A Closed Form Solution to the Single Degree of Freedom Simultaneous Localisation and Map Building (SLAM) Problem

Peter W. Gibbens, Gamini M.W.M. Dissanayake, and Hugh F. Durrant-Whyte
Australian Centre for Field Robotics
Department of Aerospace, Mechanical and Mechatronic Engineering
The University of Sydney
NSW 2006, Australia
pwg, dissa, hugh@acfr.usyd.edu.au

Abstract

This paper presents a closed form solution to the estimation-theoretic simultaneous localisation and map building (SLAM) problem. The solution is obtained by explicit solution of the differential Riccati equation associated with the n-landmark SLAM problem. The solution describes and explains the many experimental and theoretical results obtained so far in the study of the SLAM problem. Further, the solution, for the first time, allows a precise means of analysing the performance of different SLAM algorithms and enables the design of efficient SLAM systems.

1 Introduction

The solution to the simultaneous localisation and map building (SLAM) problem is, in many respects a "Holy Grail" of autonomous vehicle navigation research. The ability to place an autonomous vehicle at an unknown location in an unknown environment and then have it build a map, using only relative observations of the environment, while simultaneously using the map to navigate would indeed make such a vehicle "autonomous". A solution to the SLAM problem will be of inestimable value in a range of applications where absolute position or precise map information is unobtainable.

An estimation-theoretic or Kalman-filter based approach to the SLAM problem is adopted in this paper. The study of estimation-theoretic solutions to the SLAM problem has a long history starting with the key paper by Smith and Cheeseman [10] (see also [4, 7, 8]. However, after a long haitus, there has recently been an explosion of interest and a number of substantial results on the estimation-theoretic SLAM problem. In particular, the work by Leonard on large-scale SLAM problems [9, 6], the experimental work of Castellanos and Tardos [3, 2], and the work of Dissanayake and Durrant-Whyte on convergence [5]. All of this work shows that, while many problems remain, a solution to the estimation-theoretic SLAM problem is possible both theoretically and practically.

The key issue in solving the SLAM problem is that an error in estimated vehicle location leads to a common error in the estimated location of environment landmarks. Indeed, it is possible to show that the correlation caused by this common error between landmarks tends to unity, and thus in the limit a *perfect* relative map of landmarks can be constructed [5]. However, the need to maintain this correlation as an integral part of the SLAM solution leads to enormous computational problems, as the location of each landmark in the environment must, in theory, be updated at each step in the estimation cycle. This leads inexorably to a need to find effective map-management policies for large scale problems [9]. The key to all these problems is to understand the structure and evolution of the covariance matrix for the SLAM problem. This contains the uncertainties in estimated vehicle and landmark location together with the (all important) covariances between landmarks, and landmarks and vehicle.

The main contribution of this paper is to provide, for the first time, a closed form solution to the SLAM problem. The solution is found by explicit solution of the corresponding Riccati equation describing the structure and evolution of the SLAM covariance matrix. The solution provides a time-domain description of the evolution of the single degree-of-freedom n-landmark SLAM problem for the case where the motion of the vehicle is Brownian, the landmarks are scalar and the observations are linear. While this is a restricted case, it nevertheless exposes the structure and evolution process of more general SLAM problems. It can therefore be used as an indicator of the performance of more complex SLAM problems including multi-dimensional problems, nonlinear and asynchronous range bearing observation problems, and more complex vehicle motion models.

It can also illustrate short-term SLAM performance in two or three dimensional problems where the vehicle motion is nominally linear and the feature distribution is not changing rapidly. This is the subject of ongoing investigation.

The solution provides an explicit picture of how and why the SLAM principle works, the effect that different vehicle uncertainties have on the solution, and how the observation of different landmarks effects the convergence and steady-state behaviour of the problem. Most importantly, the solution also provides, for the first time, an explicit means of developing and analysing efficient SLAM algorithms and of designing a range of multi-sensor SLAM systems.

This paper begins by establishing and formulating the SLAM problem to be solved. The method of solution and the solution itself are then presented. The structure of the solution is then discussed and explained in terms of previously known results. An analysis of the properties of the solution is then presented. This provides a study of the convergence properties of the SLAM algorithm, the effect of vehicle model uncertainty and the effects of observation uncertainties in different landmarks on convergence rates and steady state performance. In conclusion, the paper discusses the practical implications of this solution.

2 Formulation of the SLAM Problem

The general estimation-theoretic SLAM problem considers the motion of a platform in space, taking relative observations to a number of fixed landmarks. In most practical developments, the SLAM algorithm is described in discrete time. However, for the purposes of obtaining a closed-form solution to the problem, it is most valuable to consider the associated continuous time problem. In this paper, a simple one-dimensional version of this SLAM problem is considered.

2.1 Process and Observation models

A platform is in linear motion. Generally measurements of velocity $u(t)$ are made (from an inertial measurement unit, Doppler or encoder system). These measurements are subject to errors w, assumed to be zero-mean white noise of variance q. The position $x(t)$ of the platform is constrained to a single dimension and is thus described by $\dot{x}(t) = u(t) + w$.

The environment is populated by an arbitrary number of stationary landmarks $i = 1, \cdots, n$ placed at fixed locations p_i (in the single motion direction) , so that $\dot{p}_i = 0$. A combined state vector is defined consisting of the location of the platform and the location of all landmarks in the form $\mathbf{x}(t) = [x(t), p_1, \cdots, p_n]^T$. The general state-transition equation can now be written in the usual form

$$\dot{\mathbf{x}}(t) = \mathbf{F}\mathbf{x}(t) + \mathbf{G}(\mathbf{u}(t) + \mathbf{w}(t)) \tag{1}$$

where clearly $\mathbf{F} = \mathbf{0}$ and $\mathbf{G} = [1, 0, \cdots, 0]^T$.

The platform is equipped with a sensor that can measure the relative range between the platform and landmarks. An observation z_i to the i^{th} landmark is described by $z_i(t) = p_i - x(t) + v_i$ where v_i is assumed to be a zero mean white process of variance r_i describing range measurement errors. These (continuous time) observations can be placed in a single observation model of the form

$$\mathbf{z}(t) = \mathbf{H}\mathbf{x}(t) + \mathbf{v} \tag{2}$$

where $\mathbf{z}(t) = [z_1(t), \cdots, z_n(t)]^T$, $\mathbf{v} = [v_1, \cdots, v_n]^T$, and

$$\mathbf{H} = \begin{bmatrix} -1 & 1 & 0 & \cdots & 0 \\ -1 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -1 & 0 & 0 & \cdots & 1 \end{bmatrix} \tag{3}$$

is the $n \times (n+1)$ observation matrix.

The process noise variance and observation noise variance are both assumed time-invariant and are defined by $\mathbf{Q} = \mathrm{E}[\mathbf{w}\mathbf{w}^T] = q$ and $\mathbf{R} = \mathrm{E}[\mathbf{v}\mathbf{v}^T] = \mathrm{diag}(r_1, \cdots, r_n)$. Not surprisingly, the quantities $\mathbf{Q}$ and $\mathbf{R}$ are the essential parameters governing the performance of the SLAM process.

The goal of a SLAM algorithm is to compute an estimate $\hat{\mathbf{x}}$ of the location of both the vehicle and landmarks given the complete observation signal Z^t up to time t. For a problem such as this, the estimate is normally considered to be the conditional mean; $\hat{\mathbf{x}}(t) = \mathrm{E}[\mathbf{x}(t) \mid Z^t]$. The error in this estimate is denoted $\tilde{\mathbf{x}}(t) = \hat{\mathbf{x}}(t) - \mathbf{x}(t)$. The corresponding state estimate covariance matrix is given by $\mathbf{P}(t) = \mathrm{E}[\tilde{\mathbf{x}}(t)\tilde{\mathbf{x}}^T(t) \mid Z^t]$. This covariance matrix describes the evolution of the vehicle and landmark variances together with all covariances between landmarks and vehicle and landmarks. The evolution of this covariance matrix is governed by the differential Riccati equation in the form

$$\dot{\mathbf{P}}(t) = \mathbf{F}\mathbf{P}(t) + \mathbf{P}(t)\mathbf{F}^T + \mathbf{G}\mathbf{Q}\mathbf{G}^T - \mathbf{P}(t)\mathbf{H}^T\mathbf{R}^{-1}\mathbf{H}\mathbf{P}(t) \tag{4}$$

This paper describes a solution to Equation 4 for the system described by Equations 1 and 2.

2.2 Solution Method

The differential Riccati equation is non-linear. However, there are a number of well known methods for obtaining solutions to this equation in specific cases (see [1] for a summary of methods). The method adopted here is to describe the solution to the Riccati equation in the form

$$\mathbf{P}(t) = \mathbf{U}(t)\mathbf{V}^{-1}(t) \tag{5}$$

where $\mathbf{U}(t)$ and $\mathbf{V}(t)$ are matrices satisfying the homogeneous linear differential equation and the initial condition

$$\begin{bmatrix} \dot{\mathbf{U}}(t) \\ \dot{\mathbf{V}}(t) \end{bmatrix} = \begin{bmatrix} \mathbf{F} & \mathbf{G}\mathbf{Q}\mathbf{G}^T \\ \mathbf{H}^T\mathbf{R}^{-1}\mathbf{H} & -\mathbf{F}^T \end{bmatrix} \begin{bmatrix} \mathbf{U}(t) \\ \mathbf{V}(t) \end{bmatrix} \tag{6}$$

$$\begin{bmatrix} \mathbf{U}(0) \\ \mathbf{V}(0) \end{bmatrix} = \begin{bmatrix} \mathbf{P}(0) \\ \mathbf{I} \end{bmatrix} \tag{7}$$

where $\mathbf{V}(t)$ is always invertible, $\mathbf{P}(0)$ is the initial state covariance matrix and $\mathbf{I}$ is the identity matrix with appropriate dimensions. It is straight-forward to show (by substitution) that this is equivalent to solving the Riccati equation. The transition matrix for this problem is known as the Hamiltonian matrix.

For the SLAM problem as stated, Equation 6 can be substantially simplified. In particular, $\mathbf{F} = 0$, $\mathbf{Q} = q$ and $\mathbf{R} = \text{diag}(r_1, \cdots, r_n)$, so that

$$\begin{aligned} \dot{\mathbf{U}}(t) &= \mathbf{G}\mathbf{Q}\mathbf{G}^T\mathbf{V}(t) \\ \dot{\mathbf{V}}(t) &= \mathbf{H}^T\mathbf{R}^{-1}\mathbf{H}\mathbf{U}(t) \end{aligned} \tag{8}$$

Furthermore, for the case considered, where the vehicle position estimation and landmark observations commence at time $t = 0$, the initial covariance matrix is given by $\mathbf{P}(0) = \text{diag}(0, r_1, \cdots, r_n)$.

2.3 The n Landmark Solution

Practically, a solution to Equation 8 was found by using a commercial symbolic algebra package to obtain solutions for the cases $n = 1, 2, 3$. From these, the solution for arbitrary n was hypothesised and shown to satisfy the original Riccati equation.

The general time-domain solution for the state covariance matrix for the general n-landmark scalar SLAM problem with $\mathbf{P}(0)$ as defined is given by Equation 9 (overleaf) where $D(t) = (\alpha + 1) + (\alpha - 1)e^{-2\alpha t}$ is the characteristic equation of the system, $I_T = \sum_{i=1}^{n} r_i^{-1}$ is the total Fisher information available to the filter, and $\alpha = \sqrt{qI_T}$ is the dominant time constant for the system.

3 Analysis of the Solution

The covariance matrix given in Equation 9 encapsulates in full the structure and evolution of the linear, scalar SLAM system. Section 3 of this paper provides a detailed analysis of the structure and convergence properties of this equation. Here we identify some principal features of the solution.

The SLAM solution is primarily governed by the total Fisher information I_T, the process noise covariance q and the time constant α (which is itself a function of I_T and q). The dependence on Fisher information is logical; this is a measure of the total information (per unit time) available to the filter to generate estimates of landmark and platform positions. The dependence on q is also logical; this is a measure of the loss of information from the filter through uncertain platform motion. The time constant $\alpha = \sqrt{qI_T}$ is thus precisely the ratio of information loss to information gain and therefore governs both the rate of convergence and steady state performance of the SLAM system.

Herein, for convenience, we denote the variance in estimated vehicle position as $P_{vv}(t)$, the variance in the estimated i^{th} landmark position as $P_{ii}(t)$, the covariance between the platform and the i^{th} landmark position errors as $P_{vi}(t)$, and the covariance between the i^{th} and j^{th} landmark position errors as $P_{ij}(t)$. Note also that due to symmetry $P_{vi}(t) = P_{iv}(t)$ and $P_{ij}(t) = P_{ji}(t)$.

The platform variance is clearly proportional to the process nose variance q, as are the landmark to platform cross-correlation terms. It is interesting to note that the vehicle to landmark covariances are not a function of the specific landmark; that is the correlation represents a platform to map, and not a platform to specific landmark relationship (this has been observed experimentally in earlier work).

According to Equation 9, the landmark variance consists of two components, one converges as $1/t$, the other converges exponentially. The exponential component is independent of the specific landmark and converges to the same value as the platform to landmark covariance. This component essentially describes the evolution of the "whole map" variance. The former component is a function of the specific landmark but it converges to zero. However, the convergence of this component is slow and dominates the convergence of the overall map building process. This component explains the many poor convergence rates observed in earlier experimental work. The term itself is in the form of a mutual-information gain equation and would appear to represent the *diffusion* of information from the landmark to the rest of the map.

The landmark to landmark covariance is identical for every landmark pair. This cross covariance again has two components, one which converges as $1/t$ and one which converges exponentially. The exponential term, which is the same as for the landmark variance, may be considered as the common "whole map" variance and thus converges to the same value as the corresponding platform to map term. The $1/t$ component converges to zero but, as before, dominates the convergence process.

$$
\mathbf{P}(t) = \begin{bmatrix}
\frac{q(1-e^{-2\alpha t})+\frac{2q}{\alpha}(1-e^{-\alpha t})^2}{\mathbf{D}(t)} & \frac{\frac{q}{\alpha}(1-e^{-\alpha t})^2}{\mathbf{D}(t)} & \cdots & \frac{\frac{q}{\alpha}(1-e^{-\alpha t})^2}{\mathbf{D}(t)} & \cdots \\
\frac{\frac{q}{\alpha}(1-e^{-\alpha t})^2}{\mathbf{D}(t)} & \frac{r_1(I_T - r_1^{-1})}{(t+1)I_T} + \frac{\frac{q}{\alpha}(1+e^{-2\alpha t})}{\mathbf{D}(t)} & \cdots & -\frac{1}{(t+1)I_T} + \frac{\frac{q}{\alpha}(1+e^{-2\alpha t})}{\mathbf{D}(t)} & \cdots \\
\vdots & \vdots & \ddots & \vdots & \\
\frac{\frac{q}{\alpha}(1-e^{-\alpha t})^2}{\mathbf{D}(t)} & -\frac{1}{(t+1)I_T} + \frac{\frac{q}{\alpha}(1+e^{-2\alpha t})}{\mathbf{D}(t)} & \cdots & \frac{r_i(I_T - r_i^{-1})}{(t+1)I_T} + \frac{\frac{q}{\alpha}(1+e^{-2\alpha t})}{\mathbf{D}(t)} & \cdots \\
\vdots & \vdots & & \vdots & \ddots
\end{bmatrix} \quad (9)
$$

This component is a measure of the information that has been diffused from one landmark to the others.

The most important aspect of the solution is the role of the total Fisher information I_T. This depends on the number and quality of the landmark observations. Its value will be determined by the landmarks selected for observation. The value of I_T directly affects the time constant α for a given velocity measurement variance q and hence strongly affects the convergence rate and final values of the vehicle and landmark position estimate covariances. This therefore raises a question regarding the selection of landmarks for observation. Specifically, it implies that concentrating on a select group of landmarks giving high quality observations might be more effective than the generally accepted approach of observing as many landmarks as possible.

3.1 Convergence

Figure 1 illustrates logarithmically the transient convergence rates of the vehicle variance and vehicle to landmark covariances for $\alpha = 10, 1, 0.1$. The three situations have been generated using the same value of $q = 1$ to permit direct comparison while α has been manipulated via selection of observation information $I_T = 100, 1, 0.01$ respectively. Intuitively, the transient response rate is far more rapid for larger α, leading to smaller steady-state variance. This is completely consistent with discrete simulation.

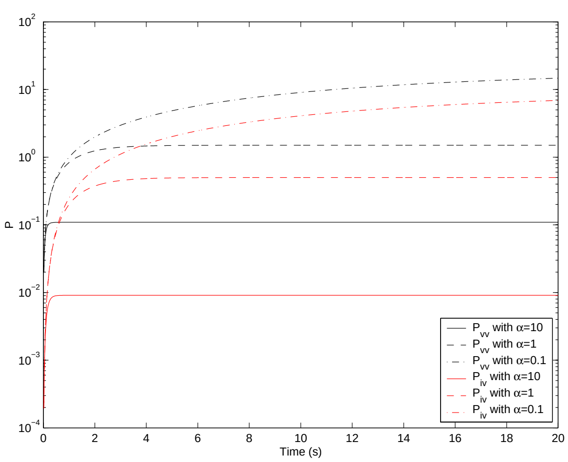

Figure 1: Effect of time constant on vehicle position covariance convergence rates

The convergence behaviour of the landmark covariances and the landmark to landmark covariances is determined by the relative decay rates of the exponential components and the inverse of time associated with the diffusion components. For small α the inverse time decay rate is comparable to the exponential decay rate. The mechanism of the convergence is through the exponential components commencing at a value $\frac{q}{\alpha^2} = \frac{1}{I_T}$ at $t = 0$, and converging toward $\frac{q}{\alpha(\alpha+1)}$. The evolution of the diffusion components is determined only by the initial value of P_{ii} and the total information I_T.

To illustrate this, consider a problem in which 3 landmarks are each observed with variance $r_i = 3$ and where $q = 1$, giving $I_T = 1$ and $\alpha = 1$. In order to look at the effect of α, consider alternative configurations of I_T obtained with $r_i = 300$ and $r_i = 0.03$ giving $\alpha = 0.1$ and $\alpha = 10$ respectively.

The responses of the example are given in Figure 2. Each of the graphs in Figure 2 shows the evolution of the diffusion terms (dotted) for the three landmarks. For a given amount of observation information I_T, they are unaffected by the value of α (i.e. independent of q). For larger α, the exponential components converge more rapidly to larger relative steady-state values (relative to r_i). The convergence of the total landmark covariance toward the steady-state values of the exponential components is obvious.

Similar behaviour occurs for the landmark-to-landmark terms, where the convergence is from zero to $\frac{q}{\alpha(\alpha+1)}$ as determined by the the exponential components. This too is evident in Figure 2.

3.2 Steady-state Behaviour

In the limit as $t \to \infty$, the steady state covariance is given by:

$$
\mathbf{P}(\infty) = \begin{bmatrix}
\frac{q}{\alpha}\frac{(\alpha+2)}{(\alpha+1)} & \frac{q}{\alpha(\alpha+1)} & \cdots & \frac{q}{\alpha(\alpha+1)} \\
\frac{q}{\alpha(\alpha+1)} & \frac{q}{\alpha(\alpha+1)} & \cdots & \frac{q}{\alpha(\alpha+1)} \\
\vdots & \vdots & \ddots & \vdots \\
\frac{q}{\alpha(\alpha+1)} & \frac{q}{\alpha(\alpha+1)} & \cdots & \frac{q}{\alpha(\alpha+1)}
\end{bmatrix} \quad (10)
$$

The steady-state convergence properties are notable in their dependencies. For large α the vehicle variance converges towards $q/\alpha = \sqrt{q/I_T}$, and for small α converges towards $2q/\alpha = 2\sqrt{q/I_T}$. All other terms converge towards $q/\alpha = \sqrt{q/I_T}$ for small α, but for large α converge towards $q/\alpha^2 = 1/I_T$ and therefore the precision of the map is independent of the vehicle uncertainty q.

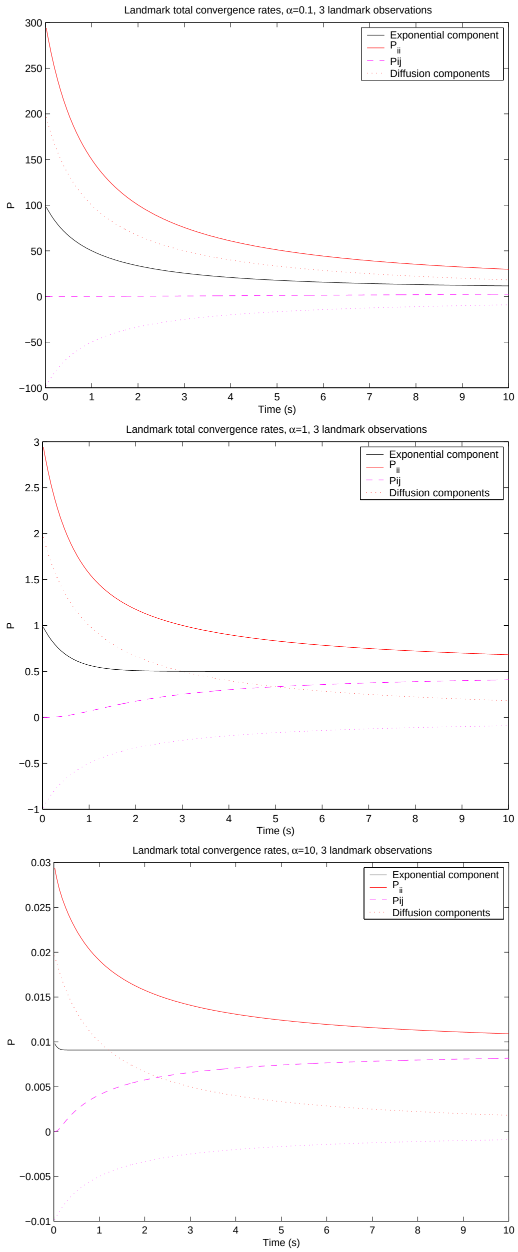

Figure 2: Effect of time constant on landmark covariance convergence rates

Figure 3 summarises the steady-state behaviour for P_{vv}, P_{iv}, P_{ii} and P_{ij}, showing the break value at $\alpha = 1$. For α given by any q and I_T combination, the steady-state precision is defined by the normalised covariances and the specific q.

The effect of total landmark information I_T is now shown by example. Consider a situation in which we

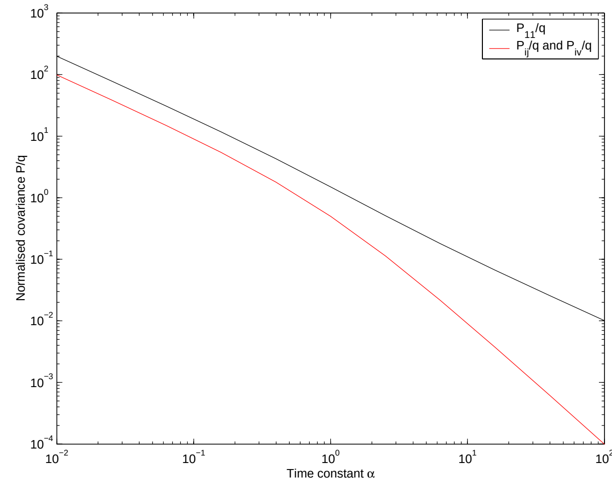

Figure 3: Effect of time constant on the normalised steady-state covariances

have three landmarks observed consecutively with identical variance $r_i = 3$. For a vehicle velocity variance $q = 1$, this gives rise to $\alpha = 1$. Figure 4 illustrates this situation.

The total observation information I_T is increased by increasing the number of landmarks consecutively observed with the same variance. Three additional cases observing 12, 48, and 300 landmarks are given corresponding to $\alpha = 2, 4$ and 10. Total $P_{ii}(t)$ behaviours are shown in Figure 4 together with the exponential component to indicate the steady-state.

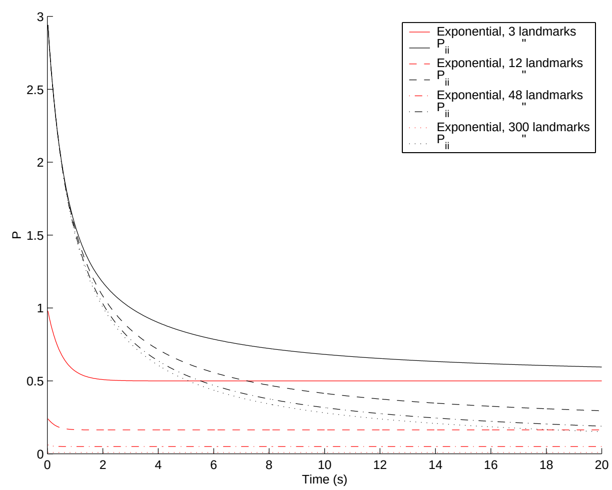

Figure 4: Effect of total information I_T on steady-state landmark covariance

It is clear from Figure 4 that although precision is improved by observing more landmarks, a situation of diminishing returns exists. Little gain is made in going from 48 to 300 landmarks compared to the move from 3 to 12 or 12 to 48. Conversely, note that if observations can be made with smaller r_i, thus increasing I_T and hence α, exponentially less landmarks are required to achieve the same performance.

There is a strong conclusion here, that SLAM performance will be enhanced by concentrating on *fewer, better landmark observations.*

4 Discussion and Conclusions

The closed form solution presented in this paper serves to elucidate the fundamental mechanics of the convergence process of the SLAM problem. In particular, an understanding of expected convergence performance and precision can be obtained a-priori simply by knowing the vehicle velocity sensor and landmark observation sensor measurement characteristics. The solution therefore provides a basis for the design of SLAM sensor systems and for the development of sensor management strategies. In this regard, it implies that SLAM performance will be enhanced by concentrating on *fewer, better landmark observations.* Although it is derived from a restricted case, this solution is a starting position for understanding the behaviour of SLAM in 2 and 3 dimensions where the additional difficulties of nonlinearities and component cross-correlations are present. The extension is via projection of sensor observations and errors into component axes and is the subject of ongoing investigation.

There are two important issues outstanding. The first relates to the convergence behaviour where the initial covariance matrix is non-trivial. Such a situation may arise where the process must be re-initialised from a previous solution, for instance when newly observed landmarks are to be added to the observation set. A closed form solution has been obtained for this situation and will be reported elsewhere.

The second point relates to practical implementations of SLAM in discrete form with sequential observations. It is clear that the convergence behaviour is heavily dependent upon on the total observation information. In the discrete form this arises from the total information accumulated over time. In order to maximise the information, the conclusion to concentrate on a few good landmarks will generalise to a requirement for frequent observations of a few good landmarks in preference to making infrequent observations of many lower quality landmarks. This indicates that an observation optimality condition will become apparent.

References

[1] B.D.O. Anderson. *Optimal Control.* Prentice Hall International, 1991.

[2] J.A. Castellanos, J.M.M. Montiel, J. Neira, and J.D. Tardos. Sensor influence in the performance of simultaneous mobile robot localisation and map building. In *Int. Symposium on Experimental Robotics*, 1999.

[3] J.A. Castellanos, J.D. Tardos, and G. Schmidt. Building a global map of the environment of a mobile robot: the importance of correlations. In *Proc. IEEE Int. Conf. Robotics and Automation*, 1997.

[4] R. Chatila and J.P. Laumond. Position referencing and consistent world modeling for mobile robots. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 138–143, 1985.

[5] M.W.M.G. Dissanayake, P. Newman, and H. Durrant-Whyte. An experimental and theoretical investigation into simultaneous localisation and map building. In *Int. Symposium on Experimental Robotics*, 1999.

[6] J.J. Leonard, R.N. Carpenter, and H.J. Feder. Stochastic mapping using forward looking sonar. In *Field and Service Robotics (FSR)*, 1999.

[7] J.J. Leonard and H.F. Durrant-Whyte. Simultaneous map building and localization for an autonomous mobile robot. In *IEEE Int. Conf. on Intelligent Robot Systems (IROS)*, 1991.

[8] J.J. Leonard and H.F. Durrant-Whyte. *Directed Sonar Navigation.* Kluwer Academic Press, 1992.

[9] J.J. Leonard and H.J. Feder. Experimental analysis of adaptive convurrent mapping and localisation using sonar. In *Int. Symposium on Experimental Robotics*, 1999.

[10] R.C. Smith and P. Cheeseman. On the representation of spatial uncertainty. *Int. J. Robotics Research*, 5(4):56–68, 1987.