ANALYTICALLY-SELECTED MULTI-HYPOTHESIS INCREMENTAL MAP ESTIMATION

Guoquan Huang, Michael Kaess, and John J. Leonard

Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology, Cambridge, MA 02139

Stergios I. Roumeliotis

Department of Computer Science and Engineering
University of Minnesota, Minneapolis, MN 55455

ABSTRACT

In this paper, we introduce an efficient maximum a posteriori (MAP) estimation algorithm, which effectively tracks multiple most probable hypotheses. In particular, due to multimodal distributions arising in most nonlinear problems, we employ a bank of MAP to track these modes (hypotheses). The key idea is that we *analytically* determine all the posterior modes for the current state at each time step, which are used to generate highly probable hypotheses for the entire trajectory. Moreover, since it is expensive to solve the MAP problem sequentially over time by an iterative method such as Gauss-Newton, in order to speed up its solution, we reuse the previous computations and *incrementally* update the square-root information matrix at every time step, while batch relinearization is performed only periodically or as needed.

Index Terms— Maximum a posteriori (MAP), QR factorization, analytical solution, multi-hypothesis tracking

1. INTRODUCTION

Nonlinear estimation problems, such as target tracking, are often addressed using linearized filters, e.g., extended Kalman filter (EKF) [1, 2]. The performance of these filters can be significantly degraded by large linearization errors as well as the inability to track multimodal distributions (which arise in most nonlinear estimation problems). To reduce linearization errors, the iterated EKF (IEKF) [2] is often used, which iterates the filter update till convergence by relinearizing the measurement function at each iteration. Alternatively, the unscented Kalman filter (UKF) [3] deterministically samples the nonlinear models around the current state estimate (i.e., statistical linearization), thus improving the linear approximation. However, any linearization-based filtering approach marginalizes all but the current state, and hence is unable to refine past linearization points.

In contrast, a batch maximum a posteriori (MAP) estimator [4] computes the estimates for the states at all time steps using all available measurements. This allows batch relinearization (i.e., recomputing the Jacobian) of the entire state trajectory, thus greatly reducing the linearization errors. However, the computational cost of the batch-MAP estimator grows unboundedly as the size of its state vector increases linearly over time. Moreover, as the EKF and its variants, the batch-MAP estimator can only track one of the potentially many modes of the posterior pdf. Even though multi-hypothesis tracking algorithms, such as the multi-hypothesis EKF (MHEKF) [5], are available to track a set of different hypotheses, in most cases the hypotheses are generated randomly, thus wasting considerable computational resources.

To address the aforementioned issues, in this paper we introduce a bank of incremental MAP (B-iMAP) estimation algorithm, which provides effective multi-hypothesis tracking and efficient incremental solutions. In particular, we propose an efficient analytical hypothesis generation scheme for MAP-based multi-hypothesis tracking. By transforming the nonlinear cost function of the one-step MAP problem for the current state into polynomial form, and employing algebraic geometry techniques [6], we determine all the posterior modes *analytically*. Each mode is used to initialize a new MAP in the bank, thus allowing to track *multiple* most probable hypotheses of the trajectory. Note that in our prior work [7] we successfully applied this method of analytically selecting hypotheses to the particular problem of range-only target tracking. Furthermore, when using the Gauss-Newton method to solve the batch-MAP problem, in general, it is too expensive to conduct the batch relinearization at every time step. To speed up the batch solver, we reuse the previously-computed Jacobians and *only* compute the one corresponding to the new available measurements, while relinearization is performed periodically or as needed in order to reduce the linearization errors. Although similar idea was used in our recent work of robot localization and mapping [8], in this paper we extend this methodology to a large class of nonlinear systems and integrate it within the analytically-selected multi-hypothesis tracking.

2. THE PROPOSED ALGORITHM

Consider the following general nonlinear system:

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

$$\mathbf{z}_k = \mathbf{h}(\mathbf{x}_k) + \mathbf{v}_k \quad (2)$$

This work was partially supported by ONR grants N00014-10-1-0936, N00014-11-1-0688, and N00014-12-10020; and was also supported by the National Science Foundation (IIS-0643680).

where $\mathbf{x}_k$ denotes the state of the system, $\mathbf{u}_k$ is the control input, and $\mathbf{w}_k$ is zero-mean white Gaussian process noise, i.e., $\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{W}_k)$; $\mathbf{z}_k$ is the measurement, corrupted by measurement noise, $\mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{V}_k)$. We aim to estimate the entire state trajectory, $\mathbf{x}_{0:k}^T = \begin{bmatrix}\mathbf{x}_0^T & \mathbf{x}_1^T & \cdots & \mathbf{x}_k^T\end{bmatrix}$, using all available information. To this end, the batch-MAP estimator is often employed to determine the entire trajectory estimate $\hat{\mathbf{x}}_{0:k|k}$ that maximizes the following posterior pdf:[1]

$$p(\mathbf{x}_{0:k}|\mathbf{z}_{0:k}) \propto p(\mathbf{x}_0)\prod_{\kappa=1}^{k} p(\mathbf{x}_\kappa|\mathbf{x}_{\kappa-1})p(\mathbf{z}_\kappa|\mathbf{x}_\kappa) \quad (3)$$

where $p(\mathbf{x}_0) = \mathcal{N}(\hat{\mathbf{x}}_{0|0}, \mathbf{P}_{0|0})$ is the prior distribution, and $\mathbf{z}_{1:k}$ denotes all the sensor measurements in the time interval $[1, k]$. In (3), we have used the fact of independent state and measurement noise and Markov system dynamics. Using the assumption of Gaussian noise, the maximization of (3) is equivalent to the minimization of the following cost function:

$$c(\mathbf{x}_{0:k}) = \frac{1}{2}||\mathbf{x}_0 - \hat{\mathbf{x}}_{0|0}||^2_{\mathbf{P}_{0|0}} + \quad (4)$$
$$\frac{1}{2}\sum_{\kappa=0}^{k-1}||\mathbf{x}_{\kappa+1} - \mathbf{f}(\mathbf{x}_\kappa, \mathbf{u}_\kappa)||^2_{\mathbf{W}_\kappa} + \frac{1}{2}\sum_{\kappa=1}^{k}||\mathbf{z}_\kappa - \mathbf{h}(\mathbf{x}_\kappa)||^2_{\mathbf{V}_\kappa}$$

where we have employed the notation, $||\mathbf{a}||^2_{\mathbf{\Lambda}} \triangleq \mathbf{a}^T\mathbf{\Lambda}^{-1}\mathbf{a}$. This clearly is a *nonlinear* least-squares problem [see (1) and (2)]. A standard iterative Gauss-Newton approach is often used for its optimization, which, however, is only able to converge to one local minimum within the basin of attraction of the initial estimate. Hence, this approach heavily depends on the quality of the initial estimate. In particular, at the ℓ-th Gauss-Newton iteration, a correction, $\delta\mathbf{x}_{0:k}^{(\ell)}$, to the current estimate, $\hat{\mathbf{x}}_{0:k|k}^{(\ell)}$, is computed by minimizing the second-order Taylor-series approximation of (4), which can be written as:

$$c(\hat{\mathbf{x}}_{0:k|k}^{(\ell)} + \delta\mathbf{x}_{0:k}) \simeq \frac{1}{2}||\delta\mathbf{x}_0||^2_{\mathbf{P}_{0|0}} + \quad (5)$$
$$\frac{1}{2}\sum_{\kappa=0}^{k-1}||\hat{\mathbf{x}}_{\kappa+1|k}^{(\ell)} - \mathbf{f}(\hat{\mathbf{x}}_{\kappa|k}^{(\ell)}, \mathbf{u}_\kappa) + \delta\mathbf{x}_{\kappa+1} - \mathbf{\Phi}_\kappa\delta\mathbf{x}_\kappa||^2_{\mathbf{W}_\kappa} +$$
$$\frac{1}{2}\sum_{\kappa=1}^{k}||\mathbf{z}_\kappa - \mathbf{h}(\hat{\mathbf{x}}_{\kappa|k}^{(\ell)}) - \mathbf{H}_\kappa\delta\mathbf{x}_\kappa||^2_{\mathbf{V}_\kappa} =: ||\mathbf{A}\delta\mathbf{x}_{0:k} - \mathbf{b}||^2$$

where we linearize the process and measurement models (1)-(2) at the current state estimate, with the system and measurement Jacabians, $\mathbf{\Phi}_\kappa = \frac{\partial\mathbf{f}}{\partial\mathbf{x}_\kappa}\Big|_{\mathbf{x}_\kappa=\hat{\mathbf{x}}_{\kappa|k}^{(\ell)}}$ and $\mathbf{H}_\kappa = \frac{\partial\mathbf{h}}{\partial\mathbf{x}_\kappa}\Big|_{\mathbf{x}_\kappa=\hat{\mathbf{x}}_{\kappa|k}^{(\ell)}}$, respectively. We now have a *linear* least-squares problem w.r.t. $\delta\mathbf{x}_{0:k}$ (5), where $\mathbf{A}$ is the Jacobian matrix obtained by appropriately stacking and weighting all the system and measurement Jacobians, and $\mathbf{b}$ is the corresponding stacked residual vector. We employ QR-factorization [9] to solve the problem of minimizing (5), i.e.,

$$\min_{\delta\mathbf{x}_{0:k}} ||\mathbf{A}\delta\mathbf{x}_{0:k} - \mathbf{b}||^2 = \left|\left|\mathbf{Q}\begin{bmatrix}\mathbf{R}\\ \mathbf{0}\end{bmatrix}\delta\mathbf{x}_{0:k} - \mathbf{b}\right|\right|^2 =$$
$$\left|\left|\begin{bmatrix}\mathbf{R}\\ \mathbf{0}\end{bmatrix}\delta\mathbf{x}_{0:k} - \mathbf{Q}^T\mathbf{b}\right|\right|^2 =: \left|\left|\begin{bmatrix}\mathbf{R}\\ \mathbf{0}\end{bmatrix}\delta\mathbf{x}_{0:k} - \begin{bmatrix}\mathbf{d}\\ \mathbf{e}\end{bmatrix}\right|\right|^2$$
$$\Leftrightarrow \min_{\delta\mathbf{x}_{0:k}} ||\mathbf{R}\delta\mathbf{x}_{0:k} - \mathbf{d}||^2 \quad (6)$$

where we have used the economic QR of $\mathbf{A}$ [9], since it in general is tall, i.e.,

$$\mathbf{A} = \mathbf{Q}\begin{bmatrix}\mathbf{R}\\ \mathbf{0}\end{bmatrix} = \begin{bmatrix}\mathbf{Q}_1 & \mathbf{Q}_2\end{bmatrix}\begin{bmatrix}\mathbf{R}\\ \mathbf{0}\end{bmatrix} = \mathbf{Q}_1\mathbf{R} \quad (7)$$

Once $\delta\mathbf{x}_{0:k}^{(\ell)}$ is found by back substitution (6), the new state estimate is updated as: $\hat{\mathbf{x}}_{0:k|k}^{(\ell+1)} = \hat{\mathbf{x}}_{0:k|k}^{(\ell)} + \delta\mathbf{x}_{0:k}^{(\ell)}$.

2.1. Analytical hypotheses

Since in general it is computationally intractable to solve the batch-MAP problem analytically [see (4)], we use the Gauss-Newton iterative method to solve it at each time step. However, as mentioned before, any iterative approach only converges to one local minimum, despite the fact that we often have multimodal posterior pdfs. To address this issue, we provide an effective way to generate multiple, most probable, hypotheses, each of which is tracked by a MAP estimator. Toward this end, we relax the problem of minimizing (4) and incrementally solve a one-step MAP problem analytically. At time-step k, by fixing the past state estimates $\hat{\mathbf{x}}_{0:k-1|k-1}$ (i.e., assuming they are optimal), we approximate (4) as follows:

$$c(\mathbf{x}_{0:k}) \simeq c(\hat{\mathbf{x}}_{0:k-1|k-1}) + \frac{1}{2}||\mathbf{x}_k - \hat{\mathbf{x}}_{k|k-1}||^2_{\mathbf{P}_{k|k-1}}$$
$$+ \frac{1}{2}||\mathbf{z}_k - \mathbf{h}(\mathbf{x}_k)||^2_{\mathbf{V}_k} \quad (8)$$

where $\mathcal{N}(\hat{\mathbf{x}}_{k|k-1}, \mathbf{P}_{k|k-1})$ is the Gaussian prior pdf for the current new state $\mathbf{x}_k$, and is computed by:

$$\hat{\mathbf{x}}_{k|k-1} = \mathbf{f}(\hat{\mathbf{x}}_{k-1|k-1}, \mathbf{u}_{k-1}) \quad (9)$$
$$\mathbf{P}_{k|k-1} = \mathbf{\Phi}_{k-1}\mathbf{P}_{k-1|k-1}\mathbf{\Phi}_{k-1}^T + \mathbf{W}_{k-1} \quad (10)$$

Now the one-step MAP problem of minimizing (8) becomes equivalent to solving the following one-step minimization problem only for the new state estimate:

$$\min_{\mathbf{x}_k} \frac{1}{2}||\mathbf{x}_k - \hat{\mathbf{x}}_{k|k-1}||^2_{\mathbf{P}_{k|k-1}} + \frac{1}{2}||\mathbf{z}_k - \mathbf{h}(\mathbf{x}_k)||^2_{\mathbf{V}_k} \quad (11)$$

In many cases (e.g., see Section 3), the optimality condition of (11) can be transformed into a polynomial system of equations, which can be solved for all the local minima (corresponding to all the modes of the posterior pdf) *analytically* using algebraic geometry techniques [6]. Once all the local minima are found, we use them along with the past state estimates as accurate initial estimates for the MAPs in the bank.

[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. $\hat{x}$ denotes the estimate of a random variable x.

2.2. Incremental QR update

When a new measurement, $\mathbf{z}_{k+1}$, becomes available,[2] we need to recompute the whole Jacobian $\mathbf{A}$, which is an expensive operation in solving the batch-MAP problem. To save computations, we reuse the previously-computed Jacobian and *incrementally* update the QR factorization directly. In particular, we augment $\mathbf{A}$ (without recomputing it) with the new weighted measurement Jacobian $\mathbf{H}_{k+1}$ [see (7)]:

$$\bar{\mathbf{A}} := \begin{bmatrix} \mathbf{A} \\ \mathbf{V}_{k+1}^{-\frac{1}{2}}\mathbf{H}_{k+1} \end{bmatrix} = \begin{bmatrix} \mathbf{Q}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{R} \\ \mathbf{V}_{k+1}^{-\frac{1}{2}}\mathbf{H}_{k+1} \end{bmatrix} \quad (12)$$

We now aim to decompose $\bar{\mathbf{A}}$ into triangular form (i.e., square-root information matrix). Since $\mathbf{A}$ was already factorized into the triangular $\mathbf{R}$, we only need to zero out the new block row of the weighted measurement Jacobian, $\mathbf{V}_{k+1}^{-\frac{1}{2}}\mathbf{H}_{k+1}$, in order to obtain the updated square-root information matrix $\bar{\mathbf{R}}$. This can be achieved efficiently, for example, by Givens or Householder QR [9]. Similarly, the corresponding new vector, $\bar{\mathbf{d}}$, can be obtained by applying the same Givens rotations or Householder reflections to the augmented residual vector, $\begin{bmatrix} \mathbf{d} \\ \mathbf{V}_{k+1}^{-\frac{1}{2}}\left(\mathbf{z}_{k+1} - \mathbf{h}(\hat{\mathbf{x}}_{k+1|k})\right) \end{bmatrix}$. It is important to note that, although relinearization is not needed at each time step when a new measurement becomes available, in order to reduce the linearization errors, we relinearize the system at the latest, and thus the best, state estimates periodically [8] or as needed (e.g., when the linearization point significantly deviates from the current state estimate).

3. EXAMPLE: RANGE-ONLY TRACKING

In this section, we apply the proposed B-iMAP presented in the preceding section to the particular problem of range-only target tracking. Consider a single sensor moving in a plane and estimating the state (position, velocity, etc.) of a moving target, by processing the available range measurements. The target state at time-step k is defined as a vector of dimension $2N$, where $N-1$ is the highest order of the time derivative of the target position described by a known stochastic motion model. It can include components such as position, velocity, and acceleration, i.e., $\mathbf{x}_k^T = \begin{bmatrix} x_{T_k} & y_{T_k} & \dot{x}_{T_k} & \dot{y}_{T_k} & \ddot{x}_{T_k} & \ddot{y}_{T_k} & \cdots \end{bmatrix}$, where $\mathbf{p}_{T_k} := \begin{bmatrix} x_{T_k} & y_{T_k} \end{bmatrix}^T$ is the target position, and $\mathbf{d}_{T_k} := \begin{bmatrix} \dot{x}_{T_k} & \dot{y}_{T_k} & \ddot{x}_{T_k} & \ddot{y}_{T_k} & \cdots \end{bmatrix}^T$ denotes all the higher-order time derivatives of the target position.

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 [2]) is known. In particular, the discrete-time state propagation equation is given by the following linear form:

$$\mathbf{x}_k = \mathbf{\Phi}_{k-1}\mathbf{x}_{k-1} + \mathbf{G}_{k-1}\mathbf{w}_{k-1} \quad (13)$$

where $\mathbf{w}_{k-1}$ is zero-mean white Gaussian noise with covariance $\mathbf{W}_{k-1}$; and the state transition matrix, $\mathbf{\Phi}_{k-1}$, and the process noise Jacobian, $\mathbf{G}_{k-1}$, depend on the motion model used [2]. We will make no further assumptions on these matrices other than that their values are known. In the case where a single sensor measures its distance to the target, the range-only measurement equation at time-step k is given by:

$$z_k = \sqrt{(x_{T_k} - x_{S_k})^2 + (y_{T_k} - y_{S_k})^2} + v_k \quad (14)$$

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

3.1. Analytically selecting hypotheses (posterior modes)

We now show how we solve (11) analytically in the case of range-only tracking, whose solutions (posterior modes) are used to generate the most probable hypotheses for the entire trajectory. By observing that the range measurement depends only on the target position, as shown in [7], we can decouple the target position $\mathbf{p}_{T_k}$ and the remaining states $\mathbf{d}_{T_k}$ in solving (11). Hence, we hereafter focus on the following minimization w.r.t. the target position, while the remaining states can be directly inferred from the optimal solution of (15).

$$\min_{\mathbf{p}_{T_k}} \frac{1}{2}||\mathbf{p}_{T_k} - \hat{\mathbf{p}}_{T_{k|k-1}}||^2_{\mathbf{P}_{\mathbf{pp}_{k|k-1}}} + \frac{1}{2}||z_k - h(\mathbf{p}_{T_k})||^2_{\sigma_k^2} \quad (15)$$

where $\mathbf{P}_{\mathbf{pp}_{k|k-1}}$ is the covariance matrix corresponding to the target position, obtained by partitioning the covariance matrix as: $\mathbf{P}_{k|k-1} := \begin{bmatrix} \mathbf{P}_{\mathbf{pp}_{k|k-1}} & \mathbf{P}_{\mathbf{pd}_{k|k-1}} \\ \mathbf{P}_{\mathbf{dp}_{k|k-1}} & \mathbf{P}_{\mathbf{dd}_{k|k-1}} \end{bmatrix}$. To solve (15) analytically, by introducing a new variable $\rho_k = h(\mathbf{p}_{T_k})$, we have the following equivalent *constrained* minimization:[3]

$$\min_{\mathbf{p}_{T_k},\, \rho_k} \frac{1}{2}||\mathbf{p}_{T_k} - \hat{\mathbf{p}}_{T_{k|k-1}}||^2_{\mathbf{P}_{\mathbf{pp}_{k|k-1}}} + \frac{1}{2}||z_k - \rho_k||^2_{\sigma_k^2} \quad (16)$$

$$\text{s.t. } \rho_k^2 = (x_{S_k} - x_{T_k})^2 + (y_{S_k} - y_{T_k})^2\ ,\ \rho_k \geq 0 \quad (17)$$

We solve (16)-(17) by employing the method of Lagrange multipliers [10]. Specifically, without loss of generality, by assuming $\mathbf{P}^{-1}_{\mathbf{pp}_{k|k-1}} := \mathbf{Diag}(s_1, s_2)$, the Lagrangian function is constructed as follows:[4]

$$\mathcal{L}(x_{T_k}, y_{T_k}, \rho_k, \lambda) = \frac{s_1}{2}(x_{T_k} - \hat{x}_{T_{k|k-1}})^2 + \frac{s_2}{2}(y_{T_k} - \hat{y}_{T_{k|k-1}})^2 + \frac{(z_k - \rho_k)^2}{2\sigma_k^2} + \lambda\left(\rho_k^2 - (x_{S_k} - x_{T_k})^2 - (y_{S_k} - y_{T_k})^2\right) \quad (18)$$

[2] The dynamic constraint (1) can be treated analogously as a measurement.

[3] Similar derivations of analytically solving (11) can be found in our previous work [7], which are briefly described here for completeness.

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

where λ is the Lagrange multiplier. Setting the derivatives of $\mathcal{L}(\cdot)$ w.r.t. the four optimization variables to zero, and performing simple algebraic manipulations, we obtain:

$$\frac{\partial \mathcal{L}}{\partial x_T} = 0 \Rightarrow x_{T_k} = \frac{s_1 \hat{x}_{T_{k|k-1}} - 2\lambda x_{S_k}}{s_1 - 2\lambda} \tag{19}$$

$$\frac{\partial \mathcal{L}}{\partial y_T} = 0 \Rightarrow y_{T_k} = \frac{s_2 \hat{y}_{T_{k|k-1}} - 2\lambda y_{S_k}}{s_2 - 2\lambda} \tag{20}$$

$$\frac{\partial \mathcal{L}}{\partial \rho_k} = 0 \Rightarrow \rho_k = \frac{z_k}{1 + 2\sigma_k^2 \lambda} \tag{21}$$

$$\frac{\partial \mathcal{L}}{\partial \lambda} = 0 \Rightarrow 0 = \rho_k^2 - (x_{S_k} - x_{T_k})^2 - (y_{S_k} - y_{T_k})^2 \tag{22}$$

We substitute (19)-(21) into (22) and multiply both sides of (22) with $(1 + 2\sigma_k^2\lambda)^2(s_1 - 2\lambda)^2(s_2 - 2\lambda)^2$, to obtain a fourth-order *univariate* polynomial in λ, i.e., $0 = f(\lambda) = \sum_{i=0}^{4} a_i \lambda^i$, where a_i, $i = 0, \ldots, 4$, are the coefficients expressed in terms of the known quantities s_1, s_2, z_k, σ_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, we compute its roots in closed form. Although, in general, there exist 4 solutions for λ and thus 4 solutions for x_{T_k}, y_{T_k} and ρ_k, as they depend injectively on λ [see (19)-(21)], we only need to consider the pairs (x_{T_k}, y_{T_k}) that correspond to real solutions for λ and to a nonnegative ρ_k [see (17)]. Moreover, since some of these solutions could be local maxima or saddle points, the second-order derivative test [10] is employed to identify the minima. In fact, it was shown in [7] that there are at most 2 local minima for the problem (15). Once all the local minima for the target position are determined, we can accordingly compute the corresponding estimates for $\mathbf{d}_{T_k}$ [7].

3.2. Range-only tracking using the B-iMAP

When applying the proposed B-iMAP to range-only tracking, the key idea is to use the analytically-computed local minima at each time step as guidance to find the most probable hypotheses of the target trajectory, and then to track them efficiently by incremental QR-factorization. Specifically, at time-step $k-1$, based on (9) and (10), we first propagate the current state estimate corresponding to the i-th hypothesis and its covariance matrix, $\hat{\mathbf{x}}^{[i]}_{k-1|k-1}$ and $\mathbf{P}^{[i]}_{k-1|k-1}$, $i = 1, 2, \ldots, m$ (m is the number of estimators in the bank at time-step $k-1$). Then, once a new range measurement becomes available, the propagated state estimate and covariance, $\hat{\mathbf{x}}^{[i]}_{k|k-1}$ and $\mathbf{P}^{[i]}_{k|k-1}$, are used as the prior in (11). Next, we use the algebraic-geometry method described in Section 3.1 to determine all the local minima of (11) analytically, denoted by $\mathbf{x}^{[j]}_k$, $1 \leq j \leq 2m$. For each of these solutions, we incrementally update the QR factorization so as to efficiently solve for the entire trajectory estimate $\hat{\mathbf{x}}^{[j]}_{0:k|k}$. Note that batch relinearization is performed periodically every τ time steps to reduce the linearization errors. In the end, we will have multiple candidates of the MAP estimate, among which the one with the least cost is selected as the best estimate for the global optimum.

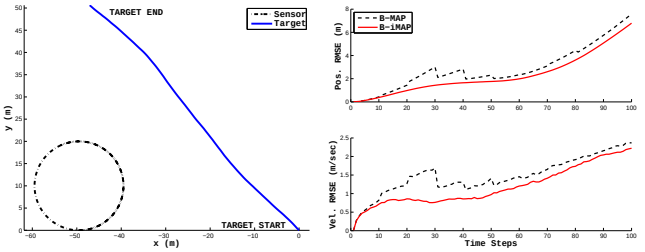

(a) Sensor/target's trajectories (b) Avg. root mean square errors

Fig. 1. Monte-Carlo simulation results show that the proposed B-iMAP achieves better accuracy than the B-MAP [7].

3.3. Numerical results

We validate the proposed B-iMAP in 100 Monte-Carlo simulations, as compared to the bank of MAP (B-MAP) [7].[5] The B-MAP performs batch relinearization at every time step, *with* marginalization of old states every 25 time steps to reduce the computational cost. In contrast, the proposed B-iMAP periodically relinearizes the system every $\tau = 10$ time steps, *without* marginalization. In both estimators, we prune out the least probable hypotheses based on the MAP costs and keep $m \leq 10$ most probable hypotheses. In this test, we adopt a zero-acceleration motion model [2] for the target [see Fig. 1(a)], and set the standard deviation of the measurement noise equal to 10% of the sensor-to-target distance. Fig. 1(b) shows that the proposed B-iMAP, in average, performs **19%** more accurately than the B-MAP, which even comes at a **65%** lower cost. Specifically, the average CPU runtime (averaging over all Monte Carlo runs and over all time steps) of our Matlab implementation on a Core2 Quad CPU is 0.0416 sec for the B-iMAP, compared to 0.1217 sec for the B-MAP. This superior performance of the B-iMAP is attributed to the efficient incremental update (see Section 2.2). This also implies that in solving a batch-MAP problem, marginalization can be substituted, or at least postponed, by the application of incremental QR update, while achieving better efficiency and accuracy.

4. CONCLUSIONS

In this paper, we have introduced a novel efficient multi-hypothesis tracking algorithm, i.e., the B-iMAP. The key idea of this approach is to analytically determine the most probable hypotheses at each time step, which is attained by transforming the one-step MAP problem into polynomial form. By reusing the previous computations, the incremental QR-factorization is employed to efficiently track all the hypotheses. We have applied the B-iMAP to the particular problem of range-only tracking and shown its superior performance in terms of both accuracy and efficiency.

[5] As shown in [7], the B-MAP which also determines most probable hypotheses analytically, outperforms the standard multi-hypothesis tracking approaches, and thus we focus only on its comparison to the proposed B-iMAP.

5. REFERENCES

[1] S. S. Blackman and R. F. Popoli, *Design and Analysis of Modern Tracking Systems*, Artech House, 1999.

[2] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*, New York: Wiley, 2001.

[3] 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.

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

[5] T. R. Kronhamn, "Bearings-only target motion analysis based on a multihypothesis Kalman filter and adaptive ownship motion control," *IEE Proc. - Radar, Sonar and Navigation*, vol. 145, pp. 247–252, Aug. 1998.

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

[7] G. P. Huang, K. X. Zhou, N. Trawny, and S. I. Roumeliotis, "A bank of MAP estimators for single-sensor range-only target tracking," in *Proc. of the American Control Conference*, Baltimore, MD, June 30–July 2, 2010, pp. 6974–6980.

[8] M. Kaess, A. Ranganathan, and F. Dellaert, "iSAM: Incremental smoothing and mapping," *IEEE Trans. on Robotics*, vol. 24, no. 6, pp. 1365–1378, Dec. 2008.

[9] G. H. Golub and C. F. Van Loan, *Matrix Computations*, The Johns Hopkins University Press, 1996.

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