A Bank of Maximum A Posteriori Estimators for Single-Sensor Range-only Target Tracking

Guoquan P. Huang, Ke X. Zhou, Nikolas Trawny, and Stergios I. Roumeliotis

Abstract— **In this paper, we study estimation consistency of single-sensor target tracking using range-only measurements. We show analytically that the cost function minimized by the iterated extended Kalman filter (IEKF) has up to three local minima, which can potentially result in inconsistency or even divergence. To address this issue, we instead propose a bank of maximum a posteriori (MAP) estimators to determine the target state-space trajectory. In particular, we use the local minima of the IEKF cost function at each time step as highly accurate initial hypotheses to start a bank of iterative nonlinear optimizations. Moreover, we employ pruning and marginalization to control computational complexity. Extensive Monte Carlo simulations show that the proposed algorithm significantly outperforms the IEKF, the unscented Kalman filter (UKF), the bank of IEKFs, the particle filter (PF), and the standard MAP, both in terms of accuracy and convergence speed.**

I. Introduction

Range-only target tracking is a classical nonlinear estimation problem and many algorithms have been developed for its solution, amongst which the extended Kalman filter (EKF) arguably remains one of the most popular estimators [1], [2], primarily due to its low cost and ease of implementation. The EKF linearizes the nonlinear measurement function around the latest state estimate. In order to reduce linearization errors, the iterated EKF (IEKF) iteratively relinearizes the measurement function until its estimate converges to a local minimum [3]. However, as shown in this paper, the IEKF (and thus the EKF) can become inconsistent and even diverge when applied to this problem.

EKF consistency has recently been investigated in the context of simultaneous localization and mapping (SLAM) [4], and cooperative localization [5]. In these cases, the filter violates the observability properties of the underlying nonlinear system and erroneously injects information along directions of the state space that are actually unobservable. However, this is not the case for range-only tracking, since the system is observable (provided appropriate sensor motions) [6].

Instead, the fundamental problem is the fact that the cost function minimized by the IEKF has multiple local minima, of which the IEKF is only able to compute one. In this paper, we *analytically* determine all (up to four) real stationary points among which are shown to be at most *three* local minima. This enables global optimization at each time step. However, a globally optimal estimate per time step does not necessarily result in a globally optimal trajectory over the entire time horizon. A conceivable method to address this issue would be to employ a bank of IEKFs that track all local minima at every time step as possible trajectory hypotheses. This approach unfortunately is still unable to ensure globally optimal estimates over the entire time history. This is due to the fact that any filtering approach inherently marginalizes out the old states prior to the current time step and thus cannot update these obsolete states once new measurements become available for improving their estimates.

An alternative approach to solve nonlinear estimation problems with multiple minima is to use a particle filter (PF) [7]. In the PF, each particle represents a hypothesis of the target state, weighted by its likelihood. If the particles sample the state space sufficiently, the PF will converge to the optimal estimate. However, the particles are usually initialized randomly in the state space. If far from a (local) minimum, their weights will decay quickly and thus lead to particle depletion. Therefore, in order to converge to meaningful estimates, the PF requires a large number of particles and thus has high computational demands.

In contrast, a maximum a posteriori (MAP) batch estimator [8] provides the best estimate of the entire time history of the state. Formulated as a nonlinear least-squares optimization problem and solved iteratively (e.g., using Gauss-Newton [9]), the estimator can attain the globally optimal estimate if well-initialized. However, as it will become evident later on, the standard approach of using the IEKF estimates to initialize the optimization is insufficient and can result in convergence to a local minimum of poor accuracy. Although it is possible in principle to solve the MAP optimization problem globally (by using algebraic geometry techniques [10]) to compute all stationary points, such an approach is in general computationally intractable due to its exponential complexity. Instead, in this paper, we propose to utilize a set of multiple, high-quality initial values to start the iterative MAP optimizations (and thus a *bank* of MAP estimators). In particular, the initial estimates are combinations of all local minima of the filtering problem found analytically at each time step and then used to incrementally solve each of the MAP problems. Since the number of possible combinations could grow exponentially over time, we employ an effective pruning scheme as well as marginalization to control computational cost and thus to meet available computational resources.

This work was supported by the University of Minnesota (DTC), and the National Science Foundation (IIS-0643680, IIS-0811946, IIS-0835637).

G. P. Huang, N. Trawny, and S. I. Roumeliotis are with the Department of Computer Science and Engineering, University of Minnesota, Minneapolis, MN 55455, USA {ghuang|trawny|stergios}@cs.umn.edu

K. X. Zhou is with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455, USA kezhou@cs.umn.edu

II. Problem Formulation

Consider a single sensor moving in a plane and estimating the state (position, velocity, etc.) of a moving target, by processing the available distance measurements. In this work, we study the case of *global* tracking, i.e., the position of the target is expressed with respect to a fixed (global) frame of reference, instead of a relative *sensor-centered* one. Hence, we hereafter assume that the position of the tracking sensor is known with high accuracy in the global frame. The state vector of the target at time-step k, is defined as a vector of dimension $2N$, where $N-1$ is the highest-order time derivative of the target position described by the motion model, and can include components such as position, velocity, acceleration, etc.:

$$\mathbf{x}_k = \begin{bmatrix} x_{T_k} & y_{T_k} & \dot{x}_{T_k} & \dot{y}_{T_k} & \ddot{x}_{T_k} & \ddot{y}_{T_k} & \cdots \end{bmatrix}^T \quad (1)$$

$$= \begin{bmatrix} \mathbf{p}_{T_k}^T & \mathbf{x}_{T_{\mathbf{d},k}}^T \end{bmatrix}^T \quad (2)$$

where $\mathbf{p}_{T_k} \triangleq \begin{bmatrix} x_{T_k} & y_{T_k} \end{bmatrix}^T$ is the target position expressed in the global frame of reference, while $\mathbf{x}_{T_{\mathbf{d},k}} \triangleq \begin{bmatrix} \dot{x}_{T_k} & \dot{y}_{T_k} & \ddot{x}_{T_k} & \ddot{y}_{T_k} & \cdots \end{bmatrix}^T$ denotes all higher-order time derivatives of the target position.

In what follows, we present the motion and measurement models for the range-only target tracking that will be used throughout the paper.

A. Motion Model

We consider the case where the target moves randomly but assume that the stochastic model describing the motion of the target (e.g., constant acceleration or constant velocity [3]) is known. In particular, the discrete-time state propagation equation is generically given by the following linear form:

$$\mathbf{x}_{k+1} = \mathbf{\Phi}_k \mathbf{x}_k + \mathbf{G}_k \mathbf{w}_k \quad (3)$$

where $\mathbf{w}_k$ is the zero-mean white Gaussian noise with covariance $\mathbf{Q}_k$. The state transition matrix, $\mathbf{\Phi}_k$, and the process noise Jacobian, $\mathbf{G}_k$, that appear in the preceding expressions depend on the motion model used [3]. In our work, these can be arbitrary, but known, matrices, since no assumptions on their properties are made.

B. Measurement Model

In this work, we particularly focus on the nonlinear case where a single sensor measures its distance to the target. Specifically, the measurement at time-step k is given by:

$$z_k = ||\mathbf{p}_{T_k} - \mathbf{p}_{S_k}|| + n_k \triangleq h(\mathbf{x}_k) + n_k \quad (4)$$

$$= \sqrt{(x_{T_k} - x_{S_k})^2 + (y_{T_k} - y_{S_k})^2} + n_k \quad (5)$$

where $\mathbf{p}_{S_k} \triangleq [x_{S_k} \;\; y_{S_k}]^T$ is the known sensor position expressed in the global frame of reference, and n_k is white zero-mean Gaussian measurement noise, with variance $\sigma_{d_k}^2$.

For use with an EKF [3], the nonlinear measurement function (4) is linearized around the current estimate in order to update its state and covariance estimate. To reduce linearization errors, the IEKF iteratively relinearizes (4) until convergence. As will become evident in the following section, the state estimate computed by the IEKF at each time step is one of the (possibly multiple) *local minima* of the corresponding cost function (6).

III. EKF Consistency

In what follows, we investigate the main cause of the inconsistency of the IEKF when applied to range-only target tracking. To do so, we first examine the nonlinear cost function minimized by the IEKF at each time step and show that it has multiple (up to three) local minima.

The nonlinear update problem solved by the IEKF at each time-step k for determining the state estimate $\hat{\mathbf{x}}_{k|k}$ can be formulated as the following minimization problem [11]:[1]

$$\min_{\mathbf{x}_k} \left[\frac{1}{2}(\mathbf{x}_k - \hat{\mathbf{x}}_{k|k-1})^T \mathbf{P}_{k|k-1}^{-1} (\mathbf{x}_k - \hat{\mathbf{x}}_{k|k-1}) + \frac{\big(z_k - h(\mathbf{x}_k)\big)^2}{2\sigma_{d_k}^2} \right] \quad (6)$$

It has been shown [11], [12] that the IEKF update corresponds to applying the Gauss-Newton method to determine *one* local minimum. However, the update problem (6) may have *multiple* local minima. By computing only one of the possible local minima (which is often not the global minimum) and using this as the updated state estimate in subsequent propagation and update steps, the IEKF will eventually become inconsistent and diverge. Before presenting a remedy for this problem, in what follows we show how to compute all stationary points of (6) *analytically*.

A. Analytic Determination of the Local Minima

Observing that the distance measurement depends *only* on the target position [see (4)], in [13] we have shown that we can decouple the target position $\mathbf{p}_{T_k}$ and the remaining states $\mathbf{x}_{T_{\mathbf{d},k}}$ when solving (6). Moreover, by introducing a new variable $d = h(\mathbf{p}_{T_k})$, the problem (6) is equivalent to the following constrained minimization problem:[2]

$$\min_{\mathbf{p}_{T_k},\, d} \left[\frac{1}{2}(\mathbf{p}_{T_k} - \hat{\mathbf{p}}_{T_{k|k-1}})^T \mathbf{P}_{\mathbf{p}\mathbf{p}_{k|k-1}}^{-1} (\mathbf{p}_{T_k} - \hat{\mathbf{p}}_{T_{k|k-1}}) + \frac{(z_k - d)^2}{2\sigma_{d_k}^2} \right] \quad (7)$$

$$\text{s.t.} \;\; d^2 = (x_{S_k} - x_{T_k})^2 + (y_{S_k} - y_{T_k})^2 \;,\; d \geq 0 \quad (8)$$

which can be solved by employing the method of Lagrange multipliers [9]. Specifically, without loss of generality, by assuming $\mathbf{P}_{\mathbf{p}\mathbf{p}_{k|k-1}}^{-1} = \mathbf{Diag}(s_1, s_2)$, the Lagrangian function

[1] Throughout this paper the subscript $\ell|j$ refers to the estimate of a quantity at time-step ℓ, after all measurements up to time-step j have been processed. $\hat{x}$ is used to denote the estimate of a random variable x, while $\tilde{x} = x - \hat{x}$ is the error in this estimate. Finally, $\mathbf{0}_{m\times n}$ and $\mathbf{1}_{m\times n}$ denote $m \times n$ matrices of zeros and ones, respectively, while $\mathbf{I}_n$ is the $n \times n$ identity matrix.

[2] Note that the size of the *nonlinear* problem has dramatically decreased, from $2N$, for (6), to a constant size of 2 for minimizing (7).

can be constructed as follows:[3]

$$L(x_{T_k}, y_{T_k}, d, \lambda) = \frac{s_1}{2}(x_{T_k} - \hat{x}_{T_{k|k-1}})^2 + \frac{s_2}{2}(y_{T_k} - \hat{y}_{T_{k|k-1}})^2 + \frac{(z_k - d)^2}{2\sigma_{d_k}^2} + \lambda\left(d^2 - (x_{S_k} - x_{T_k})^2 - (y_{S_k} - y_{T_k})^2\right) \quad (9)$$

where λ is the Lagrangian multiplier. Setting the derivatives of $L(\cdot)$ with respect to the four optimization variables to zero, and performing simple algebraic manipulations, we have:

$$\frac{\partial L}{\partial x_{T_k}} = 0 \Rightarrow x_{T_k} = \frac{s_1 \hat{x}_{T_{k|k-1}} - 2\lambda x_{S_k}}{s_1 - 2\lambda} \quad (10)$$

$$\frac{\partial L}{\partial y_{T_k}} = 0 \Rightarrow y_{T_k} = \frac{s_2 \hat{y}_{T_{k|k-1}} - 2\lambda y_{S_k}}{s_2 - 2\lambda} \quad (11)$$

$$\frac{\partial L}{\partial d} = 0 \Rightarrow d = \frac{z_k}{1 + 2\sigma_{d_k}^2 \lambda} \quad (12)$$

$$\frac{\partial L}{\partial \lambda} = 0 \Rightarrow 0 = d^2 - (x_{S_k} - x_{T_k})^2 - (y_{S_k} - y_{T_k})^2 \quad (13)$$

Substitution of (10)-(12) into (13) and multiplying both sides of (13) with $(1 + 2\sigma_{d_k}^2\lambda)^2(s_1 - 2\lambda)^2(s_2 - 2\lambda)^2$, we obtain a fourth-order univariate polynomial in λ:

$$0 = f(\lambda) = \sum_{i=0}^{4} a_i \lambda^i \quad (14)$$

where a_i, $i = 0, 1, \ldots, 4$, are the coefficients expressed in terms of the known quantities s_1, s_2, z_k, σ_{d_k}, $\hat{x}_{T_{k|k-1}}$, $\hat{y}_{T_{k|k-1}}$, x_{S_k}, and y_{S_k}. Since $f(\lambda)$ is quartic, its roots can be found in closed-form [13].

Although there exist 4 solutions for λ and thus 4 solutions for x_T, y_T and d, as they depend injectively on λ [see (10)-(12)], we only need to consider the pairs (x_T, y_T) that correspond to real solutions for λ and to a nonnegative d [see (8)]. Moreover, since some of these solutions could be local maxima or saddle points, the second derivative test [9] is employed to extract the minima. Finally, once we determine all the local minima for the target position, we can easily compute the corresponding estimates for the higher-order position derivatives [13].

Since the maximum number of local minima for the problem (6) will significantly impact the computational complexity of our proposed algorithm (see Section IV), we seek a tighter upper bound for it. In particular, based on the Finite Dimensional Mountain Pass Theorem (see Theorem 5.2, [14]), we can show the following lemma:[4]

Lemma 3.1: There are at most three local minima for the problem of minimizing (7).

Proof: The details can be found in [13]. ■

B. Limitations of the IEKF

Fig. 1 shows that the IEKF posterior estimate erroneously converges to a local (not global) minimum in a typical example. In cases like this, the IEKF will not be able to compute the best estimate for a particular time step mainly for two reasons: (i) The IEKF can only determine one of the local minima at each time step given the prior estimate and the current measurements; (ii) The linearization points of the IEKF selected at each time step (i.e., the current state estimate) in computing the filter Jacobians cannot be updated later on when more measurements become available, since the corresponding states have been marginalized. Note that even if we use the results of Section III-A to analytically compute and track all current local minima (e.g., by using a bank of IEKFs), the inability of the IEKF to recompute the Jacobians of previous time steps when new information becomes available can significantly degrade the estimation accuracy due to the impact of past linearization errors (see Section V).

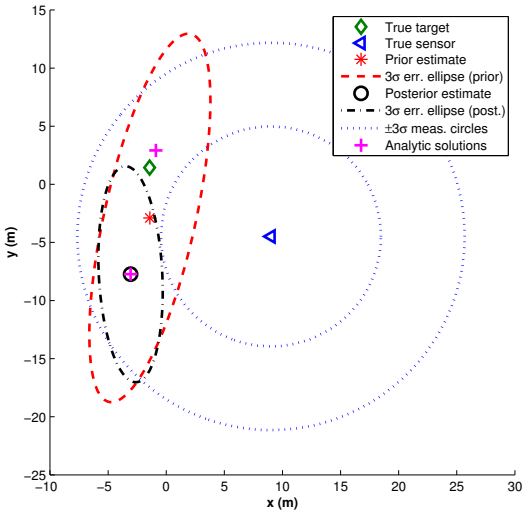

Fig. 1. Illustration of the IEKF posterior estimate converging to the local minimum further from the true state.

IV. Bank of MAP for Range-only Tracking

To address the IEKF limitations, in this section, we start by describing the MAP estimator that considers all measurements when computing the current state estimates. This is the basic component of our proposed estimation approach designed to avoid local minima (see Section IV-B).

A. MAP Estimator

The MAP estimator utilizes all available information to estimate the *entire* history of the target states, described by stacking all states in the time interval $[0, k]$ [see (1)]:

$$\mathbf{x}_K \triangleq \mathbf{x}_{0:k} = \begin{bmatrix} \mathbf{x}_0^T & \mathbf{x}_1^T & \cdots & \mathbf{x}_k^T \end{bmatrix}^T \quad (15)$$

The information used in the MAP estimator includes: (i) the prior information about the initial target state described by a Gaussian pdf with mean $\hat{\mathbf{x}}_{0|0}$ and covariance $\mathbf{P}_{0|0}$, (ii) the target motion-model [see (3)], and (iii) the sensor range measurements [see (4)].

[3] We can always diagonalize $\mathbf{P}_{\mathbf{pp}}^{-1}$ by applying a 2D rotational transformation, which does not affect distance measurements. Moreover, we here temporarily omit the positivity constraint on d, which will be used later for determining feasible solutions.

[4] In practice, we have never observed more than two local minima. Proving this analytically is part of our ongoing work.

The MAP estimator seeks to determine the entire state-space trajectory $\hat{\mathbf{x}}_K$ that maximizes the posterior pdf:

$$p(\mathbf{x}_K|z_{0:k}) \propto p(\mathbf{x}_0)\prod_{\kappa=1}^{k} p(z_\kappa|\mathbf{x}_\kappa)p(\mathbf{x}_\kappa|\mathbf{x}_{\kappa-1}) \tag{16}$$

where $z_{0:k}$ denotes all the sensor measurements in the time interval $[0,k]$. For Gaussian and independent state and measurement noise (see (3) and (4), respectively), maximization of (16) is equivalent to the minimization of the following cost function [13]:

$$c(\mathbf{x}_K) = \frac{1}{2}||\mathbf{x}_0-\hat{\mathbf{x}}_{0|0}||^2_{\mathbf{P}_{0|0}} + \sum_{\kappa=1}^{k}\frac{1}{2}||z_\kappa - h(\mathbf{x}_\kappa)||^2_{\sigma^2_{d_\kappa}} + \sum_{\kappa=1}^{k}\frac{1}{2}||\mathbf{x}_\kappa - \mathbf{\Phi}_{\kappa-1}\mathbf{x}_{\kappa-1}||^2_{\mathbf{Q}'_{\kappa-1}} \tag{17}$$

where we have employed the notations, $||\mathbf{a}||^2_{\mathbf{M}} = \mathbf{a}^T\mathbf{M}^{-1}\mathbf{a}$ and $\mathbf{Q}'_k = \mathbf{G}_k\mathbf{Q}_k\mathbf{G}_k^T$. Since $c(\mathbf{x}_K)$ is a nonlinear function, a standard approach for its optimization is to employ Gauss-Newton iterative minimization [9]. Specifically, at the ℓ-th iteration of this method, a correction, $\delta\mathbf{x}_K^{(\ell)}$, to the current estimate, $\hat{\mathbf{x}}_K^{(\ell)}$, is computed by minimizing the second-order Taylor-series approximation of the cost function which is given by:

$$c(\hat{\mathbf{x}}_K^{(\ell)}+\delta\mathbf{x}_K^{(\ell)}) \simeq c(\hat{\mathbf{x}}_K^{(\ell)}) + \mathbf{b}^{(\ell)^T}\delta\mathbf{x}_K^{(\ell)} + \frac{1}{2}\delta\mathbf{x}_K^{(\ell)^T}\mathbf{A}^{(\ell)}\delta\mathbf{x}_K^{(\ell)} \tag{18}$$

where $\mathbf{b}^{(\ell)} \triangleq \nabla c(\mathbf{x}_K^{(\ell)})$ and $\mathbf{A}^{(\ell)} \triangleq \nabla^2 c(\mathbf{x}_K^{(\ell)})$ are the Jacobian and Hessian of $c(\cdot)$ with respect to $\mathbf{x}_K$, evaluated at $\hat{\mathbf{x}}_K^{(\ell)}$, respectively (see [13]). The value $\delta\mathbf{x}_K^{(\ell)}$ that minimizes (18) is determined by solving the linear system:

$$\mathbf{A}^{(\ell)}\delta\mathbf{x}_K^{(\ell)} = -\mathbf{b}^{(\ell)} \tag{19}$$

Once $\delta\mathbf{x}_K^{(\ell)}$ is found, the new state estimate is computed as:

$$\hat{\mathbf{x}}_K^{(\ell+1)} = \hat{\mathbf{x}}_K^{(\ell)} + \delta\mathbf{x}_K^{(\ell)} \tag{20}$$

Given an initial estimate $\hat{\mathbf{x}}_K^{(0)}$, this iterative algorithm computes the local estimate for the entire target trajectory given all measurements up to time-step k. Finally, we note that this optimization problem needs to be solved at every time step (i.e., once a new measurement becomes available).

1) MAP convergence: The Gauss-Newton algorithm described above converges to the global optimum of the posterior pdf (and hence the MAP estimate), only if the initial estimate $\hat{\mathbf{x}}_K^{(0)}$ is within the region of attraction of the global optimum. This is a limitation of most iterative algorithms used for solving nonlinear optimization problems such as minimization of (17), since, in general, there exists no systematic method for determining an initial estimate $\hat{\mathbf{x}}_K^{(0)}$ that can always ensure convergence to the global optimum. Thus, the MAP estimate for range-only tracking can diverge if no accurate initial estimate is provided, which is confirmed by the simulation results presented in Section V. This is due to the fact that the cost function (17) has *multiple* local minima and the initial estimate for the current time step provided to the MAP estimator (e.g., via the IEKF propagation equations) is generally *not* within the attraction basin of the global minimum.

B. Bank of MAP Estimators

In what follows, we introduce a new bank of MAP estimators algorithm to improve accuracy and convergence speed, by tracking (most) possible trajectories and selecting the minimal cost as the final estimate.

Since, in general, it is intractable to analytically solve the MAP estimation problem for all local minima (i.e., minimization of (17) using the approach presented in Section III-A), we propose to solve it *incrementally*. In particular, we decompose the MAP estimation problem into one-step minimization problems, one for each time step [i.e., the update problem (6)], and then analytically solve them for all local minima corresponding to the current state estimate (see Section III-A). Subsequently, we propagate and refine all these local minima over time.

Specifically, at time-step k, we first propagate the current state estimate corresponding to the i-th solution and its covariance matrix, $\hat{\mathbf{x}}^{[i]}_{k|k}$ and $\mathbf{P}^{[i]}_{k|k}$, $\forall i = 1,2,\ldots,m$ (m is the number of MAP estimators in the bank at time-step k), through the target motion model (3):

$$\hat{\mathbf{x}}^{[i]}_{k+1|k} = \mathbf{\Phi}^{[i]}_k\hat{\mathbf{x}}^{[i]}_{k|k} \tag{21}$$

$$\mathbf{P}^{[i]}_{k+1|k} = \mathbf{\Phi}^{[i]}_k\mathbf{P}^{[i]}_{k|k}{\mathbf{\Phi}^{[i]}_k}^T + \mathbf{G}^{[i]}_k\mathbf{Q}_k{\mathbf{G}^{[i]}_k}^T \tag{22}$$

where the Jacobians $\mathbf{\Phi}^{[i]}_k$ and $\mathbf{G}^{[i]}_k$ are evaluated at the current estimate, $\hat{\mathbf{x}}^{[i]}_{k|k}$. Then, once a new measurement becomes available, the propagated state estimate and covariance are used as the prior in the update problem (6). We subsequently use the algebraic method presented in Section III-A to determine all analytic local minima of (6), denoted by $\check{\mathbf{x}}^{[j]}_{k+1}$, $m \leq j \leq 3m$ (see Lemma 3.1). Finally, for each of these solutions, we employ a MAP estimator that uses the latest estimates of the trajectory corresponding to this solution as the initial value, to *refine* the entire state estimates $\hat{\mathbf{x}}^{[j]}_{K+1}$ up to the current time-step $k+1$ [see (17)].

This procedure recursively evolves over time, and at every time step, generates at most $3m$ trajectory estimates. Thus, in the end, we will have multiple MAP estimates, among which the one with the least cost is reported as the best current estimate of the entire state trajectory. Fig. 2 shows the block diagram of the proposed bank of MAP estimators algorithm for the i-th estimated trajectory, while Algorithm 1 outlines its main steps.

C. Computational Cost Reduction

At this point, it is important to note that, in the worst case when multiple solutions exist for the one-step minimization problem (6) at each time step, the total number of the solutions grows exponentially with time (see Lemma 3.1). This immediately leads to the exponential growth of the number of initial estimates and thus the number of MAP estimators in the bank. Moreover, as the target continuously

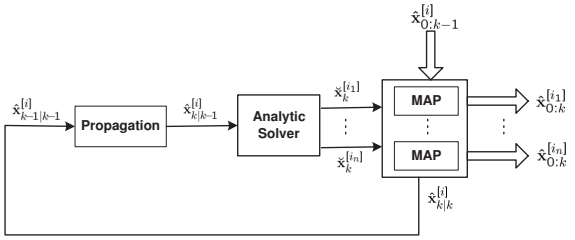

Fig. 2. Block diagram of the i-th MAP estimator in the bank ($1 \leq n \leq 3$).

Algorithm 1 Bank of MAP for Range-only Tracking

For each time-step k:

- Propagate *current* target state estimate and covariance via (21) and (22).
- Analytically determine all local minima of (6).
- For each solution, refine the corresponding *entire* state (trajectory) estimates and covariance, by employing the MAP estimator that uses the latest state estimates corresponding to this solution as the initial guess, and also compute the MAP cost (17).

In the end, select the MAP estimate in the MAP bank with the least cost as the resulting best estimate.

moves, the size of the entire state vector $\mathbf{x}_K^{[i]}$ of each MAP estimator increases linearly with time. In order to make the algorithm suitable for real-time applications, in what follows, we present an effective pruning scheme, as well as the marginalization of old states, so as to reduce the computational cost of the proposed algorithm.

1) Pruning scheme: We observe that in general if two MAP estimators in the bank have similar costs, their state estimates (i.e., the estimated trajectories) are close, too. Moreover, two MAP estimators will converge to the same minimum given that the two initial estimates are in the same basin of attraction. Therefore, we first aggregate the estimated trajectories whose costs are equal within a tolerance, and retain one representative trajectory of each such group while discarding the others. In addition, we also employ the K-means algorithm [15] to cluster the remaining estimated trajectories into two groups (i.e., inlier and outlier) based on their costs and state estimates, and remove the outlier group which have larger costs. These two steps, aggregation and clustering, are repeated,[5] until the number of MAP estimators in the bank is within the threshold, m_{max}.

2) Marginalization of old states: To further reduce the computational complexity, we also employ the marginalization of old states for each MAP estimator in the bank (see [16]). In particular, suppose that at time-step k, all available measurements have been processed by each MAP estimator in the bank. The old states $\mathbf{x_M} \triangleq \mathbf{x}_{0:m}$ are marginalized out, while the states $\mathbf{x_R} \triangleq \mathbf{x}_{m+1:k}$ remain active in the sliding window. Then, as the target keeps moving and the sensor continuously collects new measurements in the time interval $[k, k']$, the state vector of the MAP estimator is augmented by the new target states $\mathbf{x_N} \triangleq \mathbf{x}_{k+1:k'}$. Thus, at time-step k', the sliding window contains the states $\mathbf{x_R}$ and $\mathbf{x_N}$, and the optimal MAP estimate (i.e., without marginalization), is computed by minimizing a cost function similar to (17):

$$c(\mathbf{x}_{K'}) = c(\mathbf{x_M}, \mathbf{x_R}, \mathbf{x_N}) = c_m(\mathbf{x_M}, \mathbf{x_R}) + c_n(\mathbf{x_R}, \mathbf{x_N})$$

where we have decomposed the cost function into two terms: $c_m(\mathbf{x_M}, \mathbf{x_R})$ that contains all quadratic terms that involve states in $\mathbf{x_M}$ only, as well as terms involving the last state in $\mathbf{x_M}$ and the first state in $\mathbf{x_R}$; and $c_n(\mathbf{x_R}, \mathbf{x_N})$ that contains all quadratic terms that involve states in $\mathbf{x_R}$ only, states in $\mathbf{x_N}$ only, and terms involving the last state in $\mathbf{x_R}$ and the first state in $\mathbf{x_N}$. It is important to note that there is no quadratic term jointly involving states in $\mathbf{x_N}$ and $\mathbf{x_M}$, since the target states marginalized at time-step k do not participate in any measurement after that time. Thus, we have:

$$\begin{aligned} &\min_{\mathbf{x_M}, \mathbf{x_R}, \mathbf{x_N}} c(\mathbf{x_M}, \mathbf{x_R}, \mathbf{x_N}) \\ &= \min_{\mathbf{x_R}, \mathbf{x_N}} \left(c_n(\mathbf{x_R}, \mathbf{x_N}) + \min_{\mathbf{x_M}} c_m(\mathbf{x_M}, \mathbf{x_R}) \right) \end{aligned} \tag{23}$$

Next we find $\mathbf{x}_\mathbf{M}^\star = \min_{\mathbf{x_M}} c_m(\mathbf{x_M}, \mathbf{x_R})$ that *only* depends on $\mathbf{x_R}$. This results in an approximately equivalent cost function of (23), $c'_n(\mathbf{x_R}, \mathbf{x_N})$, which does *not* depend on $\mathbf{x_M}$ and whose minimization can be carried out by the Gauss-Newton method (see Section IV-A and [13]). In this process, c_m is permanently approximated by its second-order Taylor series expansion, and the marginalized states $\mathbf{x_M}$, as well as all measurements that directly involve them, can be discarded. Marginalization reduces the computational and memory requirements of each MAP estimator in the bank to a constant depending on the size of the sliding window $[k+1, k']$. This, along with pruning (see Section IV-C.1), results in *constant* computational complexity for the proposed bank of MAP estimators, compared to *linear* for the standard (non-marginalized) MAP estimator (due to the banded structure of the Hessian matrix) (see [13]).

V. Simulation Results

A series of Monte-Carlo comparison studies were conducted under various conditions, in order to validate the preceding analysis and to demonstrate the capability of the proposed bank of MAP estimators algorithm to improve the performance of range-only tracking. In the simulation tests, we considered the scenario where a single moving sensor is tracking a moving target. 50 Monte Carlo simulations were performed, and during each run, all estimators processed the same data, to ensure a fair comparison. The six estimators compared are: (i) the IEKF, (ii) the unscented Kalman filter (UKF) [17], (iii) the bank of IEKFs [tracking all local minima at each time step with the same pruning as in (vi)], (iv) the PF using 1000 particles (see [18]), (v) the standard MAP estimator using the IEKF propagated estimates as the initial guess, and (vi) the proposed bank of MAP estimators

[5]Simulation results have shown that the aggregation is so effective that most of the time there is no need to perform the clustering.

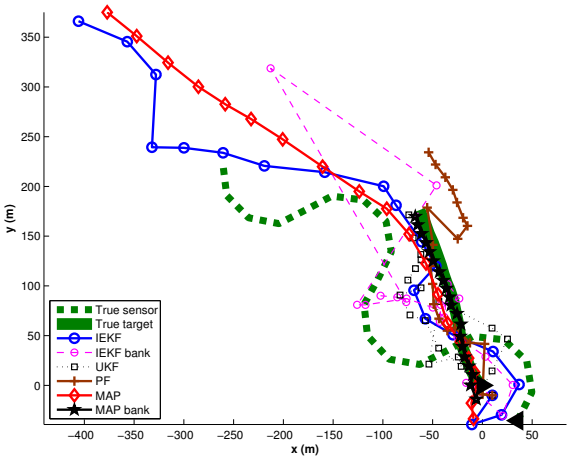

Fig. 3. The estimates of the target trajectory, obtained from one typical realization of the Monte Carlo simulations. In this plot, the dashed line corresponds to true trajectory of the sensor, the solid line to the true trajectory of the target, the solid line with circles to the IEKF, the dashed line with circles to the bank of IEKFs, the dotted line with squares to the UKF, the solid line with crosses to the PF, the solid line with diamonds to the standard MAP which uses the IEKF estimates as initial guess, and the solid line with pentagrams to the proposed bank of MAP estimators. Starting positions of the target and sensor are both marked by ▲. Note that the estimates of the MAP bank is very close to the ground truth, which makes the corresponding lines difficult to distinguish.

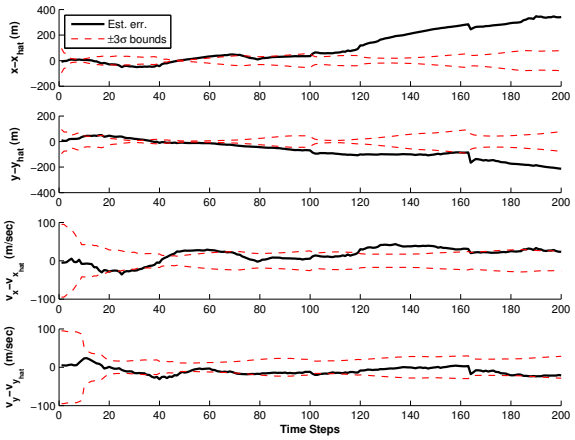

Fig. 4. The IEKF estimation errors vs. 3σ bounds in the typical realization corresponding to Fig. 3. The σ values are computed as the square-root of the corresponding diagonal element of the estimated covariance matrix. Clearly, the estimation errors exceed the 3σ bounds indicating inconsistency.

with pruning ($m_{\max} = 10$) and marginalization (sliding window of 50 time steps).

For the results presented in this section, we adopt a zero-acceleration target motion model [3]:

$$\dot{\mathbf{x}}(t) = \mathbf{F}\mathbf{x}(t) + \mathbf{G}\mathbf{w}(t) \tag{24}$$

where

$$\mathbf{F} = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \mathbf{G} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}, \mathbf{x}(t) = \begin{bmatrix} x_T(t) \\ y_T(t) \\ \dot{x}_T(t) \\ \dot{y}_T(t) \end{bmatrix}$$

and $\mathbf{w}(t) = \begin{bmatrix} w_x(t) & w_y(t) \end{bmatrix}^T$ is zero-mean, white Gaussian noise with covariance $\mathbb{E}\left[\mathbf{w}(t)\mathbf{w}(\tau)^T\right] = q\mathbf{I}_2\delta(t-\tau)$, $q = 5\left(\frac{\text{m}}{\text{sec}}\right)^2\frac{1}{\text{Hz}}$, and $\delta(t-\tau)$ is the Dirac delta function. In our implementation, we discretize the continuous-time system model (24) with time step $\Delta t = 0.1$ sec. The initial true state of the target is $\mathbf{x}_0 = \begin{bmatrix} 0 & 0 & -5 & 5 \end{bmatrix}^T$, while the initial estimate of the target state is set to $\hat{\mathbf{x}}_{0|0} = \begin{bmatrix} 10 & -10 & 0 & 0 \end{bmatrix}^T$, with initial covariance $\mathbf{P}_{0|0} = 1000\mathbf{I}_4$. Finally, the standard deviation of the distance-measurement noise is equal to 10% of the sensor-to-target distance.

Fig. 3 shows the estimates of the target trajectory for the six compared estimators, which are obtained from one typical realization of the 50 Monte Carlo simulations. It is clear that the standard IEKF estimates diverge from the true trajectory and become inconsistent. This can be seen more clearly from Fig. 4 showing that the estimation errors of the IEKF exceed the 3σ bounds. Since the standard MAP uses the IEKF propagated estimates as initial guess (which are often outside the attraction basin of the optimum), the MAP estimate diverges from the true trajectory, too. We also see that the UKF and the bank of IEKFs perform better than the IEKF in this particular simulation. Nevertheless, it is interesting to see that the PF achieves higher accuracy than the MAP estimator in this simulation, though this superior performance is not guaranteed in other runs. This can be justified by the fact that each particle in the PF essentially represents a hypothesis of the target state, and thus the PF is more likely to converge to the optimal solution. However, the increased representational power of the PF (i.e., using more particles), comes at the cost of higher computational requirements. Most importantly, the estimated trajectory of the proposed bank of MAP estimators is the one closest to the ground truth. This is attributed to the good initial estimates attained through the algebraic method (see Section III-A).

Fig. 5 shows the average root mean squared (RMS) errors of the six estimators, which are obtained by averaging over all Monte Carlo runs for each time step. As evident, the bank of MAP estimators performs substantially better than its competitors. We also note that the standard MAP estimates are as poor as those of the IEKF. One possible explanation for this is that the local minimum the MAP converges to, is not necessarily more accurate than the initial estimates (i.e., the IEKF estimates), and thus a poor initial guess could lead to the divergence of the MAP estimates.

Finally, we compared the convergence rate of the two MAP estimators, the standard MAP and the proposed MAP bank. In particular, we counted and compared the number of diverged runs as well as the number of Gauss-Newton iterations used in each realization. Note that in this simulation, the maximum number of Gauss-Newton iterations allowed is set to 100 (i.e., if the MAP iterates for more than 100 times, it is considered to be divergent). As evident from Fig. 6, the bank of MAP estimators converges in significantly more cases than the standard MAP estimator, and moreover, it converges substantially faster.

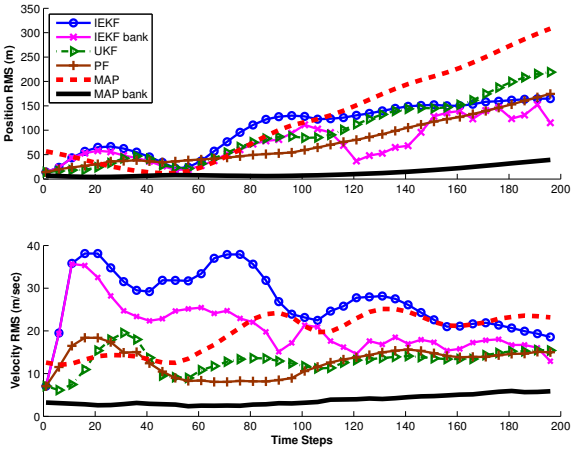

Fig. 5. Averaged RMS errors over 50 Monte Carlo runs.

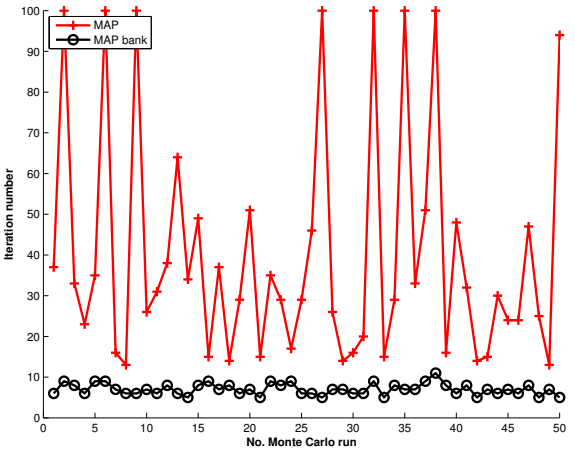

Fig. 6. Number of Gauss-Newton iterations required in each Monte Carlo run. It is clear that the MAP bank converges significantly faster than the standard MAP. Note that in this simulation, the maximum number of Gauss-Newton iterations allowed is set to 100.

VI. Conclusions and Future Work

In this paper, we have shown that the (iterated) EKF can become inconsistent and even diverge when applied to the problem of range-only target tracking. This is due to two reasons: (i) the EKF (and its variants) is only able to track a single mode of the generally multimodal posterior pdf of the target state, and (ii) the EKF marginalizes out the old states prior to the current time step and thus cannot update these states once new measurements become available for improving their estimates. Therefore, we instead employ the optimal (up to linearization errors) MAP estimator. However, the convergence of any iterative optimization algorithm employed to solve the MAP estimation problem is guaranteed only when the initial estimate lies within its basin of attraction. To improve the MAP convergence speed and accuracy, we have introduced a new bank of MAP estimators, which tracks most possible trajectories, by analytically determining all locally-optimal estimates of the current target state at every time step and subsequently employing Gauss-Newton iterations to compute the entire (locally) optimal trajectories corresponding to each candidate local minimum. In addition, we have employed pruning and marginalization for each MAP estimator in the bank to control the computational complexity. Our extensive simulation results have shown that the proposed algorithm significantly outperforms the IEKF, the UKF, the bank of IEKFs, the PF, and the standard MAP, in terms of both accuracy and convergence speed. In this paper, we have concentrated on the case in which only distance measurements are used. However, our approach is applicable to other sensing modalities (e.g., bearing-only, or range and bearing). This is the subject of our future work.

References

[1] K. Zhou and S. I. Roumeliotis, "Optimal motion strategies for range-only constrained multisensor target tracking," *IEEE Trans. on Robotics*, vol. 24, no. 5, pp. 1168–1185, Oct. 2008.

[2] K. X. Zhou and S. I. Roumeliotis, "Multi-robot active target tracking with distance and bearing observabtions," in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, St. Louis, MO, Oct. 11–15 2009, pp. 2209–2216.

[3] Y. Bar-Shalom, X. Li, and T. Kirubarajan, *Estimation with applications to tracking and navigation*. John Wiley & Sons, Inc., 2001.

[4] G. P. Huang, A. I. Mourikis, and S. I. Roumeliotis, "Analysis and improvement of the consistency of extended Kalman filter-based SLAM," in *Proc. IEEE International Conference on Robotics and Automation*, Pasadena, CA, May 19–23 2008, pp. 473–479.

[5] G. P. Huang, N. Trawny, A. I. Mourikis, and S. I. Roumeliotis, "On the consistency of multi-robot cooperative localization," in *Proc. Robotics: Science and Systems*, Seattle, WA, Jun. 28–Jul. 1 2009, pp. 65–72.

[6] T. L. Song, "Observability of target tracking with range-only measurements," *IEEE Journal of Oceanic Engineering*, vol. 24, no. 3, pp. 383–387, Jul. 1999.

[7] F. Gunnarsson, N. Bergman, U. Forssell, J. Jansson, R. Karlsson, and P.-J. Nordlund, "Particle filters for positioning, navigation and tracking," *IEEE Trans. on Signal Processing*, vol. 50, no. 2, pp. 425–437, Feb. 2002.

[8] S. Kay, *Fundamentals of Statistical Signal Processing, Vol. I - Estimation Theory*. Prentice Hall, 1993.

[9] D. P. Bertsekas, *Nonlinear Programming*. Athena Scientific, 1999.

[10] D. Cox, J. Little, and D. O'Shea, *Using Algebraic Geometry*. Springer, 2005.

[11] B. M. Bell and F. W. Cathey, "The iterated Kalman filter update as a Gauss-Newton method," *IEEE Trans. on Automatic Control*, vol. 38, no. 2, pp. 294–297, Feb. 1993.

[12] A. H. Jazwinski, *Stochastic Processes and Filtering Theory*. Academic Press, 1970.

[13] G. P. Huang, K. X. Zhou, N. Trawny, and S. I. Roumeliotis, "A bank of maximum a posteriori estimation algorithm for single-sensor target tracking with range-only measurements," MARS Lab, University of Minnesota, Minneapolis, MN, Tech. Rep., Sep. 2009. [Online]. Available: www.cs.umn.edu/∼ghuang/paper/TR_ROT.pdf

[14] Y. Jabri, *The Mountain Pass Theorem: Variants, Generalizations and Some Applications*. Cambridge University Press, 2003.

[15] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*. Wiley-Interscience, 2000.

[16] E. D. Nerurkar, S. I. Roumeliotis, and A. Martinelli, "Distributed maximum a posteriori estimation for multi-robot cooperative localization," in *Proc. IEEE International Conference on Robotics and Automation*, Kobe, Japan, May 12–17 2009, pp. 1402–1409.

[17] S. Julier, J. Uhlmann, and H. F. Durrant-Whyte, "A new method for the nonlinear transformation of means and covariances in filters and estimators," *IEEE Trans. on Automatic Control*, vol. 45, no. 3, pp. 477–482, Mar. 2000.

[18] V. Cevher, R. Velmurugan, and J. H. McClellan, "A range-only multiple target particle filter tracker," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, Toulouse, France, May 14–19 2006.