Active Online Visual-Inertial Navigation and Sensor Calibration via Belief Space Planning and Factor Graph Based Incremental Smoothing

Yair Ben Elisha* and Vadim Indelman*

Abstract— **High accuracy navigation in GPS-deprived environments is of prime importance to various robotics applications and has been extensively investigated in the last two decades. Recent approaches have shown that incorporating sensor's calibration states in addition to the 6DOF pose states may cause better performance of the system. However, these approaches typically consider a passive setting, where robot actions are externally defined. On the other hand, belief space planning (BSP) approaches account for different sources of uncertainty, thus identifying actions that improve certain aspects in inference, such as accuracy. Yet, existing BSP approaches typically do not consider sensor calibration, nor a visual-inertial SLAM setup. In this paper we contribute a BSP approach for active sensor calibration of a visual-inertial SLAM setup. For this purpose we incorporate within the belief both robot's pose and sensor calibration states while considering operation in partially unknown and uncertain environment. In particular, we leverage the recently developed concept of IMU pre-integration and develop appropriate factor graph formulation for future beliefs to facilitate computationally efficient inference within BSP. Our approach is valid for general cost functions, and can be used to identify best robot actions from a given set of candidate actions or to calculate locally-optimal actions using direct trajectory optimization techniques. We demonstrate our approach in high-fidelity synthetic simulation and show that incorporate sensors calibration state into the BSP significantly improved estimation accuracy.**

I. Introduction

Autonomous navigation in unknown or uncertain environments has been extensively investigated over the past two decades with numerous applications in robotics, including aerial GPS-denied navigation, indoor navigation and autonomous driving. Highly accurate online navigation in these, and many other, applications is of prime importance. Modern navigation systems no longer rely solely on inertial measurement units that are suspect to drift and on GPS that may be unreliable or unavailable, but calculate the navigation solution by fusing measurements captured by different on-board sensors (e.g. camera, laser sensors). When the environment is unknown or uncertain, robot localization (and navigation) involves also mapping the environment, a problem known in the navigation context as visual-inertial SLAM.

The corresponding inference problem involves tracking the probability density function (pdf) over variables of interest given available information. These variables often include navigation state, landmarks representing the mapped environment thus far, and sensor calibration parameters. The latter can represent extrinsic calibration, such as relative pose between different sensors, and also intrinsic parameters such as IMU bias and camera focal length. While some calibration parameters can be recovered in an offline fashion, due to stochasticity, some sensors (e.g. IMU) require also online calibration, without which navigation accuracy will be compromised. Online IMU calibration, however, is traditionally done considering GPS availability with pre-defined trajectories (i.e. actions) that were specifically calculated for such a setting.

Yet, actual performance depends, among other factors, on robot actions - different robot actions can often result in different estimation accuracies, especially in lack of external information (such as GPS). Thus, attaining high-accuracy navigation involves deep intertwining between inference and planning, requiring the latter to account for different sources of uncertainty. The corresponding planning approaches are known as belief space planning (BSP) methods, which have received considerable attention in recent years in the context of autonomous navigation in known, and more recently, unknown environments. However, existing BSP approaches typically do not consider sensor calibration aspects in the context of visual-inertial SLAM, while hard-coded actions that were suitable for GPS setting may perform poorly in scenarios considered herein.

In this paper we develop a BSP approach for active sensor calibration and accurate autonomous navigation considering a visual-inertial SLAM setting. Our approach is capable of calculating optimal actions for reducing estimation error within inference, reducing estimation error growth rate via IMU sensor calibration, or a combination of both. Additional typical costs in the objective function, such as reaching a goal and control effort, are naturally supported as well. Moreover, we leverage recent work that addressed visual-inertial SLAM using factor graphs and incremental smoothing [15], and use these techniques also within belief space planning. Finally, within this framework we extend the recently-developed concept of IMU pre-integration [21], that was used thus far only for information fusion and visual-inertial SLAM [15], to BSP so that longer planning horizons can be efficiently considered in presence of high-rate IMU measurements.

II. Related work

Traditional inertial navigation systems are based on the strapdown mechanism [6], in which IMU measurements are integrated into a navigation solution. Typically, navigation aiding methods apply filtering approaches for fusing measurements from other available sensors with the inertial solution. More recent approaches for information fusion in inertial navigation systems are calculating the optimal solution

*Department of Aerospace Engineering, Technion - Israel Institute of Technology, Haifa 32000, Israel.

based on a non-linear optimization involving all the unknown variables (and using all the available measurements). These approaches, also directly related to bundle adjustment (BA), are commonly used in the robotics community for solving the full simultaneous localization and mapping (SLAM) problem (see e.g. [3], [5]). In particular, [5] represents the posterior pdf using the factor graph graphical model [19], which naturally encodes the inherent sparsity of the underlying (square root) information matrices. Kaess et al. [17], [16] develop incremental smoothing and mapping (iSAM) approaches that, utilizing graphical models, allow to efficiently update the posterior with each new incoming measurement. Indelman et al. [13] introduced a related concept in the context of multi-robot collaborative navigation, however, considering the covariance form and not the information form as in [17], [16].

While iSAM is a general approach for efficient maximum a posteriori (MAP) inference given available measurements, possibly from different sensors, recent work [14], [15] examined its specific use for inertial navigation systems in presence of high rate sensors such as IMU. The recently developed concept of IMU pre-integration [21], [15], [7] allows to avoid the issue of adding factornodes to the graph at IMU rate, by summarizing consecutive IMU measurements, and incorporating into the factor graph only the summarized quantities while still supporting real time performance (see [15] for details).

While the above approaches focus on inference (e.g. estimating robot poses) given actions, different actions can lead to substantially different estimation performance, for example, in terms of accuracy. The corresponding problem of finding an optimal action(s) is an instantiation of a partially observable Markov decision process (POMDP) problem, which is computationally intractable [24]. Numerous approaches that trade-off optimality with computational complexity have been developed in recent years. These approaches are often known as belief space planning (BSP). These approaches can be segmented into several categories: point-based value iteration methods [20], [25], simulation based approaches [27], [28], sampling based approaches [9], [10], [26] and direct trajectory optimization approaches [11], [12], [29].

Yet, only few BSP approaches consider sensor calibration aspects [30], [8], [23]. For example, Hausman et al. [8] consider IMU calibration (in terms of bias), optimizing a nominal trajectory for self-calibration. However, that approach assumes GPS availability, in contrast to the problem setup considered herein. Webb et al. [30] develop a method for active calibration of extrinsic system parameters using continuous POMDP. Active calibration of extrinsic system parameters (e.g. transformation between frames) is also considered in [23]. However, neither of these works considers visual-inertial navigation systems, nor active calibration of intrinsic parameters such as IMU bias.

III. Probabilistic Formulation and Notations

Let $x_i = \begin{bmatrix} p_i & v_i & q_i \end{bmatrix}$ and $c_i = \begin{bmatrix} d_i & b_i \end{bmatrix}$ represent the navigation state vector and sensors calibration state vector, respectively, at time t_i. Here, p_i, v_i and q_i are the position, velocity and orientation (represented as a quaternion) of the body in the world frame and d_i, b_i are the gyroscope drift and accelerometers bias. Similarly we can add to the calibration state vector of other sensors parameters, e.g. camera calibration parameters. Also denote by L_i the perceived environment, e.g. in terms of 3D points, by time t_i.

We define the joint state $\Theta_k \doteq \{X_k, C_k, L_k\}$ where $X_k \doteq \{x_0, \ldots, x_k\}$ and $C_k \doteq \{c_0, \ldots, c_k\}$ represent, respectively, navigation and calibration states up to time t_k. Similarly we define $Z_{0:k} \doteq \{z_0, \ldots, z_k\}$ and $U_{0:k} \doteq \{u_0, \ldots, u_k\}$, where u_i and z_i represent control and obtained measurements at time t_i. In the following, we shall use the notation $z_i^{IMU} \in Z_i$ to denote IMU measurements and $z_i \in Z_i$ to denote camera observations.

Given these definitions, the probability density function (pdf) at time k, also referred as the *belief*, is given by

$$b\left(\Theta_k\right) = p\left(\Theta_k | Z_{0:k}, U_{0:k-1}\right), \tag{1}$$

from which the maximum a posteriori estimate

$$\Theta^\star = \arg\max_{\Theta_k} b\left(\Theta_k\right), \tag{2}$$

can be efficiently calculated while exploiting sparsity and re-using calculations [16].

The belief (1) can be explicitly written in terms of individual motion, calibration and observation models as (see e.g. [11])

$$\begin{aligned} b\left(\Theta_k\right) \propto & priors \prod_{i=1}^{k} p\left(x_i | x_{i-1}, c_{i-1}, z_{i-1}^{IMU}\right) \\ & \cdot p\left(c_i | c_{i-1}\right) p\left(z_i | \Theta_i^o\right) \end{aligned} \tag{3}$$

A. IMU observations

The probabilistic *motion model,* $p\left(x_i | x_{i-1}, c_{i-1}, z_{i-1}^{IMU}\right)$, is represented by the underlying nonlinear inertial navigation equations and the captured IMU measurements z^{IMU}

$$x_{i+1} = f\left(x_i, c_i, z_i^{IMU}\right) + w_i \;\;,\;\; w_i \sim N\left(0, \Sigma_w\right), \tag{4}$$

where w_i is an additive Gaussian noise with noise covariance matrix Σ_w, such that $p\left(x_i | x_{i-1}, c_{i-1}, z_{i-1}^{IMU}\right) \propto \exp\left[-\frac{1}{2}\|x_i - f\left(x_{i-1}, c_{i-1}, z_{i-1}^{IMU}\right)\|_{\Sigma_w}^2\right]$. The function $f\left(x_i, c_i, z_i^{IMU}\right)$ is governed by the following differential equations (see e.g. [15])

$$\begin{aligned} \dot{p_i} &= v_i \\ \dot{v_i} &= C^T\left(q_i\right)\left(a_i^m - b_i - n^a\right) - g \\ \dot{q_i} &= \frac{1}{2}\Omega\left(\omega_i^m - d_i - n^g\right) q_i \\ \dot{d_i} &= n^g \;\;,\;\; \dot{b_i} = n^a, \end{aligned} \tag{5}$$

where $C\left(q_i\right)$ is the rotation matrix obtained from the quaternion q_i, a^m and ω^m are the measured acceleration (more

precisely, specific force) and angular velocity, and $\Omega(\omega)$ is defined as

$$\Omega(\omega) = \begin{bmatrix} -\lfloor\omega\times\rfloor & \omega \\ -\omega^T & 0 \end{bmatrix}, \lfloor\omega\times\rfloor \doteq \begin{bmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{bmatrix}. \tag{6}$$

Since the IMU biases can change over time, they are modeled as random walk processes with zero-mean Gaussian noise n^g and n^a. More advanced models are outside the scope of this paper.

Observe that the true (noise free) acceleration and angular velocity of the robot are a function of the controls u_i; hence, we refer to the above IMU model as motion model.

B. Pre-Integrated IMU measurements

Unfortunately, incorporating high-rate IMU measurements into the posterior (3) and the underlying inference (2), which involves solving a nonlinear least squares (NLS) problem, does not provide real time performance in general. To avoid doing so, we adopt a recently developed concept of IMU pre-integration [21], and its application to sensor fusion in modern navigation systems [15].

The key idea is to integrate consecutive IMU measurements $\left(z_i^{IMU}, \ldots, z_j^{IMU}\right)$, between two lower-frequency measurements from other sensors such as cameras obtained at some times t_i and t_j, while expressing the resulting relative motion $\Delta x_{i\to j}$ in the body frame of x_i. Importantly, such a concept supports re-linearization, as part of the NLS iterative optimization, *without* recalculating $\Delta x_{i\to j}$ [21], [15]. Moreover, real time performance can be achieved by using intermediate $\Delta_{i\to t}$ with $t \leq j$, see [15] for details.

To summarize, we replace multiple IMU terms in the posterior (3) between any two time instances t_i and t_j with $p(x_j|x_i, \Delta x_{i\to j}, c_i)$.

C. Calibration Model

The *calibration model* $p(c_k|c_{k-1})$ is assumed to be with additive Gaussian noise with the process noise covariance matrix Σ_e, i.e.

$$c_{i+1} = g(c_i) + e_i \ , \ e_i \sim N(0, \Sigma_e), \tag{7}$$

D. Observation Model

The *observation model* $p(z_k|\Theta_k^o)$ at any time t_k can be written as

$$p(z_k|\Theta_k^o) = \prod_{l_j \in \Theta_k^o} p(z_{k,j}|x_k, l_j), \tag{8}$$

where $\Theta_k^o \subseteq \Theta_k$ are the overall involved states, $z_{k,j}$ is an (image) observation of landmark l_j, and the product accounts for all landmark observations captured at time t_k.

The observation model for each such observation is given by

$$z_{i,j} = h(x_i, l_j) + v_i \ , \ v_i \sim N(0, \Sigma_v), \tag{9}$$

where v_i is an additive Gaussian noise with noise covariance matrix Σ_v. Note that the landmark l_j can be either a random variable (SLAM framework) or a priori known (but uncertain) landmark at planning time.

E. Factor Graph Formulation

The joint pdf (1) can be represented using a graphical model known as a *factor graph* [19]. Formally, a factor graph is a bipartite graph $G_k = (\mathcal{F}_k, \mathcal{V}_k, \mathcal{E}_k)$ with two types of nodes: *factor nodes* $f_i \in \mathcal{F}_k$ and *variable nodes* $v_j \in \mathcal{V}_k \equiv \Theta_k$. Edges $e_{ij} \in \mathcal{E}_k$ can exist only between factor nodes and variable nodes, and are present if and only if the factor f_i involves a variable v_j. Each factor represents an individual term in the factorization (3) , and therefore one can write

$$p(\Theta_k|Z_{0:k}, U_{0:k-1}) \propto \prod_i f_i\left(\mathcal{V}_k^i\right), \tag{10}$$

where $\mathcal{V}_k^i$ represents an appropriate subset of variable nodes ($\mathcal{V}_k^i \subseteq \mathcal{V}_k$). See illustration in Figure 1 and [5] for further details.

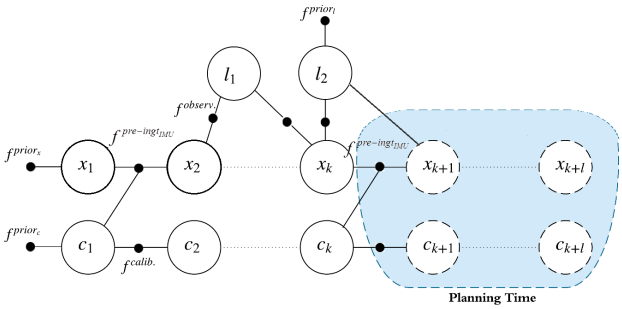

Fig. 1: Factor graph representation of Eq. (15) using pre-integrated IMU factors both for current time step (inference) and for future time steps (planning)

As seen in Figure 1, incorporating high-rate IMU measurements involves adding an appropriate IMU factor and variables at IMU rate. Instead, the concept of IMU pre-integration, mentioned in Section III-B, allows to summarize all consecutive IMU measurements between time instances t_i and t_j (e.g. determined by camera rate) into a pre-integrated delta $\triangle x_{i\to j}$, and to formulate an appropriate factor that only involves variables from these two time instances. This factor will be referred to as a pre-integrated IMU factor [15]. A single such factor, along with the variable x_j, are then added to the factor graph, instead of the numerous original IMU factors and corresponding intermediate variables. See illustration in Figure 1.

In this work, we use the principles of the pre-integrated IMU factor also within the planning problem. Therefore, future IMU measurements given control actions for L look ahead steps, are pre-integrated and added as pre-integrated IMU factor to the planning horizon factor graph.

IV. Approach

In this work we incorporate into belief space planning sensors calibration aspects such that navigation accuracy

is significantly improved while using vision-aided inertial navigation systems. Our approach extends the state of the art by incorporating into the belief (1) sensor calibration model and incorporating into the objective function, $J_k\left(b\left(\Theta_{k+L}\right), U_{k:k+L-1}\right)$ a cost function that quantifies calibration quality, possibly in combination with pose uncertainty reduction. We note that in lack of sources of absolute information (such as reliable GPS or known 3D points), accurate sensor calibration is of great importance for improving navigation accuracy. This is particularly true in GPS-deprived scenarios that involve segments of inertial navigation without vision-aiding. Moreover, we specifically consider high-rate IMU sensors and use the principle of the pre-integrated IMU factor also within the planning problem to avoid performing MAP inference (which required high computational power) at IMU rate instead of the (much lower) camera rate.

Specifically, we consider the following objective function over a planning horizon of L steps

$$J_k\left(b\left(\Theta_{k+L}\right), U_{k:k+L-1}\right) \doteq \sum_{l=0}^{L-1} \mathbb{E}\left(cf_l\left(b\left(\Theta_{k+l}\right), u_{k+l}\right)\right) + \mathbb{E}\left(cf_L\left(b\left(\Theta_{k+L}\right)\right)\right), \quad (11)$$

where $\mathbb{E}$ is the expected value and cf_l represents an immediate cost function and the expectation accounts for all possible realizations of the (unknown) future sensor observations. The optimal control is defined as

$$U^{\star}_{k:k+L-1} = \underset{U_{k:k+L-1}}{\arg\min}\, J_k\left(b\left(\Theta_{k+L}\right), U_{k:k+L-1}\right). \quad (12)$$

One can identify best robot actions (or motion plans), among those generated by existing motion planning approaches (e.g. sampling based approaches), or resort to direct optimization techniques to obtain locally optimal solutions in a timely manner [12]. In this work, we focus on the first case.

A. Inference and Recursive Formulation of the Belief

The belief at the lth look ahead time is defined, similarly to Eq. (1), as

$$b\left(\Theta_{k+l}\right) \doteq p\left(\Theta_{k+l} | Z_{0:k+l}, U_{0:k+l-1}\right), \quad (13)$$

where $Z_{0:k+l}$ and $U_{0:k+l-1}$ represent, respectively, all the measurements and controls up to (future) time t_{k+l}, while present time is t_k.

Determining optimal actions according to Eq. (12) involves evaluating the objective function J_k for different candidate actions $U_{k:k+L-1}$,given immediate costs cf_l, that we discuss in Section IV-C. This requires performing maximum a posteriori inference over the belief (13) such that

$$b\left(\Theta_{k+l}\right) = \mathcal{N}\left(\Theta^{\star}_{k+l}, \Lambda^{-1}_{k+l}\right), \quad (14)$$

where $\Theta^{\star}_{k+l}$ and Λ_{k+l} are the mean vector and information matrix, respectively.

Remark: We note that Λ_{k+l} can be calculated even though the actual values of future observations $Z_{k+1:k+l}$ are unknown, while $\Theta^{\star}_{k+l}$ is determined solely by the motion model by taking the common measurement likelihood assumption. Further details are outside the scope of this paper and can be found in [12].

The belief (13) can be recursively written as (see e.g. [10])

$$b\left(\Theta_{k+l}\right) = \eta b\left(\Theta_{k+l-1}\right) p\left(x_{k+l} | x_{k+l-1}, u_{k+l-1}, c_{k+l-1}\right) p\left(c_{k+l} | c_{k+l-1}\right) p\left(z_{k+l} | \Theta^{o}_{k+l}\right), \quad (15)$$

where η is a normalization constant and the other products are the belief at the previous step, the motion model, the calibration model and the observation model. Also, $\Theta^{o}_{k+l} \subseteq \Theta_{k+l}$ are the involved random variables in the measurement likelihood term $p\left(z_{k+l} | \Theta^{o}_{k+l}\right)$, which can be further expanded in terms of individual measurements $z_{k+l,j} \in z_{k+l}$ representing observations of 3D points l_j.

Similarly to Eq. (8), the measurement likelihood term can be written as,

$$p\left(z_{k+l} | \Theta^{o}_{k+l}\right) = \prod_{l_j \in \Theta^{o}_{k+l}} p\left(z_{k+l,j} | x_{k+l}, l_j\right), \quad (16)$$

where the product considers all the landmarks $l_j \in L_k$ expected to be observed from a future viewpoint x_{k+l}. In general, these could be landmarks mapped by planning time t_k or a priori known landmarks that correspond to known areas in the environment (if any).

B. Incorporating (Pre-Integrated) IMU Measurements

Further, since we are focusing on active visual-inertial SLAM, the motion model in Eq. (15) is actually represented by the nonlinear inertial equations and IMU measurements, see Section III-A. Thus, the motion model term $p\left(x_{k+l} | x_{k+l-1}, u_{k+l-1}, c_{k+l-1}\right)$ should be replaced by the corresponding *product* of IMU models $\prod_i p\left(x_i | x_{i-1}, c_{i-1}, z^{IMU}_{i-1}\right)$, representing the expected IMU measurements to be obtained between the time instances t_{k+l-1} and t_{k+l} given action u_{k+l-1}, i.e. $t_i \in [t_{k+l-1}, t_{k+l}]$.

As in passive visual-inertial SLAM, this would involve adding numerous factors and variables into the factor graph and the underlying MAP inference (14), thereby negatively impacting computational performance. Instead, we propose to use the concept of IMU pre-integration (Section III-B) and add a *single* probabilistic term that represents the motion model between camera times t_{k+l-1} and t_{k+l}. In other words, the motion model $p\left(x_{k+l} | x_{k+l-1}, u_{k+l-1}, c_{k+l-1}\right)$ in Eq. (15) is replaced by $p\left(x_{k+l} | x_{k+l-1}, \Delta x_{k+l-1 \to k+l}, c_{k+l-1}\right)$, where $\Delta x_{k+l-1 \to k+l}$ is the relative motion computed via pre-integrating (future) IMU measurements in the time frame $[t_{k+l-1}, t_{k+l}]$.

Note these measurements are a function of the considered actions via the nonlinear inertial navigation equations (III-A). Although the actual values of these (future) IMU observations are unknown, it is not required for calculating the posterior information matrix.

C. Choice of Cost Functions

Thus far, we showed how to incorporate within a future belief $b\left(\Theta_{k+l}\right)$ IMU measurements and sensor calibration

models (e.g. IMU bias) while considering general immediate cost functions. In this section we focus on a specific family of cost functions, cf, that utilize this new information within planning, considering different types of planning objectives. We define a general cost function [12]:

$$cf\left(b\left(\Theta_{k+l}\right),u_{k+l}\right)=\left\|E^{G}_{k+l}\Theta^{\star}_{k+l}-\Theta^{G}\right\|_{M_{\Theta}} + \left\|\zeta\left(u_{k+l}\right)\right\|_{M_u}+\underbrace{tr\left(M_{\Sigma}\Lambda^{-1}_{k+l}M^{T}_{\Sigma}\right)}_{cf^{\Sigma}\left(M_{\Sigma},\Lambda^{-1}_{k+l}\right)}, \quad (17)$$

where M_{Σ}, M_u and M_{Θ} are given weight matrices that can be *adjusted online* [11], [12], and $\zeta(u)$ is some known function that, depending on the application, quantifies the usage of control u. Θ^{G} is a predefined goal and E^{G}_{k+l} is a selection matrix, such that the matrix $E^{G}_{k+l}\Theta^{\star}_{k+l}$ contains a subset of states for which we want to impose a goal. Similarly, the matrix M_{Σ} in term cf^{Σ} can be used to choose the covariance of chosen states (position, pose, calibration etc.) from the joint covariance Λ^{-1}_{k+l}, i.e. consider only uncertainty of these states. We can consider different variations for the uncertainty cost term cf^{Σ}, such as accurate navigation or sensors self-calibration by using M_{Σ} that choose only the pose or calibration states, respectively. In this work we consider the uncertainty cost term $cf^{\Sigma^{TO}}$ that presents a 'trade-off' between two planning objectives (concurrent, or following each other): reaching a target with minimum navigation errors (i.e. *accurate navigation*) and calibrated sensors (i.e. *self-calibration*). The corresponding cost function is defined as

$$cf^{\Sigma^{TO}}\doteq tr\left(M_{\Sigma_c}\Lambda^{-1}_{k+l}M^{T}_{\Sigma_c}\right)+tr\left(M_{\Sigma_x}\Lambda^{-1}_{k+l}M^{T}_{\Sigma_x}\right), \quad (18)$$

where the terms penalize, respectively, pose uncertainty and quantify calibration quality. However, given the fact that there is a strong coupling between inertial sensors calibration to the pose uncertainty and that the belief includes both the pose and calibration states, the sensor calibration is still accounted for even if the uncertainty cost is only over position.

V. Results

We study the proposed approach in simulation considering autonomous navigation in partially known, uncertain, GPS-deprived environments. In the considered scenario, an aerial robot (e.g. quadrotor) has to autonomously navigate to a predefined goal with minimum localization uncertainty while utilizing limited prior information regarding the environment and using its onboard sensors. The latter include a downward-facing monocular camera and an IMU.

The environment comprises randomly scattered landmarks and an area without landmarks at all, where the goal is located (see Figure 3). Such a scenario was designed to highlight the importance of incorporating online sensor calibration aspects within planning, as in the area surrounding the goal no landmark observations can be obtained and inertial navigation is required; for example, this could model a dark or texture-less area. Navigation estimation error at the goal thus depends on IMU sensor calibration, in addition to the solution accuracy provided by visual SLAM while landmark observations are available. Clearly, navigation accuracy at the goal is expected to differ significantly for poorly- and well-calibrated IMU sensors.

For simplicity we assume the robot can only control its heading angle while keeping the velocity constant. The control effort, $\zeta(u)$ in Eq. (17), is therefore defined as the change in the heading angle. Such a setting induces single-axis maneuvers which are insufficient to yield a fully observable system with the considered visual-inertial SLAM framework [1], [22]; in particular, biases of accelerometer sensors cannot be calibrated in all axes.

For this reason, we consider some regions in the environment (in terms of landmarks within those regions) are a priori known to the robot with different levels of uncertainty (ranging between 10^{-5} and 10 meters) and can be used for updating the robot's belief. As will be seen, regions that are known with high-accuracy can be used for accelerometer sensor calibration and for position updates, while those with lower level of accuracy (10 meters) can only be used for position updates. See Figure 2 and supplementary material[1] [2] for further details. No other information is initially available, and thus the robot has to map the environment while navigating through it, i.e. perform SLAM, and consider appropriate aspects within the planning phase.

Given robot's trajectory from the planning phase, our simulation generates synthetic IMU measurements and camera observations, which are used to formulate appropriate pre-integrated IMU and projection factors, as described in Section III-B. Image observations are corrupted by zero-mean Gaussian noise with standard deviation of 0.5 pixels, and only landmarks within 200 meters from the sensors are observed. A basic IMU calibration model is considered, comprising accelerometers bias and gyroscopes drift of $10\,mg$ bias and $10\,^{\circ}/hr$ drift, respectively. Our Matlab implementation uses GTSAM library [4] for computationally efficient inference.

For the planning phase, at each planning time, we generate a set of candidate paths given the belief from inference, and choose the best path by calculating belief evolution and evaluating the objective function (11) for each candidate path. More specifically, similarly to [18], the landmarks seen thus far are first clustered, and then a shortest path is calculated to each cluster's center. These, together with the shortest path to the goal constitute the mentioned set of candidate paths. Belief evolution along each candidate path is performed by first simulating IMU and camera observations along the path, constructing an appropriate factor graph and recovering the posterior uncertainty via MAP inference.

A. Compared Approaches

We compare our approach, denoted as `'BSP-Calib'`, with two other methods: `'BSP'` (previous approaches, [12]) and `'Shortest-Path'`. The objective function in all the

[1] https://goo.gl/JTwGvT

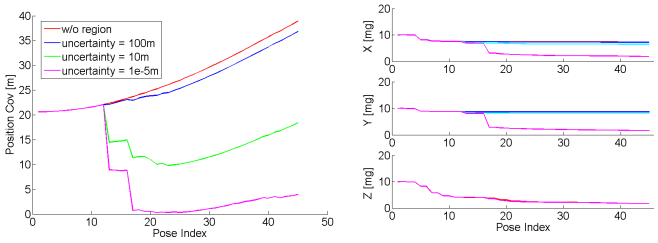

(a) **(b)**

Fig. 2: Results from basic study on known regions influence (described in more detailed in the supplementary material). The position covariance (a) and the accelerometers calibration covariance (b) are for the same scenario (map and trajectory) using different uncertainty levels of priori known regions

three methods can be described via Eq. (17) with weights appropriately adapted on the fly, as discussed below.

The `BSP` method does not incorporate sensor calibration states within the belief, which therefore describes at time k the joint pdf over past and current robot poses X_k and mapped landmark L_k. Belief evolution along a given candidate trajectory thus considers perfectly calibrated sensors, both within inference and planning. In this method, similarly to [12], we balance between position uncertainty reduction versus goal achievement by defining a soft bound $tr\left(M_{\Sigma_x}\Lambda_{k+l}^{-1}M_{\Sigma_x}^T\right) \leq \beta^x$, and set the weights $M_\Theta = \alpha_k^x$ and $M_{\Sigma_x} \doteq \sqrt{1-\alpha_k^x}\bar{M}_{\Sigma_x}$, where $\alpha_k^x \doteq 1 - \frac{tr\left(M_{\Sigma_x}\Lambda_{k+l}^{-1}M_{\Sigma_x}^T\right)}{\beta^x} \in [0,1]$ and $\bar{M}_{\Sigma_x}$ is a selection matrix that extracts the appropriate position covariance from Λ_{k+l}^{-1}. See [12] for further details.

In our approach, the `BSP-Calib` method, we consider a modification of the above to account also for the sensor calibration uncertainty while setting the weights. Specifically, we consider a soft bound β^c on sensor calibration uncertainty (accelerometer biases in our case), and let $\alpha_k^c \doteq 1 - \frac{tr\left(M_{\Sigma_c}\Lambda_{k+l}^{-1}M_{\Sigma_c}^T\right)}{\beta^c} \in [0,1]$ and $M_{\Sigma_c} \doteq \sqrt{1-\alpha_k^c}\bar{M}_{\Sigma_c}$, where $\bar{M}_{\Sigma_c}$ is defined similarly to $\bar{M}_{\Sigma_x}$. Since we now have two online-calculated parameters α_k^x and α_k^c, we set $M_\Theta = \min\left(\alpha_k^x, \alpha_k^c\right)$, and $M_{\Sigma_x} \doteq \sqrt{\alpha_k^c\left(1-\alpha_k^x\right)}\bar{M}_{\Sigma_x}$. The intuition here is as follows: as we approach any of the soft bounds β^x or β^c, i.e. $\alpha^x \to 0$ or $\alpha^c \to 0$, the goal attainment term is disabled to allow for active sensor calibration or position uncertainty reduction. However, since the environment is unknown and may include areas without landmarks (textureless or dark areas) that require inertial navigation, we prioritize sensor calibration over reduction in position uncertainty, since this would yield slower uncertainty evolution. One can observe that when the sensor is sufficiently calibrated with respect to β^c, the method will aim to actively reduce position uncertainty in case $\alpha_k^x \to 0$, or invest efforts in goal attainment. In our implementation we use $\beta^x = 150^2\, m^2$ and $\beta^c = 2$ mg. Clearly, properly balancing the weights in a multi-objective function can be delicate and better approaches could exist. Further investigation of these aspects is left to future research.

In the `Shortest-Path` method we simply set all the weight matrices to zero except of M_Θ. This method will therefore choose the shortest path to the goal without considering uncertainty aspects.

B. Simulation Results

The comparison between the three methods is presented in the following figures. The robot's trajectories are shown in Figure 3, where the cyan square is a known region with uncertainty of $10^{-5}m$, and the green square is a known region with uncertainty of $10m$. The green square located closer to start and goal locations. The unknown landmarks in the environment are the blue dots and the observed landmarks along the path are shown with '+' on them. Performance comparison between the three methods in terms of estimation uncertainty (position and accelerometers calibration) and paths length is presented in Figure 4. The presented results are the SLAM results given the chosen path in each method.

One can see in the results that our approach, `'BSP-Calib'`, yields better performance, in terms of position uncertainty, despite the fact that it is the longest path to the goal. In our approach the planning chooses to go through the known region with the lower uncertainty (cyan region) in order to calibrate the robot's accelerometers and then proceed to the goal. On the other hand, the `'BSP'` method does not consider the calibration state in the belief, and therefore chooses a shorter path that reduces the pose uncertainty by reaching the closer known region but with higher uncertainty (green region).

We note that actual performance in inference, in terms of covariance uncertainty upon reaching the goal, may be actually different than the one which was calculated within planning. While the latter only considers already-mapped (or known) landmarks and dead reckoning otherwise, the inference process updates the belief based on incoming sensor observations also of new landmarks. It is thus conceptually possible that the path determined by our approach, `'BSP-Calib'`, will yield in inference inferior results to `'BSP'` performance or `'Shortest-Path'` methods. This, for example, could be the case in an environment full with (unknown) landmarks. One possible direction to approach this gap of belief evolution in planning and inference is to consider within planning some statistics regarding landmark distribution. This direction is left to future research.

VI. Conclusions

We presented a BSP approach for active sensor calibration of a visual-inertial SLAM setup. Our BSP framework, where the belief accounts for the uncertainty in robot's IMU calibration and pose, allows the robot to select the best path that improve its state estimation. We showed performance comparison of our approach in relation to previous works for the scenario of partially unknown and uncertain environments comprising randomly scattered landmarks (including a priori known regions with different levels of uncertainty) and an area without landmarks at all. A key question that will be addressed in future research is how to incorporate such a concept within multi-robot belief space planning such that one robot can assist other robots to attain better navigation performance while autonomously navigating to different goals.

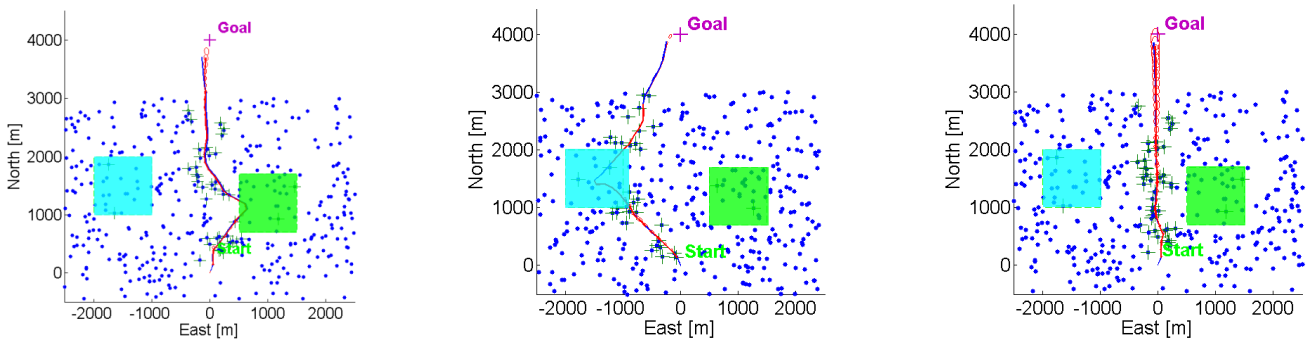

(a) **(b)** **(c)**

Fig. 3: The robot's trajectories for the three methods: (a) `'BSP'`, (b) `'BSP-Calib'` (our approach) and (c) `'Shortest-Path'`, where the cyan square is a known region with uncertainty of $10^{-5}m$, and the green square is a known region with uncertainty of $10m$. The unknown landmarks in the environment are the blue dots and the observed landmarks along the path are shown with + on them.

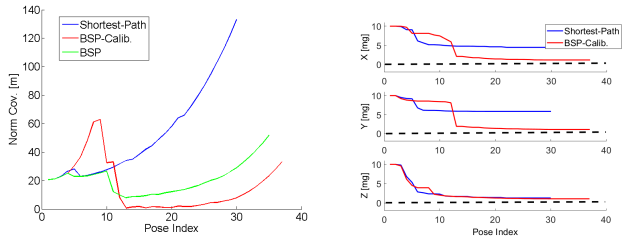

(a) **(b)**

Fig. 4: Performance comparison in terms of estimation uncertainty: (a) position covariance and (b) accelerometers calibration covariance. Calibration performance is not shown for 'BSP' as in this approach sensor calibration is not part of the belief, and is not considered in objective function. One can see that our approach, `'BSP-Calib'`, yields better performance despite the fact that it is the longest path to the goal.

REFERENCES

[1] M. W. Achtelik, S. Weiss, M. Chli, and R. Siegwart. Path planning for motion dependent state estimation on micro aerial vehicles. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 3926–3932, 2013.

[2] Y. Ben-Elisha and V. Indelman. Active online visual-inertial navigation and sensor calibration via belief space planning and factor graph based incremental smoothing - supplementary material. Technical Report ANPL-2017-02, Technion - Israel Institute of Technology, 2017.

[3] Mitch Bryson, M. Johnson-Roberson, and Salah Sukkarieh. Airborne smoothing and mapping using vision and inertial sensors. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 3143–3148, 2009.

[4] F. Dellaert. Factor graphs and GTSAM: A hands-on introduction. Technical Report GT-RIM-CP&R-2012-002, Georgia Institute of Technology, 2012.

[5] F. Dellaert and M. Kaess. Square Root SAM: Simultaneous localization and mapping via square root information smoothing. *Intl. J. of Robotics Research*, 25(12):1181–1203, Dec 2006.

[6] J.A. Farrell. *Aided Navigation: GPS with High Rate Sensors*. McGraw-Hill, 2008.

[7] Christian Forster, Luca Carlone, Frank Dellaert, and Davide Scaramuzza. On-manifold preintegration for real-time visual–inertial odometry. *IEEE Transactions on Robotics*, 2016.

[8] Karol Hausman, James Preiss, Gaurav Sukhatme, and Stephan Weiss. Observability-aware trajectory optimization for self-calibration with application to uavs. *arXiv preprint arXiv:1604.07905*, 2016.

[9] G. A. Hollinger and G. S. Sukhatme. Sampling-based robotic information gathering algorithms. *Intl. J. of Robotics Research*, pages 1271–1287, 2014.

[10] V. Indelman. Towards cooperative multi-robot belief space planning in unknown environments. In *Proc. of the Intl. Symp. of Robotics Research (ISRR)*, September 2015.

[11] V. Indelman, L. Carlone, and F. Dellaert. Towards planning in generalized belief space. In *The 16th International Symposium on Robotics Research*, Singapore, December 2013.

[12] V. Indelman, L. Carlone, and F. Dellaert. Planning in the continuous domain: a generalized belief space approach for autonomous navigation in unknown environments. *Intl. J. of Robotics Research*, 34(7):849–882, 2015.

[13] V. Indelman, P. Gurfil, E. Rivlin, and H. Rotstein. Graph-based distributed cooperative navigation for a general multi-robot measurement model. *Intl. J. of Robotics Research*, 31(9), August 2012.

[14] V. Indelman, S. Wiliams, M. Kaess, and F. Dellaert. Factor graph based incremental smoothing in inertial navigation systems. In *Intl. Conf. on Information Fusion, FUSION*, 2012.

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

[16] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert. iSAM2: Incremental smoothing and mapping using the Bayes tree. *Intl. J. of Robotics Research*, 31:217–236, Feb 2012.

[17] M. Kaess, A. Ranganathan, and F. Dellaert. iSAM: Incremental smoothing and mapping. *IEEE Trans. Robotics*, 24(6):1365–1378, Dec 2008.

[18] A. Kim and R. M. Eustice. Active visual slam for robotic area coverage: Theory and experiment. *Intl. J. of Robotics Research*, 2014.

[19] F.R. Kschischang, B.J. Frey, and H-A. Loeliger. Factor graphs and the sum-product algorithm. *IEEE Trans. Inform. Theory*, 47(2), February 2001.

[20] H. Kurniawati, D. Hsu, and W. S. Lee. Sarsop: Efficient point-based pomdp planning by approximating optimally reachable belief spaces. In *Robotics: Science and Systems (RSS)*, volume 2008, 2008.

[21] T. Lupton and S. Sukkarieh. Visual-inertial-aided navigation for high-dynamic motion in built environments without initial conditions. *IEEE Trans. Robotics*, 28(1):61–76, Feb 2012.

[22] Agostino Martinelli. Visual-inertial structure from motion: observability and resolvability. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4235–4242. IEEE, 2013.

[23] Jérôme Maye, Hannes Sommer, Gabriel Agamennoni, Roland Siegwart, and Paul Furgale. Online self-calibration for robotic systems. *The International Journal of Robotics Research*, page 0278364915596232, 2016.

[24] C. Papadimitriou and J. Tsitsiklis. The complexity of markov decision processes. *Mathematics of operations research*, 12(3):441–450, 1987.

[25] J. Pineau, G. J. Gordon, and S. Thrun. Anytime point-based approximations for large pomdps. *J. of Artificial Intelligence Research*, 27:335–380, 2006.

[26] S.I. Roumeliotis and G.A. Bekey. Distributed multi-robot localization. *IEEE Trans. Robot. Automat.*, August 2002.

[27] C. Stachniss, G. Grisetti, and W. Burgard. Information gain-based exploration using rao-blackwellized particle filters. In *Robotics: Science and Systems (RSS)*, pages 65–72, 2005.

[28] R. Valencia, M. Morta, J. Andrade-Cetto, and J.M. Porta. Planning reliable paths with pose SLAM. *IEEE Trans. Robotics*, 2013.

[29] J. Van Den Berg, S. Patil, and R. Alterovitz. Motion planning under uncertainty using iterative local optimization in belief space. *Intl. J. of Robotics Research*, 31(11):1263–1278, 2012.

[30] Dustin J Webb, Kyle L Crandall, and Jur van den Berg. Online parameter estimation via real-time replanning of continuous gaussian pomdps. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5998–6005. IEEE, 2014.