Rao-Blackwellised Inertial-SLAM with Partitioned Vehicle Subspace

Mark Euston and Jonghyuk Kim
Department of Engineering, Australian National University
Canberra, ACT 0200, Australia
e-mail: {mark.euston, jonghyuk.kim}@anu.edu.au

Abstract

This paper presents methods which enable the Rao-Blackwellised (R-B) particle filtering technique to be applicable for the airborne simultaneous localisation and mapping problem. Although R-B filter has been successfully applied to mobile/ground vehicles, its extension to flying vehicles has been impractical due to the high dimensionality involved in inertial navigation system (INS). To overcome this problem, the full INS state is further partitioned into an external state (vehicle pose) and an internal state (navigation and sensor calibration), with a particle filter being applied only to the external state. The computational complexity is further reduced by developing a hybrid R-B Inertial-SLAM. Simulation results will be presented with simulated flight data, showing reliable performances during loop-closures.

1 Introduction

Navigation (or localisation) is a fundamental, but still challenging, task in most autonomous vehicles in performing their tasks successfully and generating high-level control signals for vehicle guidance. Although satellite-based localisation systems have been widely available, they are still susceptible to signal shadings and blocking, being unreliable or completely unavailable in many robotic environments such as forestry, mining, underwater and urban canyons. An accurate mapping of environment becomes thus essential not only for a successful task operation but also for a reliable localisation within the environment. This problem has been known as Simultaneous Localisation And Mapping (SLAM) which provides a probabilistic framework to map environmental features whilst utilising them for the vehicle localisation ([Durrant-Whyte and Bailey, 2006]).

SLAM is intrinsically a high-dimensional state estimation problem, adding new features as the vehicle encompasses the environment. Extended Kalman filter (EKF) has been most popular for its real-time implementation. The key property is in maintaining the full vehicle-to-map correlation information within a covariance matrix. It however requires quadratic storage and computational complexities of $O(n^2)$, with n being the dimension of the state space. Clearly, the increasing number of features will eventually limit the real-time performance. To tackle this problem, various methods have been developed such as using compressed [Guivant and Nebot, 2001] and hierarchical map managements [Estrada and Tardos, 2005], sparsed information filter by [Thrun *et al.*, 2002], and Rao-blackwellised (R-B) particle filter by [Grisetti *et al.*, 2007] (also known as fast-SLAM in [Montemerlo *et al.*, 2004]). For more detailed discussions, refer to a recent survey paper by [Durrant-Whyte and Bailey, 2006].

In particular, the R-B particle filter SLAM offers computationally tractable particle filtering by, first, partitioning the full state into a vehicle and a map states, then applying particle filter for the former and Kalman filter for the latter. This process is called Rao-Blackwellisation, and the joint probability density function (PDF) is represented by a set of particle samples with associated conditional PDFs of map. The key benefit is the conditionally independency between map-features given the pose trajectory of the vehicle, resulting in the linear computational complexity of $\mathcal{O}(n)$. In addition, it provides an effective means to deal with nonlinearity and non-Gaussian noises within vehicle dynamics, making it highly attractive for flying vehicles.

Airborne SLAM on a fixed-wing UAV platform has been demonstrated in [Kim and Sukkarieh, 2004], showing its feasibility as a stand-alone or complementary airborne navigation system. Although there has been SLAM on a low-dynamic Blimp platform by [Jung and Lacroix, 2003], the full INS technology has never been exploited until this work. This is important since INS can provide a all-terrain navigation capability delivering the full 6 degrees-of-freedom vehicle information.

The work however was based on EKF framework, suffering quadratically increasing complexity as in ground vehicles. Considering the large-scale of airborne mapping and non-linearity, R-B particle filter SLAM becomes a natural candidate for airborne SLAM implementation. Its direct application to INS, however, is not as straightforward as in the mobile robots. One of main reasons is the high-dimensionality associated with INS, which typically has position, velocity, and attitude states, as well as sensor error states for gyroscopes and accelerometers. This leads to a dimensionality of 15 at least, while most mobile robots have only 3 for $x-y$ positions and a heading. The number of particles required increases exponentially with the state dimension hence the direct particle filtering for INS infeasible.

In this paper, the INS states are further partitioned into an external and an internal states, where the former represents the vehicle pose (position and attitude) required for mapping process, and the latter for inertial navigation (velocity) and sensor calibration for gyro and accelerometer (biases). Note that the velocity state is required to obtain position from accelerometer measurement. The external states are then estimated by a particle filter, whilst the internal and map states being estimated by a parallel Kalman filters. Unlike the conditionally independent map states, the internal states are still dependent given the pose particles. This is due to the internal dynamics within INS, causing additional computational burdens. For this problem a hybrid R-B Inertial-SLAM is developed, which has a single full INS EKF in concert with the pose-sampled Rao-Blackwellised SLAM. The simulation results will show the effectiveness of these methods.

This paper is organised as follows: Section 2 will present the problem statement with an overview of the Rao-Blackwellised filtering in Section 3. Section 4 will provide algorithms for R-B Inertial-SLAM in a Bayesian framework. Section 6 will present results using simulated flight data. Finally Section 7 will draw a conclusion with future research directions.

2 Problem Statement

The joint probability density function (PDF) for airborne SLAM system at time k conditioned on the cumulative observations is

$$p\,(\mathbf{p}, \mathbf{v}, \boldsymbol{\psi}, \mathbf{b}_a, \mathbf{b}_g, \mathbf{M} \mid \mathbf{Z}^k), \tag{1}$$

where

- INS state vector consisting of position ($\mathbf{p}$), velocity ($\mathbf{v}$), and attitude ($\boldsymbol{\psi}$) with sensor bias errors for accelerometers ($\mathbf{b}_a$) and gyroscopes ($\mathbf{b}_g$);
- the map state vector with stationary N-feature 3D positions: $\mathbf{M} = \{\mathbf{m}_1, \mathbf{m}_2, \cdots, \mathbf{m}_N\}$; and

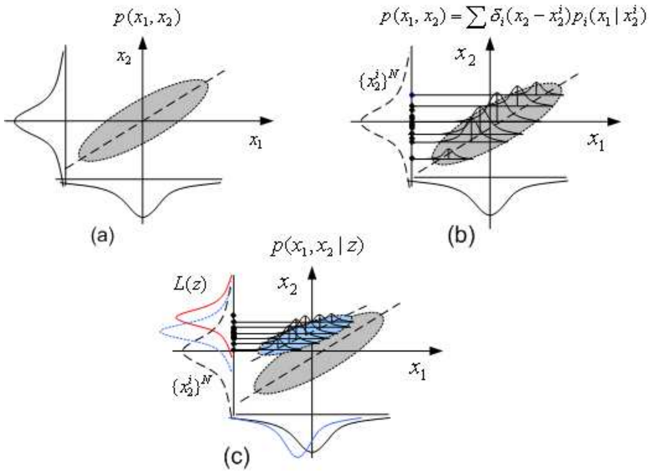

Figure 1: R-B particle filtering for a 2 dimensional state system: a) a full joint PDF. b) one of its subspaces (x_2) is approximated by a set of particles each of which carries a conditional PDF of x_1. c) With an observation on x_2, the particles are re-located, subsequently changing the full joint and marginal PDFs.

- the feature observations until current time k: $\mathbf{Z}^k = \{\mathbf{z}_1, \mathbf{z}_2, \cdots, \mathbf{z}_k\}$.

The high-dimensionality arises from both map and INS states. Thus the question is how to partition the INS states into more tractable subspaces whilst decoupling the map correlation given vehicle state.

3 Overview of Rao-Blackwellised Filter

Direct application of particle filtering for a high-dimensional system is not computationally tractable and thus not desirable. Rao-Blackwellised (R-B) filter however, provides a effective means to reduce the sample-space by factorising the full density and by applying the particle filtering only for the reduced subspace.

Figure 1 illustrates R-B filter for a bivariate estimation problem. The joint PDF $p(x_1, x_2)$ has individual marginal densities with a correlation as can be seen from the elongation of the covariance ellipsoid. In R-B filter, one of its subspaces (in this case x_2) can be approximated by a set of particle samples each of which carries a conditional density of the other state x_1. The full joint PDF is thus represented by a set of $\{\{x_2^i\}^N, p(x_1|x_2^i)\}$ collectively. Whenever, an observation occurs on the state x_2, the particle samples are weighted based on the likelihood and re-sampled accordingly [Montemerlo *et al.*, 2004]. Now the newly relocated samples with associated conditional PDFs represent the updated joint density.

In SLAM problem, although the full state vector can be partitioned to any sub groups, partitioning into a vehicle and a map is most effective due to the conditional independency given the vehicle state. The observation update needs a slight modification as the observation is related to both vehicle and map. That is, whenever an observation arrives, the particle samples are weighted and re-sampled as before, but the associated conditional PDF should be updated together assuming each particle (vehicle pose) being perfect [Grisetti *et al.*, 2007] and [Durrant-Whyte and Bailey, 2006].

4 Rao-Blackwellised Inertial-SLAM

The idea is to separate the high-dimensional INS states into two sub-states: an external pose state $\mathbf{x}_e$ which is related to the mapping, and an internal states $\mathbf{x}_i$ for navigation and inertial sensor calibration:

$$p\,(\underbrace{\mathbf{p},\boldsymbol{\psi}}_{\mathbf{x}_E},\underbrace{\mathbf{v},\mathbf{b}_a,\mathbf{b}_g}_{\mathbf{x}_I},\mathbf{M}\mid\mathbf{Z}^k) \tag{2}$$

$$= p\,(\mathbf{x}_I,\mathbf{x}_E,\mathbf{M}\mid\mathbf{Z}^k) \tag{3}$$

$$= p\,(\mathbf{x}_I,\mathbf{M}\mid\mathbf{x}_E,\mathbf{Z}^k)\times p\,(\mathbf{x}_E\mid\mathbf{Z}^k), \tag{4}$$

where the full joint PDF is factorised into a conditional PDF given the external state and a PDF for the external state.

Since the map states are conditionally independent each other given the external pose states, it further becomes

$$p\,(\mathbf{x}_k,\mathbf{M}\mid\mathbf{Z}^k) \tag{5}$$

$$= p\,(\mathbf{x}_I\mid\mathbf{x}_E)\,p\,(\mathbf{M}\mid\mathbf{x}_E,\mathbf{Z}^k)\,p\,(\mathbf{x}_E\mid\mathbf{Z}^k) \tag{6}$$

$$= \underbrace{p\,(\mathbf{x}_I\mid\mathbf{x}_E)\prod_i^N p\,(\mathbf{m}_i\mid\mathbf{x}_e,\mathbf{Z}^k)\times}_{\text{Candidate for KF}}\;\underbrace{p\,(\mathbf{x}_E\mid\mathbf{Z}^k)}_{\text{Candidate for PF}} \tag{7}$$

From this factored density, applying R-B particle filtering becomes straightforward: the PDF of external state can be represented by a particle filter and the PDF of the internal and map PDFs by analytical Kalman filters:

$$p\,(\mathbf{x}_I,\mathbf{x}_E,\mathbf{M}\mid\mathbf{Z}^k) \tag{8}$$

$$= \left[p\,(\mathbf{x}_I\mid x_E^i)\prod_j^N p\,(\mathbf{m}_j\mid x_E^i,\mathbf{Z}^k)\right]\times p\,(x_E^i\mid\mathbf{Z}^k).$$

Figure 2(a) illustrates the resulting structure of the joint PDF. Since the map features are only dependant on the pose trajectory, the map-to-map correlations are subsequently zero between map Kalman filters given the pose particles. If m-map features are used, then each particle will maintain one internal-state KF of 9-dimension and m map KFs of 3-dimension. Therefore if n particles ar used, the storage requirement will be $n[9^2+m(3^2)]$.

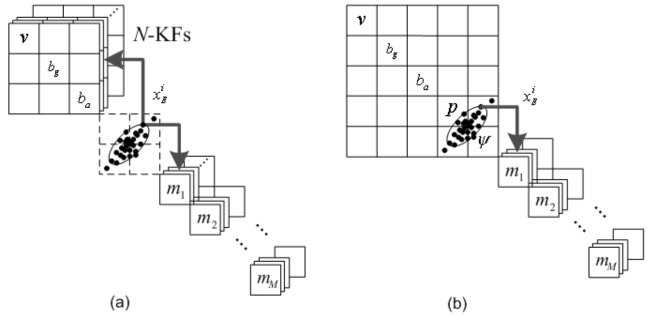

Figure 2: (a) R-B Inertial-SLAM effectively removes the map-to-map correlations but not those between the internal states. (b) Hybrid R-B Inertial-SLAM, however, maintains only one vehicle EKF, while utilising pose-particles reducing the computational complexity.

5 Hybrid R-B Inertial-SLAM

The problem in the previous method is each particle should maintain each internal-KF (which has 9 dimensionality in this case). Coupled with the high-update rates in INS (up to KHz), this still can be computationally challenging. As discussed previously, the internal states are not conditionally independent given the pose history due to the dynamic coupling within the internal states.

To relieve this complexity, the parallel internal-KFs can be merged into a single EKF while maintaining a pose-sampled particle filter. In this configuration, the pose particles are updated using vehicle-feature observations, then its marginal density should be reconstructed so that it can be propagated to the full vehicle EKF.

$$p\,(\mathbf{x}_I,\mathbf{x}_E,\mathbf{M}\mid\mathbf{Z}^k) \tag{9}$$

$$= \underbrace{p\,(\mathbf{x}_I\mid\mathbf{x}_E)\underbrace{\prod_j^N p\,(\mathbf{m}_j\mid x_E^i,\mathbf{Z}^k)\times p\,(x_E^i\mid\mathbf{Z}^k)}_{p\,(x_E^i|\mathbf{Z}^k)\longrightarrow p\,(\mathbf{x}_E|\mathbf{Z}^k)}}_{\longrightarrow p\,(\mathbf{x}_I,\mathbf{x}_E|\mathbf{Z}^k)}.$$

For n particles, the memory requirement for covariance matrix will be $[15^2+n\times m(3^2)]$. The storages requirement for INS part is constant while that of the previous method is linearly proportional to the number particles.

5.1 Prediction

The state-space model for the vehicle becomes INS model:

$$p\,(\mathbf{x}_{k+1}\mid\mathbf{x}_k)\quad\Leftrightarrow\quad\mathbf{x}_{k+1}=\mathbf{f}(\mathbf{x}_k)+\mathbf{w}_k, \tag{10}$$

with nonlinear state-transition function $\mathbf{f}(\cdot)$ (see details in Equation 15), with $\mathbf{w}_k$ being the process noise.

Since there are no direct measurement for INS state, the mean and covariance within extended Kalman filter are simply predicted using inertial measurement inputs.

In parallel, the pose particles are also propagated base on the same nonlinear model as well as using the estimated values of velocity ($\hat{\mathbf{v}}_k^n$) and gyro bias ($\hat{\mathbf{b}}_{g,k}^b$):

$$p_{k+1}^{n,i} = p_k^{n,i} + \triangle t(\hat{\mathbf{v}}_k^n + w_v^i) \quad (11)$$

$$\boldsymbol{\psi}_{k+1}^{n,i} = \boldsymbol{\psi}_k^{n,i} + \triangle t \mathbf{E}_b^n(\boldsymbol{\omega}_k^b + \hat{\mathbf{b}}_{g,k}^b + w_g^i), \quad (12)$$

where $\mathbf{E}_b^n$ is a transformation matrix between body and Euler rates and w_v^i is the velocity noise samples drawn from the uncertainty covariance $\mathcal{N}(0; \mathbf{P}_{vv})$, and w_g^i is the velocity noise samples drawn from the uncertainty covariance $\mathcal{N}(0; \mathbf{P}_{b_g})$ to make the pose particles spread out.

The prediction of the map-KFs are a simple stationary process.

5.2 Observation Update

The probabilistic observation model relates the observation to the j^{th}-feature position $\mathbf{m}_{j,k}$ and the vehicle pose $\mathbf{x}_{E,k}$:

$$p(\mathbf{z}_k \mid \mathbf{x}_k)) \quad \Leftrightarrow \quad \mathbf{z}_{j,k} = \mathbf{h}(\mathbf{x}_{E,k}, \mathbf{m}_{j,k}) + \mathbf{v}_k, \quad (13)$$

with $\mathbf{f}(\cdot)$ being the nonlinear observation function (see details in Equation 16), with $\mathbf{v}_k$ being the observation noise.

Once feature observations occur, there are three update steps to fuse these information: pose particles update, map-KFs update, and full vehicle-EKF update.

First, the pose particles are weighted based on their closenesses to the observation by using the map estimates. The weight can be further used to generate a new set of particles proportional to this weight. This process is called a re-sampling and the new particles are allocated to a uniform weight of $1/N$.

Second, the map-KFs are then updated using the standard Kalman filter algorithm with an assumption of the pose being perfect.

The last step is to propagate the information from the particle filter to the vehicle EKF. For this purpose, the density parameters, that is mean and covariance, for the pose should to be obtained. The particle distribution are in general non-Gaussian, and thus the Sum of Gaussian (SoG) representation is desirable. A single Gaussian is used though as an initial work:

$$\{x_E^i\}^N \quad \mapsto \quad \sum_j^{N=1} \mathcal{N}_j(\mathbf{x}_E^+; \mathbf{P}_E^+) \quad (14)$$

Given the updated external state, the internal state within the full vehicle EKF can be easily updated through the conditioning operation for a joint Gaussian distribution. That is,

$$\mathcal{N}\left(\left[\begin{array}{c}\mathbf{x}_I \\ \hline \mathbf{x}_E\end{array}\right]; \left[\begin{array}{c|c}\mathbf{P}_I & \mathbf{P}_{IE} \\ \hline \mathbf{P}_{IE}^T & \mathbf{P}_E\end{array}\right]\right) \stackrel{(\mathbf{x}_E^+, \mathbf{P}_E^+)}{\longrightarrow}$$

$$\mathcal{N}\left(\left[\begin{array}{c}\mathbf{x}_I + \mathbf{P}_{IE}\mathbf{P}_E^{+-}\mathbf{x}_E^+ \\ \hline \mathbf{x}_E^+\end{array}\right]; \left[\begin{array}{c|c}\mathbf{P}_I - \mathbf{P}_{IE}\mathbf{P}_E^{+-}\mathbf{P}_{EI} & \mathbf{P}_{IE} \\ \hline \mathbf{P}_{IE}^T & \mathbf{P}_E^+\end{array}\right]\right)$$

By iterating these steps, the hybrid R-B Inertial-SLAM can maintain the single vehicle EKF with a bank of parallel map-KFs, improving the computational burden in SLAM. Note that the data association problem is handled indirectly through the use of multiple hypotheses in parallel map-KFs.

6 Simulation Results

Computer simulation is performed to verify the proposed methods using a simulated flight data. The simulation parameters are shown in Table 1.

Sensor	Type	Unit	Spec.
IMU	Sampling rate	(Hz)	50
	Accel bias	(m/s^2)	0.1
	Gyro bias	$(^\circ/s)$	0.5
Range-Bearing Sensor	Sampling rate	(Hz)	50
	Range noise	(m)	$\geq$20
	Bearing noise	$(^\circ)$	0.16
	Elevation noise	$(^\circ)$	0.12

Table 1: Simulation parameters used.

In pose particles propagation, more higher process noises are used to relieve the particle depletion problem.

Figure 3 shows a 3-dimensional plot of the vehicle trajectory. In blue (dot-line) is the true trajectory and in

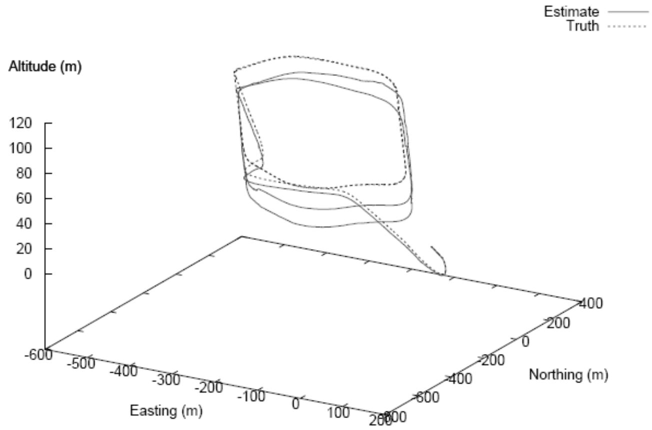

Figure 3: Simulated true vehicle trajectory (blue dot-line) with SLAM estimated one (red solid-line).

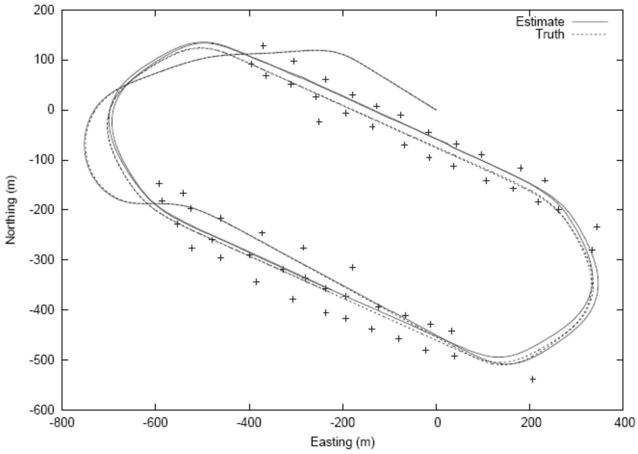

Figure 4: Estimated 2D vehicle trajectory with mean map positions.

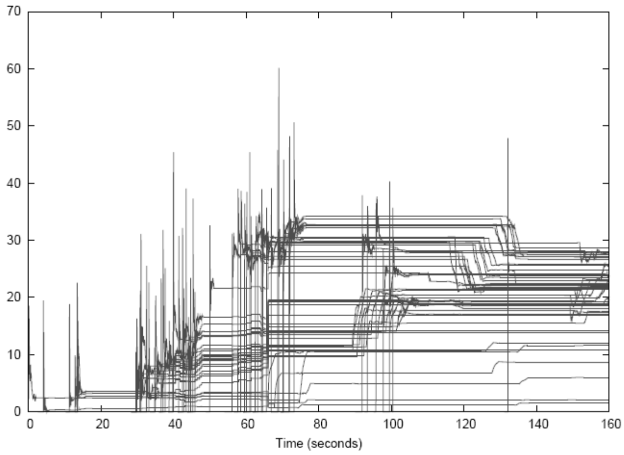

Figure 5: Evolution of average landmark position errors.

red (solid-line) is the filters estimate. It can be seen that the altitude estimate is falling below the trajectory. Figure 4 shows a 2-dimensional above-view of the trajectory. In blue is the true trajectory and landmarks. In red is the estimated trajectory. the estimated trajectory is in good agreement with the true trajectory.

Figure 5 shows the evolution of the landmark errors. The error is the Euclidean distance of the landmarks estimated position from its true position. It can be seen that up until loop closure (at 66 seconds) new landmarks are being registered. The final cluster of landmarks (registered from 55 to 73 seconds) have much larger errors than the cluster registered at 30 to 45 seconds. This is due to the inherited vehicle position error. From 120 to 135 seconds, the final cluster of landmarks is revisited and their estimates improve by more than 5 metres.

The vehicle-EKF is used solely for estimating the internal vehicle states. Figure 6 shows the vehicle-EKF

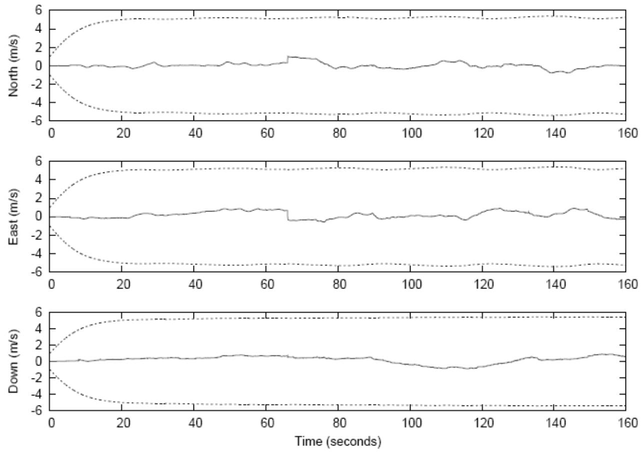

Figure 6: Estimated velocity error with 1-σ uncertainty bound within the vehicle-EKF.

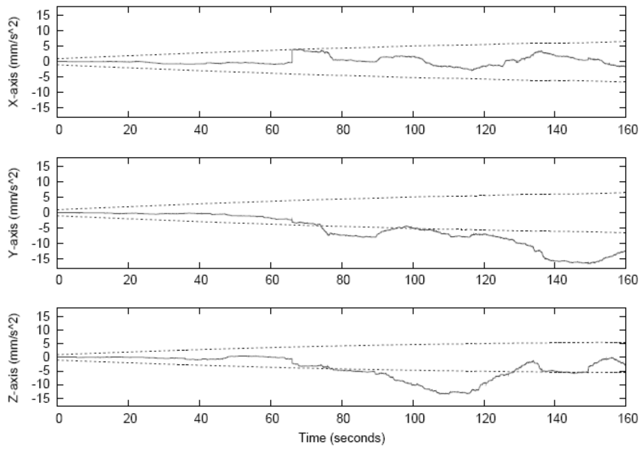

Figure 7: Estimated accelerometer bias error with 1-σ uncertainty bound within the vehicle-EKF.

velocity error along with their $1-\sigma$ uncertainty bounds. A correction can be seen at loop closure. The estimates stay within $2m/s$. Figure 7 shows the vehicle-EKF accelerometer bias error in $milli - metres/s^2$. The result is accurate but it needs more filter-tunings for convergence. Figure 8 shows the vehicle-EKF gyro-bias error in $milli-radians/s$. A significant correction can be seen at loop closure (66 seconds).

7 Conclusions

Although Rao-Blackwellised SLAM has been successful for mobile/ground robotics, its application to Inertial-SLAM system still suffers the high dimensionality in the vehicle state due to the use of inertial navigation system. This paper fills this gab by developing a hybrid R-B Inertial-SLAM which partitions the vehicle state into external and internal states and utilises the conditional

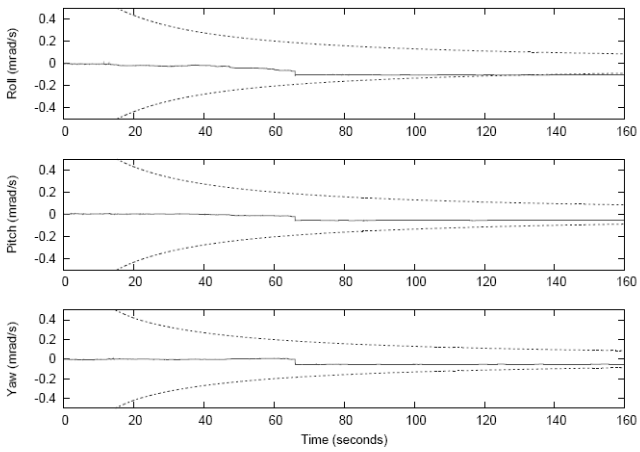

Figure 8: Estimated gyro bias error with 1-σ uncertainty bound within the vehicle-EKF.

independency between map-to-map. Simulation results with 50 pose particles showed reliable performances during loop-closures, thus significantly improving the performance. A more filter tuning is needed to confirm the bias observability and its application for the real-flight data is being conducted.

References

[Durrant-Whyte and Bailey, 2006] Hugh Durrant-Whyte and Tim Bailey. Simultaneous Localisation and Mapping (SLAM): Part I The Essential Algorithms. 2006.

[Estrada and Tardos, 2005] Neira J. Estrada, C. and J.D. Tardos. Hierarchical SLAM: Realtime accurate mapping of large environments. *IEEE Transactions on Robotics*, 21(4):588–596, 2005.

[Grisetti *et al.*, 2007] Giorgio Grisetti, Cyrill Stachniss, and Wolfram Burgard. Improved Techniques for Grid Mapping with Rao-Blackwellized Particle Filters. 23, 2007.

[Guivant and Nebot, 2001] J. Guivant and E. Nebot. Optimization of the simultaneous localisation and map building algorithm for real-time implementation. *IEEE Transactions on Robotics and Automation*, 17(3):242–257, 2001.

[Jung and Lacroix, 2003] I. Jung and S. Lacroix. High resolution terrain mapping using low altitude aerial stereo imagery. In *The Ninth IEEE International Conference on Computer Vision*, 2003.

[Kim and Sukkarieh, 2004] J. Kim and S. Sukkarieh. Autonomous Airborne Navigation in Unknown Terrain Environments. *IEEE Transactions on Aerospace and Electronic Systems*, 40(3):1031–1045, July 2004.

[Montemerlo *et al.*, 2004] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. Fast-SLAM 2.0: An improved particle filtering algorithm for simul- taneous localization and mapping that provably converges. In *International Joint Conference on Artificial Intelligence*, pages 1151–1156, 2004.

[Thrun *et al.*, 2002] S. Thrun, D. Koller, Z. Ghahramani, H. Durrant-Whyte, and A.Y. Ng. Simultaneous Mapping and Localization With Sparse Extended Information Filters. In *International Workshop on Algorithmic Foundations of Robotics*, 2002.

A INS state-space model

The vehicle state model is:

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k) + \mathbf{w}_k \Leftrightarrow \tag{15}$$

$$\begin{bmatrix} \mathbf{p}^n_{k+1} \\ \boldsymbol{\psi}^n_{k+1} \\ \hline \mathbf{v}^n_{k+1} \\ \mathbf{b}^b_{a,k+1} \\ \mathbf{b}^b_{g,k+1} \end{bmatrix} = \begin{bmatrix} \mathbf{p}^n_k + \triangle t \mathbf{v}^n_k \\ \boldsymbol{\psi}^n_k + \triangle t \mathbf{E}^n_b [\boldsymbol{\omega}^b_k + \mathbf{b}^b_{g,k} + \mathbf{w}_g] \\ \hline \mathbf{v}^n_k + \triangle t (\mathbf{C}^n_b [\mathbf{f}^b_k + \mathbf{b}^b_{a,k} + \mathbf{w}_a] + \mathbf{g}^n) \\ \mathbf{b}^b_{a,k} \\ \mathbf{b}^b_{g,k} \end{bmatrix}$$

where $\mathbf{C}^n_b$ and $\mathbf{E}^n_b$ are a body-to-navigation frame transformation and a Euler-rate conversion matrix, respectively [Kim and Sukkarieh, 2004].

B Observation model

A range/bearing sensor provides 3D observation model:

$$\begin{aligned} \mathbf{z}_{j,k} &= \mathbf{h}(\mathbf{x}_{E,k}, \mathbf{m}_{j,k}) + \mathbf{v}_k \\ &= (\mathbf{g}_2 \circ \mathbf{g}_1)(\mathbf{x}_{E,k}, \mathbf{m}_{j,k}), \end{aligned} \tag{16}$$

where $\circ$ is a composite-function operator, and $\mathbf{g}_2$ and $\mathbf{g}_1$ are

$$\mathbf{g}_2(k) = \begin{bmatrix} \sqrt{x^2 + y^2 + z^2} \\ \tan^{-1}(y/x) \\ \tan^{-1}(z/\sqrt{x^2 + y^2}) \end{bmatrix}$$

$$\mathbf{g}_1(k) = \mathbf{C}^s_b \mathbf{C}^b_n \left[\mathbf{m}_{i,k} - \mathbf{p}_k - \mathbf{C}^n_b \mathbf{p}^b_{sb} \right].$$

where $\mathbf{C}^s_b$ is a sensor-to-body transformation matrix.