the robotics literature, mainly in visual-inertial navigation systems. This technique consists in computing the integral of sets of raw inertial readings during a pre-processing step. This effectively reduces the number of unknowns to be estimated in DN and improves the conditioning number of the normal equations, without affecting the accuracy of the solution. We extend the formulation available in the literature, e.g., in (Forster et al., 2017), accounting for certain effects which are normally neglected in robotics, or with low-grade IMUs, while they are significant for longer/larger/faster missions or more precise IMUs: a position dependent gravity model, the Earth rotation and the Coriolis force.

In Section 6. we compare the two presented methodologies in a emulated scenario, considering a real world fixed-wing UAV flight as a model to generate synthetic inertial and GNSS observations at different IMU rates. The RMS of the position and orientation error is then compared for different setups of the DN estimation and the IMU pre-integration step. The results suggest that IMU pre-integration can safely substitute the conventional method in the considered case. However, we notice some issues in estimating inertial biases in case too many inertial readings are integrated. We conclude and present some direction for further investigations in Section 7..

NOTATION

In the derivations presented in this work we will often need to mix between continuous time quantities and discrete samples of these. A short preamble is therefore appropriate to clarify the notation. We first define a function $t(j) : \mathbb{N} \rightarrow \mathbb{R}$, abbreviated with t_j which gives the timestamp of the j -th sample of a given quantity. As another shortcut, $\Delta t \doteq t_j - t_{j-1}$. Continuous time quantities will be denoted without index, e.g., r^e is the continuous time position of the body frame. If needed, the dependency with respect to time will be made explicit, e.g., $r^e(\tau)$. The j -th sample of r^e will be denoted using a subscript on the left: ${}_j r^e = r^e(t_j)$.

INERTIAL OBSERVATION MODELS IN E-FRAME

Here we recall the well known inertial observation models in the Earth-centered, Earth-fixed frame (ECEF), e hereafter, in continuous time. Accelerometers and gyroscopes are assumed to yield measurements with respect to a quasi-inertial, geocentric frame i with celestial (space-fixed) orientation. The e frame is rotating with respect to i with angular velocity ω_{ie}^e , and inertial sensors measurements are expressed in body, or IMU, frame b . thus the observation models for the gyroscope and accelerometer readings, z_ω and z_f , are:

$$z_\omega = \omega_{ib}^b + b_\omega + \xi_\omega = \left[R_b^{eT} \omega_{ie}^e \wedge R_b^e \right]^\vee + \omega_{eb}^b, \quad (1)$$

$$z_f = f^b + bf + \xi_f = R_b^{eT} (a^e + 2\omega_{ie}^e \wedge v^e - g^e(r^e)), \quad (2)$$

where all the symbols are defined as follows:

1. ω_{bc}^a is the angular velocity of frame c with respect to frame b , expressed in frame a ,

2. the operator $(\cdot)^\wedge$ transforms a 3D vector into the corresponding skew-symmetric matrix, i.e.:

$$\omega^\wedge = \Omega = \begin{bmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{bmatrix}. \quad (3)$$

Its inverse is the $(\cdot)^\vee$ operator.

3. f^b is the specific force in body frame,
4. b_ω and bf are two bias terms whose stochastic characterization depends on the specific IMU at hand,
5. ξ_ω and ξ_f are two white noise terms,
6. r^e is the position of the body frame with respect to e ,
7. v^e and a^e are the linear velocity and acceleration of b with respect to e , so that $v^e \doteq \dot{r}^e$, $a^e \doteq \dot{v}^e$,
8. R_b^a is the orientation of the frame b with respect to a , so that $r^a = R_b^a r^b$.
9. $g^e(r^e)$ is a gravity model, defined as a function of r^e , since in general the gravity vector depends on the body frame position with respect to e . In this work we use the following one:

$$g^e(r^e) = \tilde{g} \frac{\nabla s}{\|\nabla s\|} - (\omega_e^i)^\wedge r^e, \\ s = r^{eT} \text{diag}([a, a, b]) r^e - 1 \quad (4)$$

In other words, we consider a nominal gravity vector of magnitude \tilde{g} , corrected for the centripetal acceleration due to Earth rotation, and orthogonal to a given ellipsoid, e.g., the WGS-84 one, defined in terms of its semi-major and semi-minor axes a and b . In Equation 4, ∇ is the gradient operator. More complex gravity models could be also employed, e.g., see (Heiskanen and Moritz, 1981).

DISCRETE INERTIAL OBSERVATION MODELS IN DYNAMIC NETWORKS

In Dynamic Networks, the goal is to determine the navigation solution (i.e., the trajectory), and, optionally, other parameters, such as inertial sensor biases, fusing all the available sensor readings in a single step. In this work we focus on the fusion of raw inertial and GNSS position observations. The problem formulation of DN allows to fuse raw inertial observation with other sensors, such as, e.g., cameras and laser scanners.

A vector of unknowns is composed of positions and orientations of the body frame at specific, discrete, time instants t_j . These are chosen to be the ones of the IMU readings, which are assumed to be synchronized with GNSS positions and equally spaced in time. Additional unknowns may be included to model inertial bias terms. For the sake of simplicity, in this work we restrict to two time-invariant bias vectors, one for the accelerometers triad, bf , and one for the gyroscopes one, b_ω . Refer to (Cucci et al., 2017a) for more complex noise models in DN. More precisely, the vector of unknowns is defined as

$$\mathbb{X} = [{}_1 r^e, {}_1 R_b^e, {}_2 r^e, {}_2 R_b^e, \dots, {}_N r^e, {}_N R_b^e, bf, b_\omega]. \quad (5)$$

The IMU readings are related to the unknown body frame positions and orientations in \mathbb{X} replacing the continuous time

derivatives with finite differences in Equation 1 and 2:

$${}_jv^e \doteq \frac{{}_jr^e - {}_{j-1}r^e}{\Delta t}, \quad (6)$$

$${}_ja^e \doteq \frac{{}_jv^e - {}_{j-1}v^e}{\Delta t} = \frac{{}_jr^e - 2{}__{j-1}r^e + {}_{j-2}r^e}{\Delta t^2}, \quad (7)$$

$${}_j\omega_{eb}^e \doteq \frac{1}{\Delta t} \log \left[{}_{j-1}R_b^{eT} {}_jR_b^e \right]^\vee. \quad (8)$$

Higher order finite differencing schemes can be used, e.g., see (Bruton et al., 1999), yet the proposed ones suffice in most cases. In Equation 8, $\log(\cdot)$ is the logarithmic map in the special orthonormal group $SO(3)$: see (Strasdat, 2012, Chapter 2.4) for an excellent introduction of Lie Groups for navigation and mapping.

For each inertial reading, and for each GNSS position observation we can formulate one instance of the appropriate observation model. An estimate for \mathbb{X} can then be obtained minimizing the residuals associated to each observation model in least-squares sense. This approach is well known and the reader can refer to (Colomina and Blázquez, 2004, Glennie et al., 2011, Cucci et al., 2017b).

Two severe issues arise when the rate of IMU readings increases, which are discussed in the following.

Number of unknowns

To fuse all the available sensor readings in a single adjustment step, all the positions and orientations required to formulate the observation models need to be available in \mathbb{X} . In modern formulations based on Lie Groups, e.g., (Strasdat, 2012), one body frame pose is represented by a 3D position vector and a 3×3 rotation matrix, totaling twelve unknowns. Since modern inertial sensors yield readings up to 1 – 2 KHz, this results in millions of unknowns to be determined simultaneously for navigation problems of typical size. Such large estimation problems do not pose severe issues for offline processing, provided that efficient linear algebra methods are used to fully exploit the sparsity of the problem, e.g., (Davis, 2006). However, real time (or close to real time) processing becomes challenging and sub-optimal solution, such as windowing schemes (Cucci, 2014) need to be used.

Ill-conditioning of the normal equations

Consider the following simplified 1D example: we would like to determine the x coordinate of a moving body fusing readings from an unbiased accelerometer and a GNSS receiver. According to the Dynamic Network approach, the unknowns are $\mathbb{X} = [{}_1x, \dots, {}_jx, \dots, {}_Nx]$, with $t(j)$ being the timestamp of accelerometer readings. We can formulate Equation 2 for the accelerometer readings:

$${}_jz_f + \xi_f = \frac{1}{\Delta t^2} ({}_{j-1}x - 2{}_jx + {}_{j+1}x), \quad \forall j \in \{2, \dots, N-1\}, \quad (9)$$

where ${}_jz_f$ is the accelerometer reading at t_j , for simplicity $\xi_f \sim \mathcal{N}(0, \sigma_f^2/\Delta t)$ and σ_f^2 is power spectral density of the accelerometer white noise. Suppose now that two position observations are available at $t_1 = 0$ and $t_N = 1$. The corresponding model is

$${}_iz_x + \xi_x = {}_ix, \quad \forall i \in \{1, N\}, \quad (10)$$

where ${}_iz_x$ is the position reading at time t_i and $\xi_x \sim \mathcal{N}(0, \sigma_x^2)$.

Given the timestamps of position readings, the number of unknowns can be related to the accelerometer rate:




Figure 1. Conditioning number of $A^T \Omega A$ as a function of $1/\Delta t$ and $r = \sigma_x/\sigma_f$.

$N = 1/\Delta t$, meaning that a higher rate accelerometer requires the determination of more positions, as discussed in Section 3.1. We have $N - 2$ observations from accelerometer readings and 2 observations from the GNSS. This means that this problem has no redundancy, yet a solution always exists and it is unique.

There are different methods to solve this estimation problem, however, we would like to follow the Dynamic Network approach: the observation models in Equation 9 and 10 are grouped together in the form $A\mathbb{X} = b$ and the estimates for ${}_jx$ are obtained in least-squares sense solving the normal equations $(A^T \Omega A)\mathbb{X} = A^T \Omega b$, where Ω is a weight matrix encoding the uncertainty of the sensor readings.

The presented problem always admits a unique solution, however it might prove difficult to determine it numerically with the DN approach: it can be shown that the conditioning number of $H = A^T \Omega A$ is a function of Δt and of σ_x/σ_f . More precisely, the higher is the rate of the accelerometer, and the lower is the standard deviation of the white noise affecting accelerometer readings, the more the conditioning number increases. This is depicted in Figure 1. The conditioning number of H encodes the stability of the solution with respect to small changes in the observation vector, b . Note that $\text{cond}(\lambda A) = \text{cond}(A)$, so the absolute values of σ_x and σ_f are not important, only their ratio counts.

The empirical conclusions that can be derived from this simplified example are well known in numerical analysis, e.g., see (Grossmann et al., 2007, Chapter 4.9.1), and can be extended to the general 3D inertial/GNSS fusion case. Such considerations suggest that any DN solver employing the formulation in Section 3. will break when a sufficiently high rate or precise IMU is considered, as the conditioning number of H increases. The actual breaking point of a specific algorithm depends on the technique employed to solve the normal equations and its sensibility to high conditioning numbers.

One solution for this issue is given by regularization. For instance, in Tikhonov regularization \mathbb{X} is obtained solving $(A^T \Omega A + \lambda I)\mathbb{X} = A^T \Omega b$. This solves any ill-conditioning problem in $A^T \Omega A$ at the price of introducing a bias in the solution. Such bias can be made acceptable by choosing an appropriate value for λ , e.g., according to (Hansen and OLeary, 1993).

IMU PRE-INTEGRATION

In this section we present a technique known as *IMU pre-integration*. This mitigates the issues arising with higher rate IMUs by reducing the number of the unknowns and by increasing the time spacing between them. The core idea is to compute certain integrals of groups of accelerometer and gyroscope readings with respect to time during a pre-processing step, before solving the DN. Such integrals form a new type of observables available at a substantially lower frequency with respect to the original IMU readings, but incorporating all the information available from the inertial sensors. Rigorous observation models are derived for those observables to be used in Dynamic Networks. This allows to exploit the full potential of high rate IMUs without incurring in the limitation of bloating number of unknowns or ill-conditioning of the normal equations.

IMU pre-integration is well known in robotics literature: it was first introduced in (Lupton and Sukkariéh, 2009, Lupton and Sukkariéh, 2012) and it is now at the core of modern visual-inertial systems, e.g., (Forster et al., 2017). Our formulation extends this approach in two senses:

1. the poses are determined with respect to the e frame, considering the Earth rotation, the apparent Coriolis force and a position dependent gravity model,
2. differently from (Lupton and Sukkariéh, 2012), we do not estimate body frame velocities explicitly, further reducing the number of unknowns in \mathbb{X} .

In the following, we will first develop the observation models for pre-integrated IMU readings. More precisely, we will introduce three quantities, ΔR , Δr_+ and Δr_- for which an estimate can be obtained computing the integral of raw accelerometer and gyroscope readings in a pre-processing step. For such quantities, we will rigorously derive the observation models that relate those to the body frame poses available in \mathbb{X} . Later on, we will discuss the details of how the pre-integration step is performed.

We assume that IMU readings are available at $t(j)$, and then groups of M IMU samples are integrated with respect to time as discussed later on in Section 5., so that, differently with respect to Equation 5, we need to maintain body frame poses only at $t(k)$, with $k \in \{1, M, 2M, \dots\}$. M can range from tents to hundred of samples, and its effect is discussed in Section 5.1.

Observation models for pre-integrated gyroscope readings

In the following we relate the integral of M raw gyroscope readings to two successive orientations available in \mathbb{X} , ${}_k R_b^e$ and ${}_{k+M} R_b^e$.

We first decompose R_b^e in $R_i^e R_b^i$. Since the angular velocity of e with respect to i is assumed to be constant for the purpose of navigation,

$$R_e^i(t_{k+M}) = R_e^i(t_k) \exp\left(\omega_{ie}^i \wedge \Delta t\right). \quad (11)$$

Regarding R_b^i , if we neglect the effects of coning (Savage, 1998), we have that

$$R_b^i(t_{k+M}) \approx R_b^i(t_k) \exp\left(\int_{t_k}^{t_{k+M}} \omega_{ib}^b(\tau) \wedge d\tau\right). \quad (12)$$

Expanding $R_b^e = R_i^e R_b^i$ using Equation 11 and 12 we obtain

$${}_{k+M} R_b^e = \exp\left(-\omega_{ie}^i \wedge \Delta t\right) {}_k R_b^e \exp\left(\underbrace{\int_{t_k}^{t_{k+M}} \omega_{ib}^b(\tau) \wedge d\tau}_{{}_k \Delta R}\right). \quad (13)$$

The highlighted term ${}_k \Delta R$ involves an integral of the angular velocity of the body frame with respect to the inertial frame. An estimate for ${}_k \Delta R$, ${}_k \hat{\Delta R}$, can be computed integrating gyroscope measurements, ${}_j z_\omega$, between t_k and t_{k+M} , provided that any bias present in the measurements is accounted for, as we will discuss in Section 5.. Equation 13 can be rearranged to obtain the desired observation model for ${}_k \hat{\Delta R}$:

$${}_k \hat{\Delta R} + \xi_{\Delta R} \wedge = {}_k R_b^{eT} \exp\left(\omega_{ie}^i \wedge \Delta t\right) {}_{k+M} R_b^e. \quad (14)$$

Equation 14 encodes the comparison between two rotation matrices, so it is rewritten in form of a vector as

$$\log\left[{}_k \hat{\Delta R}^T {}_k R_b^{eT} \exp\left(\omega_{ie}^i \wedge \Delta t\right) {}_{k+M} R_b^e\right]^\vee = 0_{3 \times 1} + \xi_{\Delta R} \quad (15)$$

Note the similarity between Equation 15 and the observation models for *relative orientation control* introduced in (Blázquez and Colomina, 2012).

Observation models for pre-integrated accelerometer readings

In the following, we will relate two different integrals of M accelerometer readings to three successive positions and one orientation available in \mathbb{X} , ${}_{k+h} r^e$, with $h \in \{-M, 0, M\}$ and ${}_k R_b^e$.

First note that, if we neglect the effects of sculling (Savage, 1998), the two following relations hold since the velocity $v^e \doteq \dot{r}^e$ and the linear acceleration $a^e \doteq \dot{r}^e$:

$$\begin{aligned} {}_{k+M} r^e &= {}_k r^e + M \Delta t {}_k v^e + \iint_{t_k}^{t_{k+M}} a^e(\tau) d\tau, \\ {}_{k-M} r^e &= {}_k r^e - M \Delta t {}_k v^e + \iint_{t_k}^{t_{k-M}} a^e(\tau) d\tau. \end{aligned}$$

Summing member-wise, we obtain

$${}_{k+M} r^e - 2 {}_k r^e + {}_{k-M} r^e = \iint_{t_k}^{t_{k+M}} a^e(\tau) d\tau + \iint_{t_k}^{t_{k-M}} a^e(\tau) d\tau. \quad (16)$$

Recalling Equation 2, we can expand the first of the two double integrals as the sum of three terms:

$$\iint_{t_k}^{t_{k+M}} a^e(\tau) d\tau = A + B + C, \quad (17)$$

where:

$$A = \int_{t_k}^{t_{k+M}} R_b^e(\tau) f^b(\tau) d\tau$$

$$= \exp\left(-\omega_{ie}^i \wedge \Delta t\right) \underbrace{k R_b^e \int_{t_k}^{t_{k+M}} \exp\left(\int_{t_k}^{\tau} \omega_{ib}^b(v) \wedge dv\right) f^b(\tau) d\tau}_{k \Delta r_+} \quad (18)$$

This has been obtained by decomposing R_b^e in $R_i^e R_b^i$ and proceeding as for Equation 13. An estimate for $k \Delta r_+$, $k \hat{\Delta} r_+$, can be computed integrating gyroscope and accelerometer measurements between t_k and t_{k+M} , provided that any bias present in the measurements is accounted for, as we will discuss in Section 5..

For B , we have that

$$B = - \int_{t_k}^{t_{k+M}} 2\omega_{ie}^e \wedge v^e(\tau) d\tau \approx 2\omega_{ie}^e \wedge k v^e \frac{(M\Delta t)^2}{2}, \quad (19)$$

where the continuous body frame velocity v^e has been replaced with $k v^e$, obtained finite differencing body frame positions as in Equation 6. This corresponds to the *mean* velocity between t_k and t_{k+M} . This approximation is tolerable for sufficiently small M since ω_{ie}^e is small.

Finally, for C we have that

$$C = \int_{t_k}^{t_{k+M}} g^e(r^e(\tau)) d\tau \approx g^e(k r^e) \frac{(M\Delta t)^2}{2}, \quad (20)$$

where the continuous body frame position r^e has been replaced with its sample at time k , $k r^e$. This is acceptable as $\partial g^e / \partial r^e \approx 0$.

We can go through the same reasoning for the second of the double integrals in Equation 16, expanding it as done for the first in Equation 18, 19 and 20. Expanding both double integrals in Equation 16, we obtain the observation model for pre-integrated accelerometer readings:

$$k \hat{\Delta} r_+ + k \hat{\Delta} r_- + \xi_{\Delta r_+} + \xi_{\Delta r_-} =$$

$$= k R_b^{eT} \exp\left(\omega_{ie}^i \wedge M\Delta t\right) \left[k_{+M} r^e - 2 k r^e + k_{-M} r^e + \right.$$

$$\left. + (2\omega_{ie}^e \wedge k v^e - g^e(k r^e)) (M\Delta t)^2 \right]. \quad (21)$$

COMPUTING DELTAS

In this section we present an algorithm to obtain $k \hat{\Delta} R$, $k \hat{\Delta} r_+$ and $k \hat{\Delta} r_-$, as well as the covariance matrices for the respective noise terms $\xi_{\Delta R}$, $\xi_{\Delta r_+}$ and $\xi_{\Delta r_-}$. The original quantities for which we look for an estimate are reported in the following for the sake of clarity:

$$k \Delta r_+ = \int_{t_k}^{t_{k+M}} \exp\left(\int_{t_k}^{\tau} \omega_{ib}^b(v) \wedge dv\right) f^b(\tau) d\tau$$

$$k \Delta r_- = \int_{t_k}^{t_{k-M}} \exp\left(\int_{t_k}^{\tau} \omega_{ib}^b(v) \wedge dv\right) f^b(\tau) d\tau$$

$$k \Delta R = \exp\left(\int_{t_k}^{t_{k+M}} \omega_{ib}^b(\tau) \wedge d\tau\right) \quad (22)$$

We will obtain estimates for such quantities by numerical integration of raw accelerometer and gyroscope readings, employing the Runge-Kutta second order recursive integration scheme described in the following. Recall from Section 4. that inertial readings are available at times t_j , with $j \in \{1, \dots, N\}$, but, if pre-integration is used, in the Dynamic Network state vector, \mathbb{X} , poses are estimated only at times t_k , with $k \in \{1, M, 2M, \dots\}$. Therefore, raw inertial readings at t_j , with $j \in \{k, \dots, k+M-1\}$ will be integrated to compute $k \Delta r_+$ and $k \Delta R$, while $k \Delta r_-$ will be computed using $j \in \{k-1, \dots, k-M\}$ (note the reversed order). In the following we present the full algorithm to evaluate $k \Delta r_+$, $k \hat{\Delta} R$ and the covariance matrices for the noise terms $\xi_{\Delta r_+}$ and $\xi_{\Delta R}$. Regarding $k \hat{\Delta} r_-$ and $\xi_{\Delta r_-}$, the process is similar and it is left to the reader. Let

$$j X = [j \Delta r_+, j \Delta v_+, j \Delta R, j b f, j b \omega], \quad (23)$$

$$k X = [0_{3 \times 1}, 0_{3 \times 1}, I_{3 \times 3}, \overline{b f}, \overline{b \omega}], \quad (24)$$

be a state vector at time j and its initial value for $j = k$, where $\overline{b f}$, $\overline{b \omega}$ are the best known values for the inertial biases available at time k . Note that such estimates may not be available, in which case $0_{3 \times 1}$ is used, see Section 5.1 later on. For each time j with $j \in \{k, \dots, k+M-1\}$ we employ the accelerometer and gyroscope readings, $j z_f$, and $j z_\omega$ to evaluate the state transition function \mathcal{F} :

$$j_{+1} X = \mathcal{F}(j X) : \begin{cases} j \Delta r_+ + j \Delta v_+ \Delta t + j \Delta R (j z_f - j b f) \frac{\Delta t^2}{2} \\ j \Delta R (j z_f - j b f) \Delta t \\ j \Delta R \exp[(j z_\omega - j b \omega) \wedge] \Delta t \\ j b f \\ j b \omega \end{cases} \quad (25)$$

Note the similarity between \mathcal{F} and the process models conventionally used in Kalman filters for GNSS/Inertial integration (biases are assumed constant in the integration period). The main difference is that no gravity or Earth rotation compensation is done in Equation 25 and raw specific force and angular velocity readings are integrated as prescribed by Equation 22. The compensation for those terms is done rigorously at observation model level, as discussed in the previous section, see in Equation 21. After M steps, the estimates for $k \Delta r_+$ and $k \Delta R$, $k \hat{\Delta} r_+$ and $k \hat{\Delta} R$ can be taken from the corresponding components of $k_{+M} X$.

The covariance matrices for $\xi_{\Delta r_+}$ and $\xi_{\Delta R}$ is obtained as a function of the assumed variance of inertial readings by means of linear covariance propagation. Let $j F$ and $j G$ be

$$j F = \frac{\partial \mathcal{F}(j X)}{\partial j X}, \quad j G = \frac{\partial \mathcal{F}(j X)}{\partial [j z_f, j z_\omega]}, \quad (26)$$

and the following recursive transition function and initial conditions:

$$k R = 0_{15 \times 15}, \quad (27)$$

$$j_{+1} R = j F j R j F^T + j G Q j G^T, \quad (28)$$

where Q is a 6×6 , time-invariant, diagonal matrix composed with the assumed variances of the accelerometer and gyroscope readings. The covariance matrices for $\xi_{\Delta r_+}$ and $\xi_{\Delta R}$ are obtained from the proper diagonal blocks of $k_{+M} R$.




Figure 2. The UAV flight used to generate synthetic sensor observations. This corresponds to a real-world corridor mapping mission and includes specific eight-shaped maneuvers that introduce the necessary dynamics to enhance the observability of inertial biases.

Estimating Biases

In Equation 22, ΔR , Δr_+ and ΔR_- have been defined as a function of angular velocities and specific forces. However, inertial sensors yield only biased measurements of such quantities. The state transition function \mathcal{F} accounts for such biases (see Equation 25), yet up to now we have not clarified how the values for \bar{bf} and $\bar{b\omega}$ can be obtained, as required in the initialization of X , see Equation 24.

One way to proceed is to examine the dependence of the estimates ${}_k\Delta\hat{r}^+$, ${}_k\Delta\hat{r}^-$, ${}_k\Delta\hat{R}$ on bf and $b\omega$. Let us expand ${}_k\Delta\hat{r}^+$ in Taylor series with respect to bf and $b\omega$. First note that the Jacobian matrix of the final state ${}_{k+M}X$, from which $\Delta\cdot$ terms are extracted, with respect to the initial conditions, is given by:

$${}_{k+M}J = \frac{\partial {}_{k+M}X}{\partial {}_kX} = \prod_{i=k}^{k+M} {}_iF, \quad (29)$$

Therefore

$$\begin{aligned} {}_k\Delta\hat{r}^+(bf) &= {}_k\Delta\hat{r}^+ + \frac{\partial {}_{k+M}\Delta\hat{r}^+}{\partial \bar{bf}}(bf - \bar{bf}) + \\ &+ \frac{\partial {}_{k+M}\Delta\hat{r}^+}{\partial \bar{b\omega}}(b\omega - \bar{b\omega}), \end{aligned} \quad (30)$$

where the partial derivatives are the proper blocks of ${}_{k+M}J$, as defined in Equation 29. In other words, ${}_{k+M}X$ has been obtained integrating inertial readings using a fixed, a-priori value for bf and $b\omega$, or zero if no a-priori knowledge is available. Equation 30 gives us a way to obtain a corrected value for ${}_k\Delta\hat{r}^+$ every time a new estimate of the inertial biases becomes available, without having to recompute the integral. The same reasoning can be made for ${}_k\Delta\hat{r}^-$ and ${}_k\Delta\hat{R}$.

The observation models for the pre-integrated inertial readings defined in Equation 15 and 21 can now be modified making the dependency of the $\Delta\cdot$ terms on the unknowns present in \mathbb{X} explicit. This allows the Dynamic Network solver to estimate the biases.

In Equation 30 a first order Taylor expansion has been used. However, \mathcal{F} is highly non-linear and such linear propagation will be inaccurate if M is large or if \bar{bf} and $\bar{b\omega}$ substantially differ from the values of bf and $b\omega$ in \mathbb{X} . In this case, it might be needed to re-compute the $\Delta\cdot$ terms or to employ a higher order Taylor expansion. This fact will be explored later on in the experimental section.

EXPERIMENTAL EVALUATION

In this section we compare the conventional inertial measurements processing in DN, as described in Section 3. with respect to the IMU pre-integration technique presented in Section 4. and 5.1 in a simulated scenario.

We consider a typical UAV flight for corridor mapping for which we would like to determine the navigation solution using inertial and GNSS position observations. We consider different IMU rates and for each rate we compare the conventional processing and different choices of M in IMU pre-integration.

To generate the simulated observations, we proceed as follows. A flight of a real fixed-wing UAV is considered as a model to generate synthetic observations, so that those would encode realistic motion patterns. A navigation solution (computed by other means) is available for the chosen flight and is depicted in Figure 2. We downsample the solution to 25 Hz and we use these poses as control points for a third order spline, as described in (Lovegrove et al., 2013). Such spline defines a continuous curve in SE(3), with continuous derivatives up to order two. We can then sample this spline at arbitrary time instants, obtaining analytically all the kinematic properties of the body frame needed to evaluate the continuous time models for inertial sensors. From these samples, the nominal ω_{ib}^i and f^b are derived and later corrupted by noise and bias as in Equation 1 and 2.

We consider the following different test cases: 100 Hz IMU readings and 200 Hz IMU readings, both with GNSS position observations at 1 Hz. For those, we compare the method presented in Section 3. with IMU-preintegration and different choices of M , so that the effective pose rate becomes 25 Hz, 10 Hz, and 5 Hz. These correspond to different choices of M depending on the IMU rate. For each test case, we compare the processing methods in Monte Carlo fashion, generating 100 different noise realizations. All the measurements are corrupted with white noise. For the inertial sensors, we assume the following power spectral densities: $41 \mu\text{g}/\sqrt{\text{Hz}}$ for the accelerometers and $1 \text{ deg}/(\text{h}\sqrt{\text{Hz}})$ for the gyroscopes, which correspond to a typical InterSense NavChip IMU. For the GNSS position observations, we consider a stand-alone operation mode and $2.5/3 \text{ m}$ horizontal and $5/3 \text{ m}$ vertical white noise standard deviation. The inertial sensors are also affected by a constant, “turn on” bias, which we draw uniformly within $[-8, 8] \text{ mg}$ for the accelerometers and within $[-720, 720] \text{ deg/h}$ for the gyroscopes. A box plot for the position and orientation RMS

obtained in each Monte Carlo run, for each processing method, is shown in Figure 3.

From the experimental results we can make the following points:

1. As expected, the conventional processing scheme, as in Section 3., yields consistently better RMS when the IMU rate is increased from 100 to 200 Hz. In our experiments, this improvement may be lower than what could be obtained in real world cases: the process of generating synthetic observations from a continuous spline based on 25 Hz control points intrinsically smooths the motion and eliminates effects such as vibrations and other high frequency components.
2. The processing with pre-integrated IMU readings equals the conventional scheme when M is small, e.g., 4 or 10. Note that when $M = 1$ the two techniques are virtually equivalent.
3. When M is large, e.g., $M \geq 20$, the performances of pre-integration deteriorate. This is essentially due to the difficulty of determining the biases of inertial sensors. Indeed, when M is large the first order Taylor expansion in Equation 30 does not capture the highly non linear dependency of ${}_k\hat{\Delta}r^+$ with respect to $b\omega$ and bf .

To confirm point three, we examine the bias estimation error for the IMU pre-integration cases. A box plot for $bf - bf_{\text{true}}$ in each Monte Carlo run is shown in Figure 4. It is possible to see that the bias estimation error considerably increases with M , and appears to converge to 0 as M decreases, which gives us an explanation for the degraded performances of pre-integration schemes shown in Figure 3. As a further confirmation, we repeat the experiments without adding the constant bias to the synthetic measurements. We find that this way, the position and orientation RMS for pre-integration schemes do not depend on M . These plots are omitted for space reasons.

CONCLUSIONS

In this work we compared two different methods for handling raw inertial measurements in Dynamic Networks. The first method consist in formulating an observation model for each measurement. This results in large estimation problems, and ill-conditioned normal equations, as the IMU rate increases. The second method, known in the robotics literature as IMU pre-integration, mitigates such issues while performing comparably to the first method. However, difficulties in estimating the inertial biases arise when the number of pre-integrated IMU measurements increases. This calls for an extension of the bias estimation method presented in Section 5.1. One possible solution could be considering higher order Taylor expansions in Equation 30.

REFERENCES

Blázquez, M. and Colomina, I., 2012. Relative INS/GNSS aerial control in integrated sensor orientation: models and performance. *ISPRS Journal of Photogrammetric Engineering and Remote Sensing* 67, pp. 120–133.

Bruton, A., Glennie, C. and Schwarz, K., 1999. Differentiation for high precision GPS velocity and acceleration determination. *GPS Solutions* 2(4), pp. 7–22.

Colomina, I. and Blázquez, M., 2004. A unified approach to static and dynamic modelling in photogrammetry and remote sensing. *ISPRS - The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* 35, pp. B1.

Cucci, D. A., 2014. A general sensor-fusion and parameters self-calibration framework with applications in mobile robotics. phdthesis, Politecnico di Milano (Italy). <https://www.politesi.polimi.it/handle/10589/98423>.

Cucci, D. A. and Matteucci, M., 2014. Position tracking and sensors self-calibration in autonomous mobile robots by gauss-newton optimization. In: *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, IEEE, pp. 1269–1275.

Cucci, D. A., Clausen, P., Skaloud, J. and Matteucci, M., 2017a. A general approach to time-varying parameters in pose-graph optimization. In: *Navigation Conference (ENC), 2017 European*, IEEE, pp. 265–271.

Cucci, D. A., Rehak, M. and Skaloud, J., 2017b. Bundle adjustment with raw inertial observations in UAV applications. *ISPRS Journal of Photogrammetry and Remote Sensing* 130, pp. 1–12.

Davis, T. A., 2006. Direct methods for sparse linear systems. Vol. 2, Siam.

Forster, C., Carlone, L., Dellaert, F. and Scaramuzza, D., 2017. On-manifold preintegration for real-time visual-inertial odometry. *IEEE Transactions on Robotics* 33(1), pp. 1–21.

Glennie, C., Skaloud, J., Rouzaud, D. and Baumann, C., 2011. Dynamic network adjustment for rigorous integration of passive and active imaging observations into trajectory determination. WO Patent App. PCT/CA2011/000,334.

Grossmann, C., Roos, H.-G. and Stynes, M., 2007. Numerical treatment of partial differential equations. Vol. 154, Springer.

Hansen, P. C. and OLeary, D. P., 1993. The use of the l-curve in the regularization of discrete ill-posed problems. *SIAM Journal on Scientific Computing* 14(6), pp. 1487–1503.

Heiskanen, W. A. and Moritz, H., 1981. Physical geodesy. Institute of Physical Geodesy, Technical University.

Indelman, V., Williams, S., Kaess, M. and Dellaert, F., 2013. Information fusion in navigation systems via factor graph based incremental smoothing. *Robotics and Autonomous Systems* 61(8), pp. 721–738.

Lovegrove, S., Patron-Perez, A. and Sibley, G., 2013. Spline fusion: A continuous-time representation for visual-inertial fusion with application to rolling shutter cameras. In: *BMVC*.

Lupton, T. and Sukkarieh, S., 2009. Efficient integration of inertial observations into visual SLAM without initialization. In: *Intelligent Robots and Systems (IROS), 2009 IEEE/RSJ International Conference on*, IEEE, pp. 1547–1552.

Lupton, T. and Sukkarieh, S., 2012. Visual-inertial-aided navigation for high-dynamic motion in built environments without initial conditions. *IEEE Transactions on Robotics* 28(1), pp. 61–76.

Savage, P. G., 1998. Strapdown inertial navigation integration algorithm design part I: Attitude algorithms. *Journal of guidance, control, and dynamics* 21(1), pp. 19–28.

Skaloud, J. and Lichti, D., 2006. Rigorous approach to bore-sight self calibration in airborne laser scanning. *ISPRS Journal of Photogrammetry and Remote Sensing* 61, pp. 47–59.

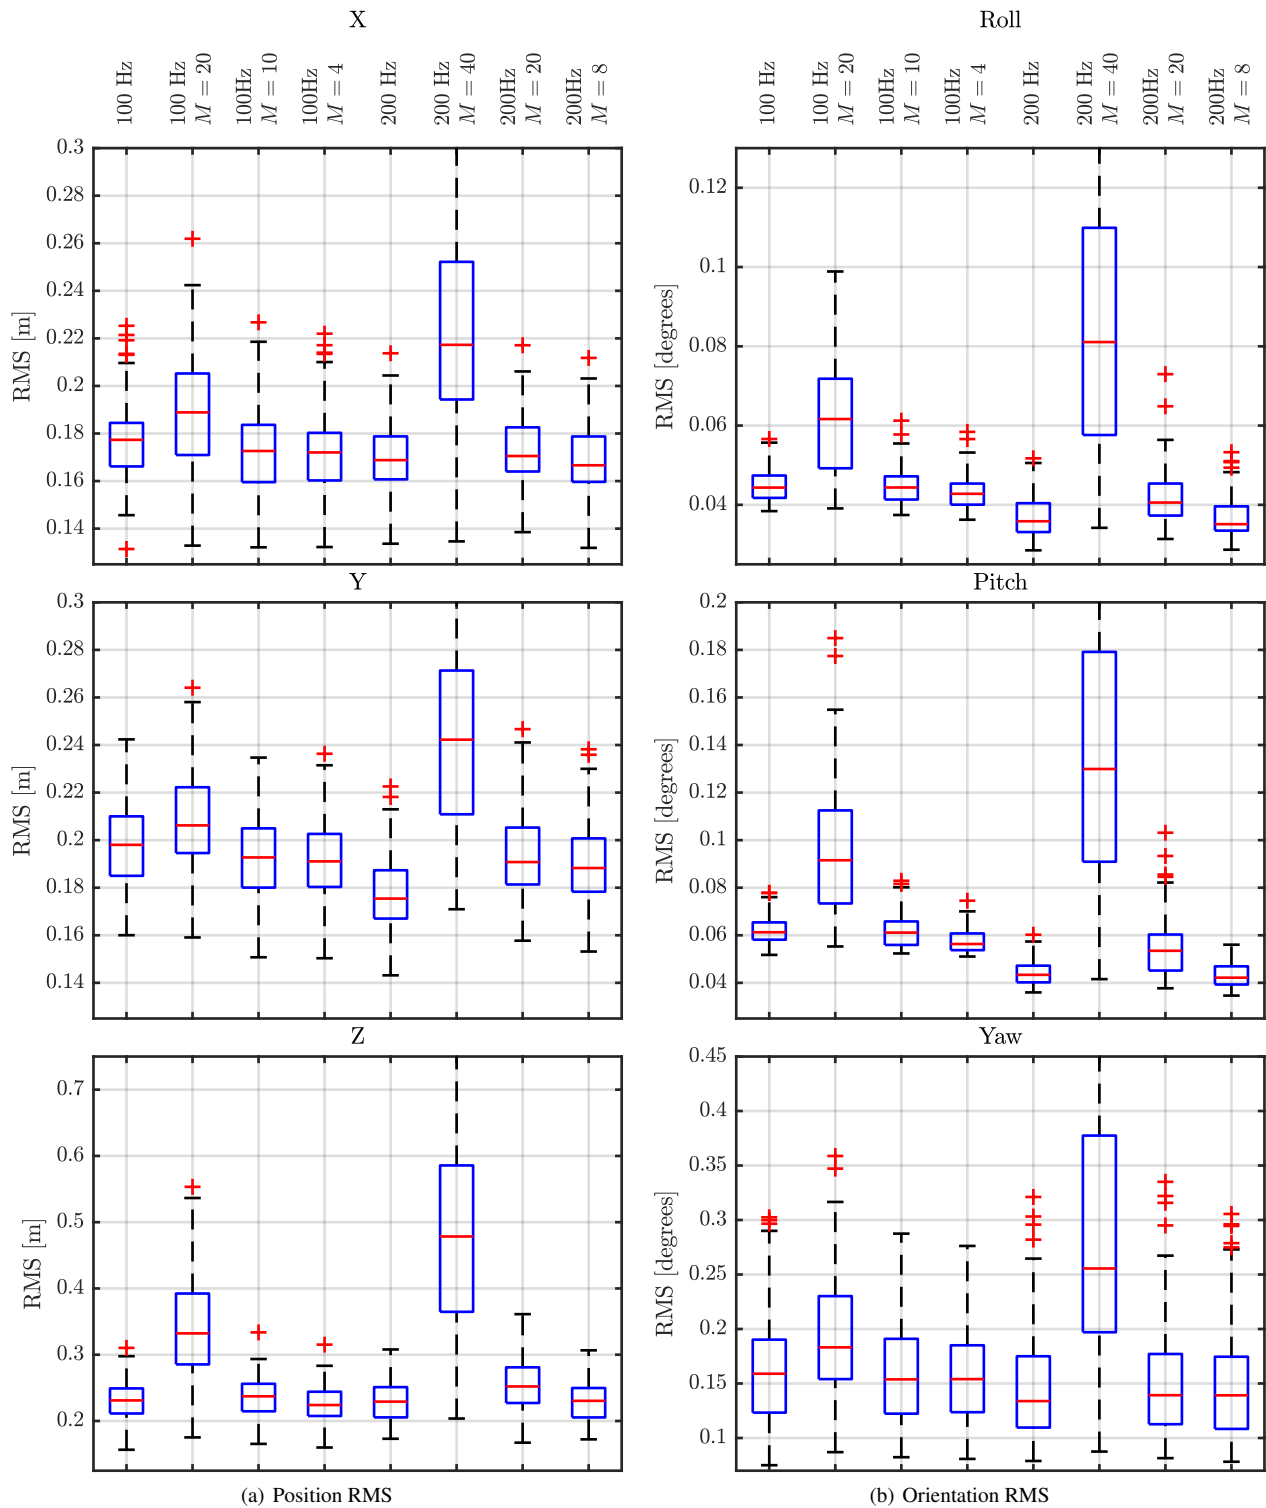


Figure 3. Box plot for the position and orientation RMS obtained in each Monte Carlo run, for each processing method.




Figure 4. Box plot for the accelerometer bias estimation error in each Monte Carlo run, for each processing method.

Strasdat, H., 2012. Local accuracy and global consistency for efficient SLAM. Ph.D. thesis, Imperial College London (UK).

Titterton, D. and Weston, J., 1997. Strapdown inertial navigation technology. Part of IEE radar, sonar, navigation and avionics series, Stevenage, U.K.