Point-mass filter and Cramer-Rao bound for Terrain-Aided Navigation

Niclas Bergman
Department of Electrical Engineering
Linköping University, S-581 83 Linköping, Sweden
www: http://www.control.isy.liu.se
email: niclas@isy.liu.se

LiTH-ISY-R-1987

November 5, 1997

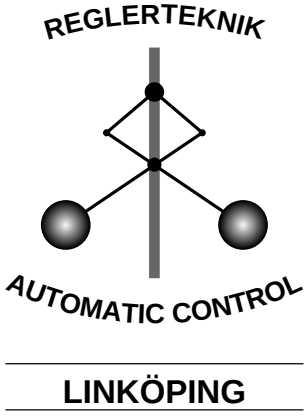

Technical reports from the Automatic Control group in Linköping are available as UNIX-compressed Postscript files by anonymous ftp at the address 130.236.20.24 (ftp.control.isy.liu.se).

Point-mass filter and Cramer-Rao bound for Terrain-Aided Navigation

Niclas Bergman Lennart Ljung
Fredrik Gustafsson
Dept. of Electrical Engineering, Linköping University, Sweden
E-mail: {niclas,ljung,fredrik}@isy.liu.se

Abstract

The nonlinear estimation problem in navigation using terrain height variations is studied. The optimal Bayesian solution to the problem is derived. The implementation is grid based, calculating the probability of a set of points on an adaptively dense mesh. The Cramer-Rao bound is derived. Monte Carlo simulations over a commercial map shows that the algorithm, after convergence, reaches the Cramer-Rao lower bound.

1 Introduction

Modern, high accuracy, navigation systems are configured around an inertial navigation system (INS). Inertial navigation is based on dead-reckoning. From an initial position and velocity estimate, measurements of the aircraft movement are simulated forward in time to continuously yield position and velocity estimates of the aircraft. Due to initial errors and measurement errors, the position estimate of the INS will drift away from the actual position of the aircraft. Terrain-aided navigation (TAN) is the concept of using terrain height variations below the aircraft to render a position estimate that is used to bound the error in the INS. In practice, the integration is often handled by letting an extended Kalman filter estimate the errors in the INS using the terrain related position estimates as measurements. For a general background on aircraft navigation, we refer to [13] and [16].

Consider Figure 1, the idea is to have a digital map, digital terrain elevation database (DTED), on board the aircraft with samples of the terrain elevation, in known, fixed, positions. Flying over an area, the aircraft altitude over mean sea-level is measured with a barometric sensor and the ground clearance is measured with a radar altimeter, pointing downward. A measurement of the terrain elevation is derived from the difference between the altitude and the ground clearance. This derived measurement is compared to the stored values in the DTED and a position estimate is generated. The generation of the estimate must be done with extreme

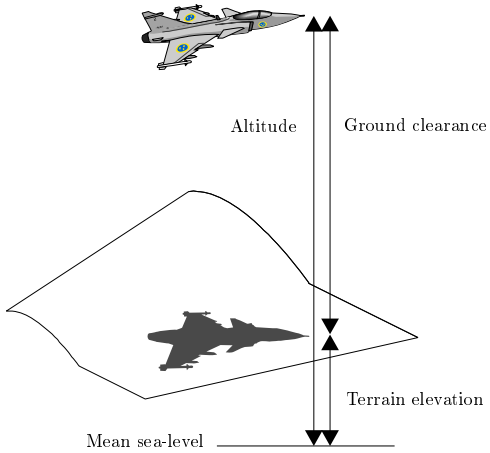

Figure 1: The terrain-aided navigation principle.

care since the terrain might have similar characteristics in several areas. Due to the unstructured and varying characteristic of the terrain, this is a nonlinear estimation problem where the standard nonlinear estimation techniques such as local linearization fail to perform well. Instead of simplifying the model of the problem through linearization this work focuses on solving the complete problem analytically and approximate the implementation of the solution instead of the model. The problem is stated and described in Section 2 and the analytical Bayesian solution is derived in Section 3. The implementation is a point-mass approximation of the posterior, described in Section 4. The Cramer-Rao bound for the estimation problem is derived in Section 5. In Section 6, the algorithm is evaluated using Monte Carlo simulations over a genuine terrain map and a simulator verified against field tests. The result of the simulations is compared with the Cramer-Rao lower bound. Finally, conclusions are drawn in the last section.

2 Problem Description

The estimation problem associated with TAN is to match measurements of terrain elevation with the

DTED. The database consists of samples of photographically generated measurements of terrain elevation in a mesh with 50 m spacing, the value in any point between the stored values is found through interpolation. Assuming that the pressure sensor produces unbiased measurements of the aircraft absolute altitude, the problem is limited to two dimensions, finding the position in the planar geometry of the DTED map.

The performance of the matching of terrain elevation measurements with the DTED depends highly on the type of terrain in the area. Flat terrain gives little, or no, information about the aircraft position. More variations in the terrain give better performance of the matching algorithm. Rough, but repetitive, terrain can give several well matched positions in an area, making it hard to distinguish between several, well matching tracks.

2.1 Estimation model

Let the $(2,1)$-vector stochastic process $\mathbf{x}_t$ denote the position in the DTED, and let $\mathbf{y}_t$ denote the measurement of the terrain elevation. The relation between the measurements and the aircraft position is modeled as,

$$\mathbf{y}_t = h(\mathbf{x}_t) + \mathbf{e}_t. \tag{1}$$

Here, $h(\mathbf{x}_t)$ is the DTED-value below the aircraft, and the additive measurement error, $\mathbf{e}_t$, is assumed white with distribution $p_{\mathbf{e}_t}(x)$. This additive noise models both the errors in the radar altimeter, the current altitude estimate errors and errors due to the DTED map not having perfect resemblance with the actual terrain. In the simulations we will use a Gaussian distribution for the measurement error, but there are some reasons why the Gaussian assumptions might fail, cf. [3] for further details. The reason for the choice of Gaussian distribution is that the Cramer-Rao bound can be derived using this distribution. Let $\mathcal{N}(m, P)$ denote the n-dimensional Gaussian distribution with mean vector m and covariance matrix P,

$$\mathcal{N}(m,P) = \frac{1}{\sqrt{(2\pi)^n\,|P|}}\, e^{-\frac{1}{2}(x-m)^T P^{-1}(x-m)},$$

the distribution of $\mathbf{e}_t$ is

$$p_{\mathbf{e}_t}(x) = \mathcal{N}(0, R).$$

Note that no distinction in notation is used for vectors and scalars, $\mathbf{y}_t$ is scalar while the state vector, $\mathbf{x}_t$, has dimension two. Bold faced notation is used to denote stochastic entities.

The INS gives the relative movement, u_t, between two measurements. The drift in the INS position is modeled by random walk in the state transition equation,

$$\mathbf{x}_{t+1} = \mathbf{x}_t + u_t + \mathbf{v}_t. \tag{2}$$

The system noise is white, has distribution

$$p_{\mathbf{v}_t}(x) = \mathcal{N}(0, Q)$$

and is independent of the measurement noise $\mathbf{e}_t$. The initial state has the distribution

$$p_{\mathbf{x}_0}(x) = \mathcal{N}(\bar{x}_0, P_0).$$

and is independent of both $\mathbf{e}_t$ and $\mathbf{v}_t$.

Equations (1) and (2) describe the TAN estimation model,

$$\begin{aligned} \mathbf{x}_{t+1} &= \mathbf{x}_t + u_t + \mathbf{v}_t \\ \mathbf{y}_t &= h(\mathbf{x}_t) + \mathbf{e}_t \end{aligned} \qquad t = 0, 1, \ldots \tag{3}$$

The terrain matching algorithm should estimate the states of this model using measurements of $\mathbf{y}_t$.

2.2 Terrain matching algorithms

The best known TAN algorithms are TERCOM (terrain contour matching) and SITAN (Sandia inertial terrain-aided navigation). TERCOM is batch oriented and correlates gathered terrain elevation profiles with the map periodically [2, 8, 16]. SITAN uses a modified version of an extended Kalman filter (EKF) in its original formulation [10]. Both these algorithms have been reported to be successful in a number of applications. However, when flying over fairly flat, or over very rough terrain or when the aircraft is highly maneuverable, they do in general not perform well. Several modifications of the SITAN approach have been proposed. In order to overcome divergence problems in the filter estimates parallel EKFs have been used in, e.g., [9, 5]. Generally, these divergence problems originate from the local approximation schemes failing to model the aircraft and terrain accurately. One more recent and different approach that tries to deal with these problems is VATAN [7]. In VATAN the Viterbi algorithm is applied to the TAN problem, yielding a maximum a posteriori position estimate.

3 The Bayesian Solution

The basic problem of estimation is to find out as much as possible about $\mathbf{x}_t$ from observations made of the related set of measurements

$$\mathcal{Y}^t = \{y_i\}_{i=0}^t$$

This problem is often posed as an optimization problem, finding the best estimate using some performance criterion on the deviation from the true state. Using a statistical view of the problem, the probability density function (PDF) for the states gives all information one can ask for regarding the characteristics of the states. That is, the conditional *a posteriori* density function $p_{\mathbf{x}_t|\mathcal{Y}^t}(x)$ summarizes everything there is to know about the states $\mathbf{x}_t$ given the collected measurements. Thus the estimation problem could be posed as the problem of determining the a posteriori density. This is generally referred to as the Bayesian approach [11].

Given the posterior density, one suitable point estimate is the conditional mean,

$$\hat{x}_{t|t} = \int x\, p_{\mathbf{x}_t|\mathcal{Y}^t}(x)\ dx. \tag{4}$$

Note that the posterior density should be unimodal in order to give accurate estimates. It can be shown that the conditional mean is also the estimate that minimizes the mean square error of the estimator [1]. This is why (4) is sometimes referred to as the minimum mean square estimate (MMSE).

The Bayesian solution is based on the following expression for the conditional PDF

$$p_{\mathbf{z}|\mathbf{w}}(z|w) = \frac{p_{\mathbf{z},\mathbf{w}}(z,w)}{p_{\mathbf{w}}(w)}. \tag{5}$$

Restated, the joint PDF for the two stochastic variables can be expressed in there conditional PDF

$$p_{\mathbf{z},\mathbf{w}}(z,w) = p_{\mathbf{z}|\mathbf{w}}(z|w)\, p_{\mathbf{w}}(w)\,. \tag{6}$$

Applying these formulas to (3) a recursion for the posterior is found. The recursion consists of one measurement update and one time update.

Assume that $p_{\mathbf{x}_t|\mathcal{Y}^{t-1}}(x)$ is known, split the measurement set $\mathcal{Y}^t$ in a prior part and a new part,

$$\mathcal{Y}^t = \mathbf{y}_t, \mathcal{Y}^{t-1}.$$

Using (5) and (6), the measurement update is

$$\begin{aligned} p_{\mathbf{x}_t|\mathcal{Y}^t}(x) = p_{\mathbf{x}_t|\mathbf{y}_t,\mathcal{Y}^{t-1}}(x) &= \frac{p_{\mathbf{x}_t,\mathbf{y}_t|\mathcal{Y}^{t-1}}(x,y_t)}{p_{\mathbf{y}_t|\mathcal{Y}^{t-1}}(y_t)} \\ &= \frac{p_{\mathbf{y}_t|\mathbf{x}_t,\mathcal{Y}^{t-1}}(y_t|x)\, p_{\mathbf{x}_t|\mathcal{Y}^{t-1}}(x)}{p_{\mathbf{y}_t|\mathcal{Y}^{t-1}}(y_t)}. \end{aligned} \tag{7}$$

Note that $p_{\mathbf{y}_t|\mathcal{Y}^{t-1}}(y_t)$ is a constant that need not be calculated since the resulting PDF must integrate to unity. Equation (1) says that knowing $\mathbf{x}_t$ the measurement $\mathbf{y}_t$ has the same distribution as $\mathbf{e}_t$, apart from the mean value, i.e., the likelihood is

$$p_{\mathbf{y}_t|\mathbf{x}_t,\mathcal{Y}^{t-1}}(y_t|x) = p_{\mathbf{e}_t}(y_t - h(x))\,,$$

inserted to (7) the measurement update step is completed.

The joint PDF for the states at time t and time $t+1$ can be expressed using (6). Marginalizing on the states $\mathbf{x}_t$, the time update is found to be

$$\begin{aligned} p_{\mathbf{x}_{t+1}|\mathcal{Y}^t}(x) &= \int_{\mathbf{R}^2} p_{\mathbf{x}_{t+1},\mathbf{x}_t|\mathcal{Y}^t}(x,\chi)\ d\chi \\ &= \int_{\mathbf{R}^2} \underbrace{p_{\mathbf{x}_{t+1}|\mathbf{x}_t,\mathcal{Y}^t}(x|\chi)}_{p_{\mathbf{v}_t}(x-\chi-u_t)} p_{\mathbf{x}_t|\mathcal{Y}^t}(\chi)\ d\chi, \end{aligned}$$

where in the last equality (2) says that knowing $\mathbf{x}_t$ the PDF for $\mathbf{x}_{t+1}$ equals the PDF for $\mathbf{v}_t$ with properly adjusted mean value. This completes the time update step.

With no measurement information available, the knowledge about the states is summarized in $p_{\mathbf{x}_0}(x)$ which initializes the recursion. Through induction the following theorem has been proved.

Theorem 1 *The Bayesian recursion for the TAN estimator is*

$$\begin{aligned} p_{\mathbf{x}_t|\mathcal{Y}^t}(x) &= \frac{1}{c} p_{\mathbf{e}_t}(y_t - h(x))\, p_{\mathbf{x}_t|\mathcal{Y}^{t-1}}(x) \\ p_{\mathbf{x}_{t+1}|\mathcal{Y}^t}(x) &= \int_{\mathbf{R}^2} p_{\mathbf{x}_t|\mathcal{Y}^t}(\chi)\, p_{\mathbf{v}_t}(x-\chi-u_t)\ d\chi, \end{aligned}$$

where $c = \int_{\mathbf{R}^2} p_{\mathbf{e}_t}(y_t - h(x))\, p_{\mathbf{x}_t|\mathcal{Y}^{t-1}}(x)\ dx$. *Initializing the recursion with* $p_{\mathbf{x}_0|y_{-1}}(x) = p_{\mathbf{x}_0}(x)$.

Note that, for the case of a linear measurement equation and Gaussian noise the theorem above coincides with the Kalman filter [12, 1].

4 Implementation

Each iteration of the Bayesian solution in Theorem 1 consists of a multiplication, a linear convolution and an integration. Further, to produce an estimate like (4) yet another integral has to be evaluated. Due to the unstructured nonlinearity $h(\cdot)$ these integrations are in general impossible to solve in closed form which makes the problem infinite dimensional. The implementation must therefore be approximate, and the problem has become one of function approximation, finding a finite description for the posterior $p_{\mathbf{x}_t|\mathcal{Y}^t}(x)$.

Applying a uniform grid to the state space is one among several ways to perform this function approximation. The main advantage with quantizing the state space is that the calculations becomes simple, the integrals turn into sums, making it feasible to update many grid points. Several quantization approaches to the nonlinear estimation problem have been proposed in the literature. The earliest reference is [6]. Later references involve the p-vector approach in [17] and a slightly different approach, presented in [15, 14], using a piecewise constant approximation to the density functions.

The grid approximation can be viewed as a bed-of-nails or a point-mass approximation to the PDF, therefore the implementation will be labeled the point-mass filter (PMF).

4.1 The point-mass filter

Since the state dimension is two and we use a uniform grid, the sampled version of the posterior $p_{\mathbf{x}_t|\mathcal{Y}^t}(x)$

will, after truncation of small sample values, be defined by a matrix containing the sampled values, a grid denseness variable, and a reference point for the location of the matrix over the map. The measurement update consists of interpolating the DTED to find the terrain elevation in the nonzero grid points of the point-mass matrix, evaluating the "likelihood matrix" $p_{\mathbf{e}_t}(y_t - h(x))$, and performing an element-wise matrix multiplication. The normalization is equivalent to a summation. The time update convolution becomes a discrete convolution between the measurement updated matrix of point-masses and a sampled version of the PDF for $\mathbf{v}_t$. The point estimate (4) is a weighted summation of the matrix of point-masses, yielding the center of mass of the point-mass approximation.

The measurement update will amplify the grid values at the grid points that coincide with the measurements and attenuate the ones that does not fit well with the gathered terrain elevation profile. The time update convolution will smooth this effect using a two dimensional moving average. The normalization will keep the grid values bounded in each iteration of the algorithm. After some measurements have been processed there will be grid points with very small values, a truncation of values less than ε times the average grid value is introduced right after the measurement update. Note that the point-mass matrix usually will have big holes with zero mass, yielding an effective way to follow several parallel tracks in the map. When the number of nonzero grid points falls below some threshold N_0 the point-mass mesh is interpolated to double denseness, increasing the approximation accuracy. Likewise an upper limit N_1 defines when to decimate the point-mass matrix.

5 The Cramer-Rao lower bound for TAN

The Cramer-Rao bound, P_t, sets a lower limit on the error covariance for an unbiased estimator,

$$C = \mathsf{E}\big[(\mathbf{x}_t - \hat{\mathbf{x}}_{t|t-1})(\mathbf{x}_t - \hat{\mathbf{x}}_{t|t-1})^T\big] \geq P_t. \qquad (8)$$

As it turns out, under the assumption of Gaussian noise, the Cramer-Rao lower bound will coincide with the Kalman filter solution to the model (3) where the nonlinear measurement equation has been replaced with the gradient of $h(\cdot)$, evaluated at the true state value. The result is summarized in the following theorem, see [4] for details.

Theorem 2 *The Cramer-Rao lower bound for the one step prediction of the states in* (3) *satisfies the recursion,*

$$P_{t+1} = P_t - P_t H_t (H_t^T P_t H_t + R_t)^{-1} H_t^T P_t + Q_t,$$

initiated with P_0. H_t *is the gradient of* $h(\cdot)$ *evaluated at the true state value at time* t,

$$H_t = \frac{\partial h(x_t)}{\partial x_t}.$$

The Cramer-Rao bound thus is a function of the noise levels and the true state sequence.

6 Evaluation

We will evaluate the algorithm for Monte Carlo simulations, comparing the mean square error with the Cramer-Rao lower bound. For results on more realistic simulations using the PMF cf. [3].

Performing M Monte Carlo simulations with identical noise characteristics, over a sequence of n measurements the mean square error (MSE) for each t is

$$\mathrm{MSE}_t^M = \frac{1}{M}\sum_{i=1}^{M}(x_t - \hat{x}_{t|t-1}^i)^T(x_t - \hat{x}_{t|t-1}^i),$$

where $\hat{x}_{t|t-1}^i$ is the one step predictor of the state at time t in Monte Carlo run i. The sum of the diagonal elements in the error covariance matrix, C, is the expected squared error, using (8) the following inequality holds in the limit

$$\lim_{M\to\infty} \mathrm{MSE}_t^M \geq \mathrm{tr}\, P_t.$$

Taking the square root on both sides above yields an inequality for the root mean square (RMS) error.

In order to test the algorithm for different terrain characteristics, two separate simulations were performed, one over an area near a lake with very smooth and flat terrain, followed by one simulation over a mountain area with rough terrain containing high peaks and deep valleys. Figures 2 depict parts of the smooth and the rough simulation area respectively. Both these areas are taken from a genuine, commercial DTED and the simulation tracks have been generated in a simulator that has been verified against actual flight test data.

Table 1 summarizes the simulation parameters used in both simulation cases, and Table 2 summarizes the filter parameters. The result of the Monte Carlo simulations are shown in Figures 3 and 4, for the smooth and the rough case respectively. The plots show the Monte Carlo averaged RMS error of the PMF estimates, the root of the mean square error plus one standard deviation of the mean square error, the maximal error obtained in the Monte Carlo runs for each time instant and, finally, the Cramer-Rao lower bound for the simulation track. The results show that the PMF indeed is efficient after convergence of the algorithm, and that

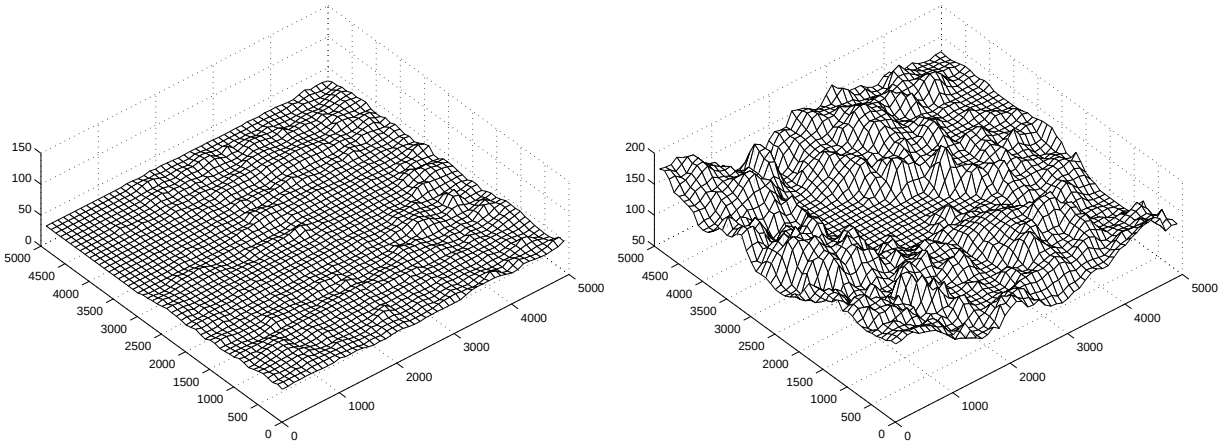

Figure 2: Parts of the smooth and rough simulation areas. Axes are labeled in meters.

Sample interval	$T = 0.103$ sec.
System noise	$Q = \left[\begin{smallmatrix} 4 & 0 \\ 0 & 4 \end{smallmatrix}\right]$
Measurement noise	$R = 16$
Initial covariance	$P_0 = \left[\begin{smallmatrix} 200^2 & 0 \\ 0 & 200^2 \end{smallmatrix}\right]$
Track length	$n = 300$
Monte Carlo runs	$M = 500$

Table 1: Simulation parameters.

Initial grid	50 m
Resampling limits	$N_0 = 1000$, $N_1 = 5000$
Truncation parameter	$\varepsilon = 0.001$

Table 2: Filter parameters.

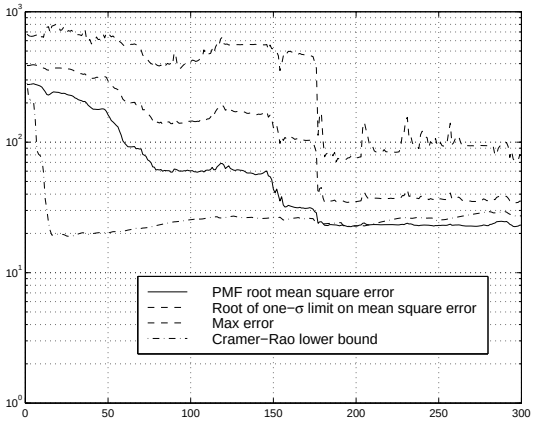

Figure 3: Monte Carlo root mean square error for smooth simulation area.

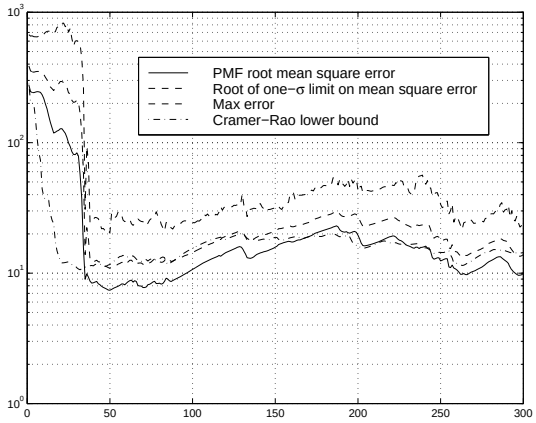

Figure 4: Monte Carlo root mean square error for rough simulation area.

the performance, as expected, depends strongly on the amount of variation in $h(\cdot)$. As seen in Figure 3, during the first half of the simulation track the PMF is not at all near the Cramer-Rao lower bound. But after convergence of the PMF estimate, the RMS error stays very close to the bound. The one-σ limit indicates that there is a great variation between different Monte Carlo runs, this is also seen in the large maximum error.

The rough terrain case in Figure 4 shows much better performance with a convergence rate near that of the Cramer-Rao bound. After convergence the RMS error and the Cramer-Rao bound show equal performance. The one-σ limit is close to the bound which shows that the spread between the different Monte Carlo runs is small. Likewise, the maximum error is much smaller than in the smooth case.

The poor performance during the settling phase is due to the quantization errors being large when using a sparse grid to sample the posterior. Once the grid

denseness has decreased however, the PMF is efficient, meeting the Cramer-Rao lower bound. After convergence the grid denseness is less than four meters.

7 Conclusions

A point-mass implementation of the Bayesian solution to the nonlinear estimation problem in TAN has been evaluated. Monte Carlo simulations, using a flight test verified simulator and commercial terrain database, show nearly optimal performance after convergence of the algorithm as it reaches the Cramer-Rao lower bound. The algorithm performs well after convergence both over varying and flat terrain.

References

[1] B.D.O. Anderson and J.B. Moore. *Optimal Filtering*. Prentice Hall, Englewood Cliffs, NJ., 1979.

[2] W. Baker and R. Clem. Terrain contour matching [TERCOM] primer. Technical Report ASP-TR-77-61, Aeronautical Systems Division, Wright-Patterson AFB, Aug. 1977.

[3] N. Bergman. A Bayesian approach to terrain-aided navigation. In *Proc. of SYSID'97, 11th IFAC Symposium on System Identification*, pages 1531–1536. IFAC, 1997.

[4] N. Bergman. On the Cramer-Rao bound for terrain-aided navigation. Technical Report LiTH-ISY-R-1970, Dept. of EE, Linköpings University, 1997.

[5] Drayton D. Boozer and J. Richard Fellerhoff. Terrain-Aided Navigation Test Results in the AFTI/F-16 Aircraft. *Journal of The Institute of Navigation*, 35(2):161–175, Summer 1988.

[6] R.S. Bucy and K.D. Senne. Digital synthesis of nonlinear filters. *Automatica*, 7:287–298, 1971.

[7] R. Enns and D. Morrell. Terrain-aided navigation using the Viterbi algorithm. *Journal of Guidance, Control and Dynamics*, 18(6):1444–1449, November-December 1995.

[8] J. P. Golden. Terrain contour matching (TERCOM): a cruise missile guidance aid. In T. F. Wiener, editor, *Image Processing for Missile Guidance*, volume 238, pages 10–18. The Society of Photo-Optical Instrumentation Engineers (SPIE), July 1980.

[9] J.A. Hollowell. Heli/SITAN: A terrain referenced navigation algorithm for helicopters. Las Vegas, Mar. 1990. IEEE Pos.,Loc. and Nav. Symp. (PLANS).

[10] L.D. Hostetler. Optimal terrain-aided navigation systems. Palo Alto, CA, Aug. 1978. AIAA Guidance and Control Conference.

[11] A. Jazwinski. *Stochastic Process and Filtering Theory*, volume 64 of *Mathematics in Science and Engineering*. Academic Press, New York, 1970.

[12] R. E. Kalman. A new approach to linear filtering and prediction problems. *Trans. AMSE, J. Basic Engineering*, 82:35–45, 1960.

[13] M. Kayton and W. Fried, editors. *Avionics Navigation Systems*. John Wiley, 2nd edition, 1997.

[14] S.C. Kramer and H.W. Sorenson. Bayesian parameter estimation. *IEEE Transactions on Automatic Control*, 33:217–222, 1988.

[15] S.C. Kramer and H.W. Sorenson. Recursive Bayesian estimation using piece-wise constant approximations. *Automatica*, 24:789–801, 1988.

[16] G. M. Siouris. *Aerospace Avionics Systems*. Academic Press, 1993.

[17] H. W. Sorenson. Recursive estimation for nonlinear dynamic systems. In J. C. Spall, editor, *Bayesian Analysis of Time Series and Dynamic Models*, volume 94 of *Statistics, textbooks and monographs*, chapter 6, pages 127–165. Marcel Dekker inc., New York, 1988.