

Conditions for Suboptimal Filter Stability in SLAM

Teresa Vidal-Calleja, Juan Andrade-Cetto and Alberto Sanfeliu
Institut de Robòtica i Informàtica Industrial, UPC-CSIC
Llorens Artigas 4-6, Barcelona, 08028 Spain
{tvidal, cetto, sanfeliu}@iri.upc.es

Abstract— In this article, we show marginal stability in SLAM, guaranteeing convergence to a non-zero mean state error estimate bounded by a constant value. Moreover, marginal stability guarantees also convergence of the Riccati equation of the one-step ahead state error covariance to at least one psd steady state solution. In the search for real-time implementations of SLAM, covariance inflation methods produce a suboptimal filter that eventually may lead to the computation of an unbounded state error covariance. We provide tight constraints in the amount of decorrelation possible, to guarantee convergence of the state error covariance, and at the same time, a linear-time implementation of SLAM.

I. INTRODUCTION

Simultaneous Localization and Map Building (SLAM) in mobile robotics has been an active research topic for over fifteen years. A Kalman filter (KF) based approach to the SLAM problem is adopted in this paper [1], [2]. One advantage in using such optimal state estimator, is that it is possible to show the convergence properties of SLAM; at least, for the linear case.

Under the typical fully correlated SLAM measurement model, it is not possible to obtain a zero mean state error estimate, unless partial observability is corrected [3]. Instead, one may obtain a constant bounded state estimate, dependant on the initial filter conditions. The reason being, that the filter used, one in which the vehicle and landmark estimates are stacked in the same state vector, is marginally stable. However, marginal stability does guarantee convergence of the state error covariance Riccati equation to at least one psd solution. The typical KF SLAM algorithm produces a relative map with zero uncertainty [2].

To speed up the performance of the algorithm, some authors have proposed the use of covariance inflation methods for the decorrelation of the state error covariance [4], subject to suboptimality of the filter. Adding pseudo-noise covariance for the landmark states is equivalent to making the system controllable. However, full decorrelation of a partially observable system might lead to filter instability [5]. In this communication we show how to diagonalize only part of the state error covariance to obtain a suboptimal filter that is both linear in time, and stable, at the same time.

The paper is structured as follows. Section II presents the model of the system and the Extended KF based

algorithm to the SLAM problem. Section III shows the convergence properties of both the state error and state error covariance for the partially observable SLAM system. In Section IV we show how as a consequence of having a partial controllability, filtering of the landmark state estimates is terminated after a small number of iterations, i.e., their corresponding Kalman gain terms tend to zero. In subsection IV-A we show a situation in which the filter becomes unstable during covariance inflation. In subsection IV-B we introduce a method for covariance decorrelation that preserves the stability of the filter. Furthermore, we show in subsection IV-C another solution for a stable covariance inflation algorithm, consisting on first recovering full observability prior to decorrelating the entire state error covariance matrix. Conclusions are presented in Section V.

II. KALMAN FILTER BASED SLAM

A. System Model

Formally speaking, the motion of the robot and the measurement of the map features are governed by the discrete-time state transition model

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, \mathbf{v}_k) \quad (1a)$$

$$\mathbf{z}_k = \mathbf{h}(\mathbf{x}_k, \mathbf{w}_k) \quad (1b)$$

The state vector \mathbf{x}_k contains the pose of the robot $\mathbf{x}_{r,k}$ at time step k , and a vector of stationary map features \mathbf{x}_f , i.e.,

$$\mathbf{x}_k = \begin{bmatrix} \mathbf{x}_{r,k} \\ \mathbf{x}_f \end{bmatrix} \quad (2)$$

The input vector \mathbf{u}_k is the vehicle control command, and \mathbf{v}_k and \mathbf{w}_k are Gaussian random vectors with zero mean and covariance matrices \mathbf{V} and \mathbf{W} , respectively, representing on the one side unmodeled robot dynamics and system noise; and measurement noise on the other side.

B. Algorithm

Provided the set of observations $Z^k = \{\mathbf{z}_1, \dots, \mathbf{z}_k\}$ was available for the computation of the current map estimate $\mathbf{x}_{k|k}$, the expression

$$\mathbf{x}_{k+1|k} = \mathbf{f}(\mathbf{x}_{k|k}, \mathbf{u}_k, \mathbf{0}) \quad (3)$$

gives an a priori noise-free estimate of the new locations of the robot and map features after the vehicle control command \mathbf{u}_k is input to the system. Similarly,

$$\mathbf{z}_{k+1|k} = \mathbf{h}(\mathbf{x}_{k+1|k}, \mathbf{0}) \quad (4)$$

This work was supported by the Spanish Council of Science and Technology under project DPI 2001-2223.

constitutes a noise-free a priori estimate of sensor measurements.

Given that the landmarks are considered stationary, their a priori estimate is simply $\mathbf{x}_{f,k+1|k} = \mathbf{x}_{f,k|k}$; and the a priori estimate of the map state error covariance showing the increase in robot localization uncertainty is

$$\mathbf{P}_{k+1|k} = E[\tilde{\mathbf{x}}_{k+1|k} \tilde{\mathbf{x}}_{k+1|k}^\top] \quad (5)$$

$$= \mathbf{F}\mathbf{P}_{k|k}\mathbf{F}^\top + \mathbf{G}\mathbf{V}\mathbf{G}^\top \quad (6)$$

The Jacobian matrices \mathbf{F} and \mathbf{G} contain the partial derivatives of \mathbf{f} with respect to \mathbf{x} and \mathbf{v} , evaluated at $(\mathbf{x}_{k|k}, \mathbf{u}_k, \mathbf{0})$.

Assuming that a new set of landmark observations \mathbf{z}_{k+1} coming from sensor data has been correctly matched to their map counterparts, one can compute the error between the measurements and the estimates with $\tilde{\mathbf{z}}_{k+1|k} = \mathbf{z}_{k+1} - \mathbf{z}_{k+1|k}$. This error aids in revising both the map and robot locations. The a posteriori state estimate is

$$\mathbf{x}_{k+1|k+1} = \mathbf{x}_{k+1|k} + \mathbf{K}\tilde{\mathbf{z}}_{k+1|k} \quad (7)$$

and the Kalman gain is computed with

$$\mathbf{K} = \mathbf{P}_{k+1|k}\mathbf{H}^\top\mathbf{S}^{-1} \quad (8)$$

where \mathbf{S} is termed the measurement innovation matrix,

$$\mathbf{S} = \mathbf{H}\mathbf{P}_{k+1|k}\mathbf{H}^\top + \mathbf{W} \quad (9)$$

and \mathbf{H} contains the partial derivatives of \mathbf{h} with respect to \mathbf{x} evaluated at $(\mathbf{x}_{k+1|k}, \mathbf{0})$.

Finally, the a posteriori estimate of the map state error covariance must also be revised once a measurement has taken place. It is revised with the Joseph form to guarantee positive semi-definiteness.

$$\mathbf{P}_{k+1|k+1} = (\mathbf{I} - \mathbf{K}\mathbf{H})\mathbf{P}_{k+1|k}(\mathbf{I} - \mathbf{K}\mathbf{H})^\top + \mathbf{K}\mathbf{W}\mathbf{K}^\top \quad (10)$$

III. CONVERGENCE

Substituting the linearized version of (3) in (7), we may rewrite the KF in the one-step ahead prediction form

$$\mathbf{x}_{k+1|k} = (\mathbf{F} - \mathbf{K}\mathbf{H})\mathbf{x}_{k|k-1} + \mathbf{K}\mathbf{z}_k \quad (11)$$

and with the appropriate substitutions, using (5) and (11), the corresponding prediction error dynamics becomes

$$\tilde{\mathbf{x}}_{k+1|k} = (\mathbf{F} - \mathbf{K}\mathbf{H})\tilde{\mathbf{x}}_{k|k-1} + \mathbf{G}\mathbf{v}_k - \mathbf{K}\mathbf{w}_k \quad (12)$$

In general, only for a stable matrix $\mathbf{F} - \mathbf{K}\mathbf{H}$, the estimation error will converge to a zero mean steady state value. However, in SLAM, $\mathbf{F} - \mathbf{K}\mathbf{H}$ is marginally stable, thus the steady state error estimate is bounded to a constant value, subject to the filter initial conditions. To show $\mathbf{F} - \mathbf{K}\mathbf{H}$ marginally stable, consider the one landmark monobot. $\mathbf{F} = \mathbf{I}$, $\mathbf{G} = [1 \ 0]^\top$, and $\mathbf{H} = [-1 \ 1]$. For any value of

$$\mathbf{P} = \begin{bmatrix} \sigma_r^2 & \rho\sigma_r\sigma_f \\ \rho\sigma_r\sigma_f & \sigma_f^2 \end{bmatrix} \quad (13)$$

the Kalman gain, computed with (8), is

$$\mathbf{K} = \frac{1}{s} \begin{bmatrix} -\sigma_v^2 \\ \sigma_w^2 \end{bmatrix} \quad (14)$$

where

$$s = \sigma_r^2 + \sigma_f^2 - 2\rho\sigma_r\sigma_f + \sigma_w^2 \quad (15)$$

is the innovation variance. Consequently,

$$\mathbf{F} - \mathbf{K}\mathbf{H} = \frac{1}{s} \begin{bmatrix} -\sigma_v^2 + s & \sigma_v^2 \\ \sigma_w^2 & -\sigma_w^2 + s \end{bmatrix} \quad (16)$$

with eigenvalues

$$\left\{ \begin{array}{l} 1, \\ \frac{1}{s}(s - \sigma_v^2 - \sigma_w^2) \end{array} \right\} \quad \text{and} \quad s \neq 0$$

One of the eigenvalues being on the unitary circle yields marginal stability, i.e., constant bounded non-zero mean error state estimate convergence. Moreover, the marginal stability of $\mathbf{F} - \mathbf{K}\mathbf{H}$ guarantees at least one *psd* steady state solution to the Riccati equation for the one-step ahead state error covariance [6]

$$\mathbf{P}_{k+1|k} = (\mathbf{F} - \mathbf{K}\mathbf{H})\mathbf{P}_{k|k-1}(\mathbf{F} - \mathbf{K}\mathbf{H})^\top + \mathbf{G}\mathbf{V}\mathbf{G}^\top + \mathbf{K}\mathbf{H}\mathbf{W}\mathbf{H}^\top\mathbf{K}^\top \quad (17)$$

IV. PARTIAL CONTROLLABILITY

The dynamics of the model assume the landmarks are fixed elements, for which no process noise is considered. Therefore, their associated noise covariance (its determinant) will asymptotically tend to zero [2]. The filter gain for the landmark states will also tend to zero.

Given that

$$\mathbf{F} = \begin{bmatrix} \mathbf{F}_{\mathbf{x}_r} \\ \mathbf{I} \end{bmatrix}$$

and

$$\mathbf{G} = \begin{bmatrix} \mathbf{G}_v \\ \mathbf{0} \end{bmatrix},$$

the controllability matrix for such plant is

$$\mathcal{C} = [\mathbf{G}\mathbf{V}^{1/2} \mid \dots \mid \mathbf{F}^{\dim \mathbf{x}-1}\mathbf{G}\mathbf{V}^{1/2}] \quad (18)$$

Consequently, the dimensionality of the controllable subspace, spanned by the column space of \mathcal{C} is $\text{rank } \mathcal{C} = \dim \mathbf{x}_r$, regardless of the number of landmarks in the map. Obviously, the only controllable states are the ones associated with the vehicle motion.

When a stochastic system is partially controllable, such as in the case of SLAM, the Gaussian noise sources \mathbf{v}_k do not affect all of the elements of the state space. That is, some states are uncorrupted by the noise. The diagonal elements of \mathbf{P} corresponding to these uncorruptible states will be driven to zero by the Kalman filter, and once this happens, these estimates will remain fixed and no further observations will alter their values. Figure 1 shows two simulations for a linear SLAM case, a monobot with one and two landmarks. The purpose is to show the evolution of the localization errors for both the monobot and the landmarks, and the reduction of the landmark part of the Kalman gain, due to the uncontrollability of the system. The only way to remedy this situation is to add a positive

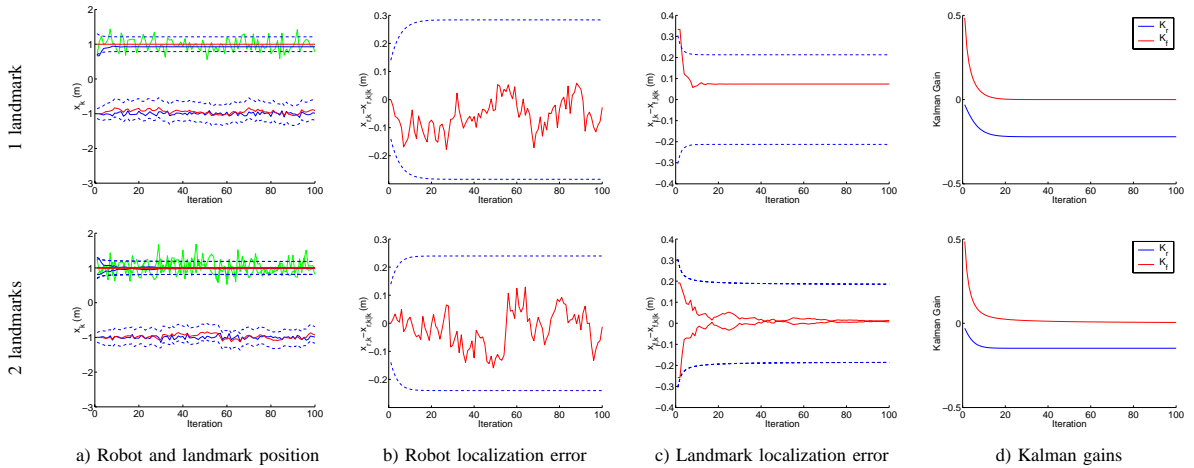


Fig. 1. Partially observable SLAM for a monobot during Brownian motion with 100 iterations. The vehicle location is indicated by the darkest curve at the $-1m$ level in the first column of plots. In the same set of figures, and close to it, is a lighter curve indicating the vehicle location estimate as computed by the filter, along with 2σ bounds on such estimate shown as dotted lines. The dark straight lines at the $1m$ level indicate the landmark location estimates; and the lighter curves are noise corrupted signals of sensor measurements. Also shown, are a pair of dotted lines for 2σ bounds on the landmark location estimates. The second column of plots shows the vehicle location error only, and its corresponding variance, also on the form of 2σ dotted bounds. See how the localization error has non-zero mean due to partial observability, an undesirable feature in Kalman filtering. The third column shows non-zero mean landmark state estimate errors. And, the last column shows the Kalman filter gains both for the vehicle and landmark revision terms. The Kalman gains for the revision of the landmark estimates rapidly tend to zero, the reason being that these states are uncontrollable.

definite pseudo-noise covariance to those uncorruptible states [7].

A. $O(N)$ but unstable partially observable SLAM

A clever way to add pseudo-noise to the model is by diagonalizing the state error covariance matrix [4], [5], [8]. The result is a suboptimal filter that will compute inflated estimates for the vehicle and landmark covariances, that has the computational advantage of being uncorrelated. The addition of a covariance term $\Delta\mathbf{P}$ to the a priori state covariance estimate

$$\mathbf{P}_{k+1|k} = \mathbf{F}\mathbf{P}_{k|k}\mathbf{F}^T + \mathbf{G}\mathbf{V}\mathbf{G}^T + \Delta\mathbf{P} \quad (19)$$

is equivalent to providing a new form to the plant noise Jacobian $\mathbf{G}' = \begin{bmatrix} \mathbf{G} & \mathbf{I} \end{bmatrix}$

$$\mathbf{P}_{k+1|k} = \mathbf{F}\mathbf{P}_{k|k}\mathbf{F}^T + \mathbf{G}' \begin{bmatrix} \mathbf{V} \\ \Delta\mathbf{P} \end{bmatrix} \mathbf{G}'^T \quad (20)$$

$\Delta\mathbf{P}$ may be chosen, for example, such as to minimize the trace of a resulting block diagonal \mathbf{P} in (19) (see [5]).

Choosing a full rank $\Delta\mathbf{P}$ is equivalent to having noise input to more states than those that can be observed with the filter. In this case, because of partial observability, both vehicle and landmark variance estimates become unbounded. Figure 2 shows this for the same monobot experiment as in the previous simulation. This phenomena was first observed in [5] using relative maps.

Not only both the vehicle and landmark state estimation variances become unbounded. Also, thanks to the full controllability of the system, the Kalman gain for the revision of the landmark states is greater than zero; but still, does not converge to a steady state value. We believe that the addition of pseudo-noise should be performed only at most, in the amount of states equal to the dimension of the observable subspace.

B. $O(N)$ and stable partially observable SLAM

One solution to the problem of instability during covariance inflation, is to decorrelate only the landmark state estimates, and to preserve all vehicle to landmark correlations.

$$\Delta\mathbf{P} = \begin{bmatrix} \mathbf{0} & \\ & \mathbf{V}_f \end{bmatrix} \quad (21)$$

such that $\mathbf{P}_f + \mathbf{V}_f$, the map part of the state error covariance, is block diagonal.

Figure 3 shows a partially observable monobot under Brownian motion for which only the landmark part of the state error covariance matrix has been decorrelated. The algorithm does converge to a steady state solution under this circumstances, and still can be implemented in real time. The one landmark case is identical than the original case, since a linear one landmark map is already diagonal (scalar actually).

For the two-landmark case, the landmark variance estimate is greater than the optimal solution shown in the third column in Figure 1. That is, the covariance has been inflated during decorrelation. Furthermore, now that the system is controllable, the Kalman gains for the landmark state estimates do not become zero, and they converge to a steady state value.

Moreover, we can see experimentally, that the covariance inflation suboptimal partially observable SLAM converges only when

$$\text{rank } \Delta\mathbf{P} \leq \text{rank } \mathcal{O} \quad (22)$$

C. $O(N)$ and stable fully observable SLAM

Considering the fully observable case [3], even if we add pseudo-noise to the vehicle as well as to the landmark states, the covariance will reach a steady-state value, and the Kalman gain will not be zero, at least, in the linear

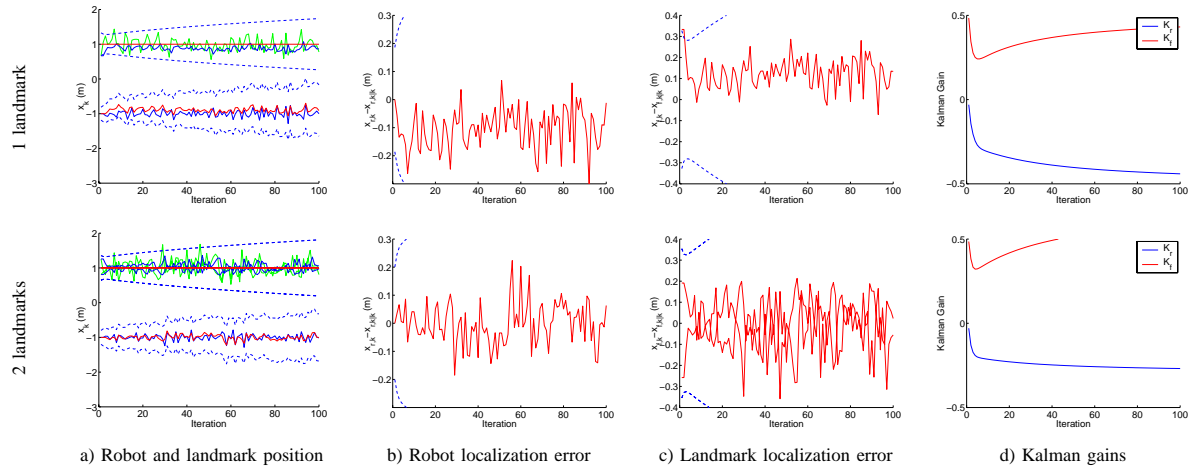


Fig. 2. Partially observable SLAM for a Brownian motion monobot with 100 iterations. The entire state error covariance is decorrelated with the minimal trace solution [5]. By decorrelating the entire state error covariance matrix, the covariance estimates become unbounded.

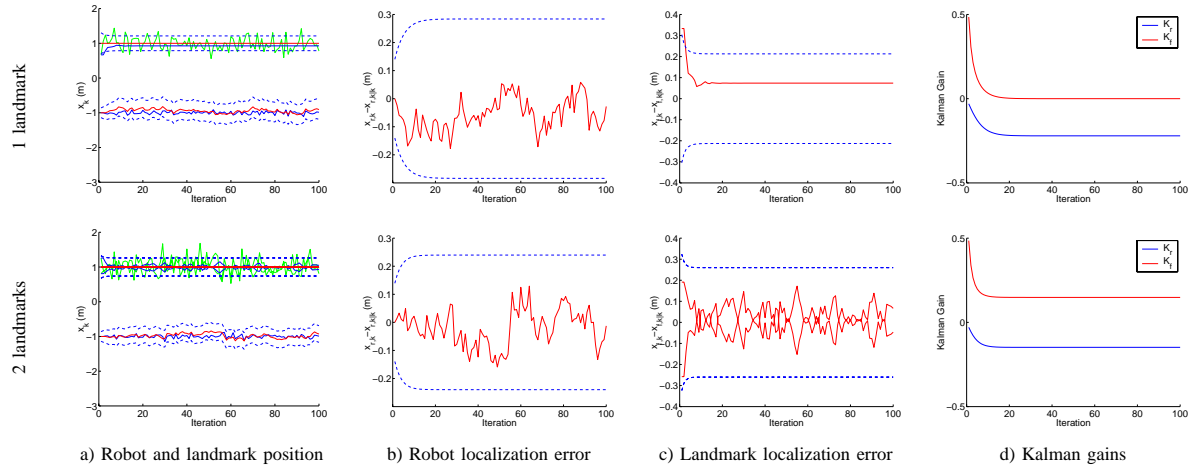


Fig. 3. Partially observable SLAM for a Brownian motion monobot with 100 iterations. The state error covariance is decorrelated only for the landmark part of the state vector, with the minimal trace solution. By decorrelating only the map part of the state error covariance matrix, we preserve filter stability.

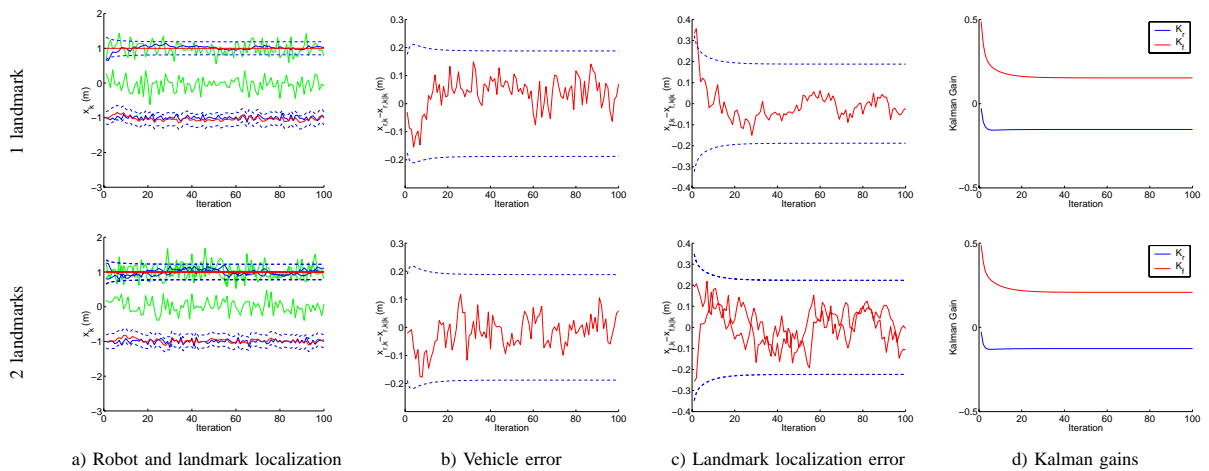


Fig. 4. Fully observable SLAM for a Brownian motion monobot with 100 iterations. The entire state error covariance is decorrelated with the minimal trace solution. In the linear case, it is possible to decorrelate the entire state error covariance matrix, and still preserve filter stability.

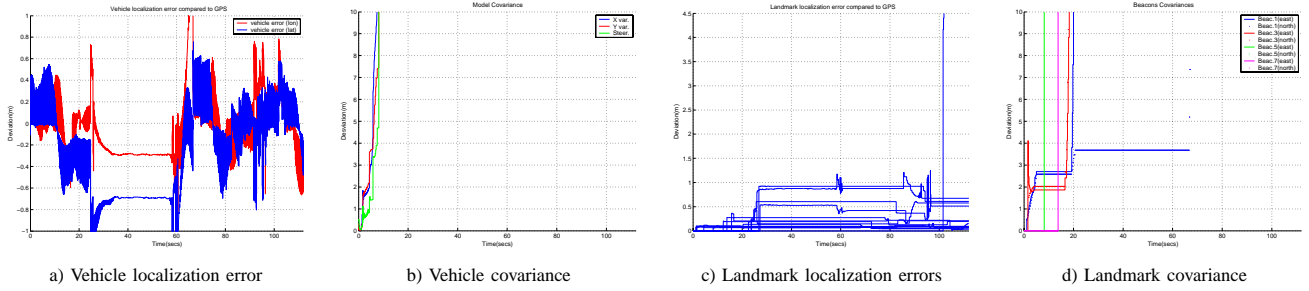


Fig. 5. Partially observable SLAM for a car-like vehicle at the University of Sydney Car Park. The entire state error covariance matrix is decorrelated with the minimal trace solution [5].

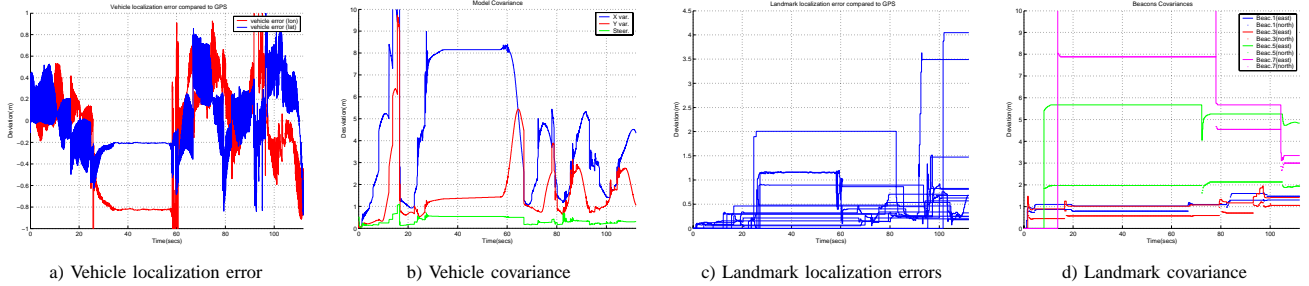


Fig. 6. Partially observable SLAM for a car-like vehicle at the University of Sydney Car Park. Only the map part of the state error covariance matrix is decorrelated with the minimal trace solution. By adding controllability to as many states as those that are observable, the filter remains stable, and the estimated covariances remain bounded.

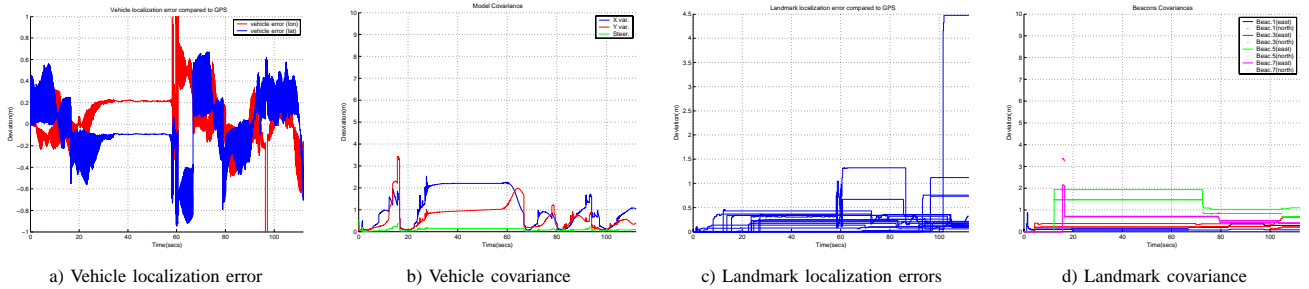


Fig. 7. Fully observable SLAM for a car-like vehicle at the University of Sydney Car Park. Only the map part of the state error covariance is decorrelated with the minimal trace solution. Full observability guarantees independence of the filter initial conditions, and an accurate absolute map is obtained, with smaller covariance estimates than its relative counterpart.

case. Figure 4 shows this results diagonalizing the whole state error covariance (not only the landmark part of \mathbf{P}).

In this latter experiment, the state error variances reach lower values than those in the partially observable case. The solution of the Riccati equation is now independent of the initial covariance estimate $\mathbf{P}_{0|0}$.

We have observed experimentally however, that with a nonlinear vehicle model, it is best to also decorrelate only the map part of the state error covariance, even in the fully observable case.

D. Experimental Results

We show now results on a series of experiments for nonlinear vehicle with an also nonlinear measurement model, using the ACFR - University of Sydney database [9]. The test run used corresponds to a car-like vehicle at the University Car Park. The landmarks used are tree trunks, as measured with a laser range finder. The reconstructed

maps are compared to GPS ground truth for accuracy. The first experiment corresponds to a typical partially observable SLAM run, in which the entire state error covariance is being decorrelated as discussed in Section IV-A. Figure 5 plots results on this run, showing in column b) and d) unbounded covariances both for the vehicle and landmark state estimates, due to the naïve covariance inflation method used.

The second experiment corresponds to the same partially observable SLAM conditions, but decorrelating only the map part of the state error covariance. Adding pseudo-noise to the landmark states during the inflation procedure amounts to making the system controllable; and doing so for as many states as those observable, produces both vehicle and landmark bounded state covariances estimates. This is shown in Figure 6, columns b) and d). Figure 8 column a) shows the actual vehicle path and landmark location estimates recovered by the algorithm, compared

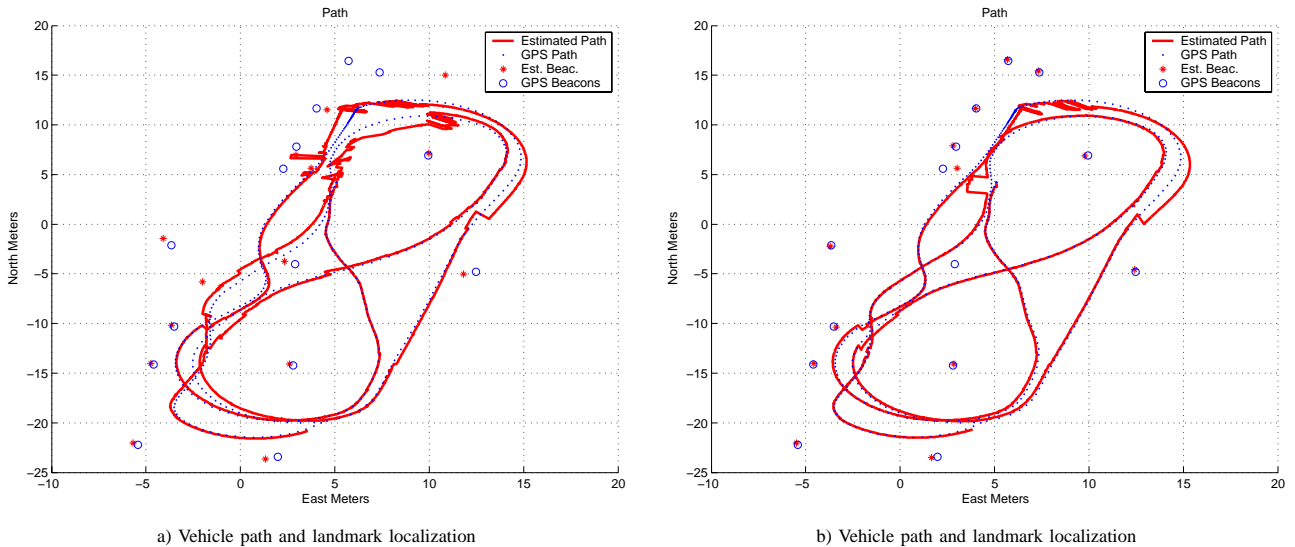


Fig. 8. Vehicle path and landmark location estimates, compared to GPS ground truth for an a) partially observable suboptimal SLAM run, and a b) fully observable suboptimal SLAM run; both with decorrelation of only the map part of the state error covariance matrix.

to GPS ground truth for the beacons. Note that even when the “relative” map is consistent [2], it is slightly rotated and shifted from the actual beacon locations. The amount of this shift depends on the initial vehicle uncertainty, i.e., the initial filter conditions, and can be seen in Figure 6, column c).

The last experiment shown corresponds to a fully observable SLAM run (using the first observed beacon as an anchor [3]), and also decorrelating only the map part of the state error covariance. In this case, the vehicle and landmark covariance estimates do not depend on the initial filter conditions, and thus are significantly reduced. This is shown in columns b) and d) in Figure 7. The absolute landmark estimate error is also significantly reduced, as shown in Figure 7, column c). Figure 8 column b) shows the actual vehicle path and landmark estimates as recovered by the filter. The beacon shown in the center of the plot is used as an anchor to the map, and no state estimate is computed for it. This last map was obtained with a suboptimal linear-time SLAM algorithm that has both bounded covariance estimates, and independence on the filter initial conditions; thus producing a fast and accurate absolute map.

V. CONCLUSIONS

A unit norm eigenvalue for the matrix $\mathbf{F} - \mathbf{K}\mathbf{H}$ makes the state error estimate converge to a non zero mean constant bounded value in the linear case SLAM. Marginal stability of such partially observable system produces also at least one *psd* solution to the steady state Riccati equation for the covariance error, provided the initial conditions of \mathbf{P} are also *psd*.

Suboptimal techniques to improve the speed of the algorithm include covariance inflation methods to diagonalize the state error covariance matrix. These techniques may lead to instability if pseudo-noise is added in a higher state dimensionality than what can be observed. We propose to

diagonalize only the map part of the state error covariance, thus guaranteeing convergence of \mathbf{P} , and at the same time obtaining an $O(N)$ algorithm.

REFERENCES

- [1] T. Duckett and U. Nehmzow, “Mobile robot self-localization and measurement of performance in middle-scale environments,” *Robot. Auton. Syst.*, vol. 24, no. 1-2, pp. 57–69, Aug. 1998.
- [2] M. W. M. G. Dissanayake, P. Newman, S. Clark, H. F. Durrant-Whyte, and M. Csorba, “A solution to the simultaneous localization and map building (SLAM) problem,” *IEEE Trans. Robot. Automat.*, vol. 17, no. 3, pp. 229–241, Jun. 2001.
- [3] J. Andrade-Cetto and A. Sanfeliu, “The effects of partial observability in SLAM,” in *Proc. IEEE Int. Conf. Robot. Automat.*, New Orleans, Apr. 2004, pp. 397–402.
- [4] J. E. Guivant and E. M. Nieto, “Solving computational and memory requirements of feature-based simultaneous localization and mapping algorithms,” *IEEE Trans. Robot. Automat.*, vol. 19, no. 4, pp. 749–755, Aug. 2003.
- [5] S. J. Julier, “The stability of covariance inflation methods for SLAM,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Las Vegas, Oct. 2003, pp. 2749–2754.
- [6] T. Kailath, A. H. Sayed, and B. Hassibi, *Linear Estimation*, ser. Information and System Sciences Series, T. Kailath, Ed. Upper Saddle River: Prentice Hall, 2000.
- [7] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York: John Wiley & Sons, 2001.
- [8] J. E. Guivant and E. M. Nieto, “Optimization of simultaneous localization and map-building algorithm for real-time implementation,” *IEEE Trans. Robot. Automat.*, vol. 17, no. 3, pp. 242–257, Jun. 2001.
- [9] E. Nebot, J. Guivant, and J. Nieto, “ACFR, experimental outdoor dataset,” 2002. [Online]. Available: <http://www.acfr.usyd.edu.au/homepages/academic/enebot/dataset.htm>