The PMHT: Solutions for Some of its Problems

Monika Wieneke Wolfgang Koch

wieneke@fgan.de w.koch@fgan.de www.fgan.de

FGAN-FKIE, Neuenahrer Str. 20, D-53343 Wachtberg, Germany

Abstract: Tracking multiple targets in a cluttered environment is a challenging task. The Probabilistic Multi-Hypothesis Tracker (PMHT) is an efficient approach to cope with it. Linearity in the number of targets and measurements is the main motivation for its continued study. Unfortunately, the PMHT has not yet shown its superiority in terms of a better track-lost statistics. Furthermore, the problem of track extraction is not satisfactorily solved. This work focuses on mitigating the PMHT's main problems and presents the integration of a sequential likelihood-ratio test for track extraction.

The PMHT works on a sliding data-window. For each window-position it iteratively applies a Kalman smoother using synthetic measurements. In [WRS02] three properties are made responsible for the PMHT's problems in track maintenance and its small practical acceptance: **Non-Adaptivity**, **Hospitality** and **Narcissism**.

We derive new synthesis weights governed by innovation covariances to make the PMHT adaptive. To avoid hospitality we introduce a spurious measurement representing a missing detection. Finally we resize the estimation errors of the start iteration to rouse up the PMHT from its narcissism. To introduce our notations we continue by a formal description of tracking.

The tracking **Scenario** is defined as follows: A sensor observes S point targets in its field of view (FoV). It generates measurements $\mathcal{Z} = \boldsymbol{z}_{1:T} = \{\boldsymbol{z}_t, N_t\}_{t=1}^{T}$ for a period $[1:T]$. The sensor output at a scan t not only consists of the set of measurements $\boldsymbol{z}_t$ but also of the number of measurements N_t. Thus we model measured data as a pair $\{\boldsymbol{z}_t, N_t\}$. Measurements $\boldsymbol{z}_t^n \in \mathbb{R}^2$ with $n \in [1:N_t]$ are assumed to be Cartesian position data. The spurious measurement $n = 0$ denotes a missing detection as already mentioned.

The **Task** of tracking consists in estimating the kinematic states $\mathcal{X} = \boldsymbol{x}_{1:T}$ of the observed targets. The states $\boldsymbol{x}_t^s \in \mathbb{R}^6$ with $s \in [1:S]$ comprise position, velocity and acceleration. Difficulties arise from unkown associations $\mathcal{A} = \boldsymbol{a}_{1:T} = \{\boldsymbol{a}_t\}_{t=1}^{T}$ of measurements to targets. They are modeled as random variables $\boldsymbol{a}_t = \{a_t^n\}_{n=0}^{N_t}$ that map each measurement $n \in [0:N_t]$ to one of the targets $s \in [0:S]$ by assigning $a_t^n = s$. The target $s = 0$ is a spurious planar target representing clutter. It corresponds to the FoV.

Mathematically expressed the optimization problem $\arg\max_{\mathcal{X}} p(\mathcal{X}|\mathcal{Z})$ is to be solved. **Expectation-Maximization** is an efficient method for this task.

The remainder of this work starts with a brief explanation of Expectation-Maximization whereas we introduce new posterior weights. Section 2 continues with the derivation of the PMHT and presents our ideas to solve its problems. In section 3 we explain the integration of likelihood-ratio testing into the PMHT.

1 Expectation-Maximization

Expectation-Maximization (EM) is an iterative method for localizing posterior modes. At each iteration, EM first calculates posterior weights $p(\mathcal{A}|\mathcal{Z}, \mathcal{X}^l)$. They define an optimal lower bound $\mathcal{Q}(\mathcal{X}; \mathcal{X}^l)$ of $p(\mathcal{X}|\mathcal{Z})$ at the current guess $\mathcal{X}^l$. In contrast to the conventional approach we regard the estimation error covariances $\mathcal{P}^l$ as intrinsic information of $\mathcal{X}^l$. To make it explicit we use **extended posterior weights** $p(\mathcal{A}|\mathcal{Z}, \mathcal{X}^l, \mathcal{P}^l)$ in our work (eqn. 1).

$$\mathcal{Q}(\mathcal{X}; \mathcal{X}^l) = \log p(\mathcal{X}) + \sum_{\mathcal{A}} \log(p(\mathcal{A}, \mathcal{Z}|\mathcal{X}))\, p(\mathcal{A}|\mathcal{Z}, \mathcal{X}^l, \mathcal{P}^l) \tag{1}$$

As $\mathcal{Q}(\mathcal{X}; \mathcal{X}^l)$ is expressed as an expectation, the step is called E-Step. In the following M-Step, EM maximizes the bound with respect to the free variable $\mathcal{X}$. How this is done depends on the application. The PMHT is the application of EM to the tracking problem. It results in estimates $\boldsymbol{x}_t^s$ for each target $s \in [1:S]$ at each time $t \in [1:T]$. Covariance matrices P_t^s occur as by-product. We interpret them as estimation error covariances of $\boldsymbol{x}_t^s$.

2 Derivation of the modified PMHT

The $\mathcal{Q}$-Function contains all available information: the statistical models of detection process, measurement process and target dynamics. A series of calculations is required to make the information visible. We pass on deriving dynamics and sensor model and proceed with our new posterior weights. Disposed readers are referred to [SL95].

Adaptive Posterior Weights $\boldsymbol{w_t^{lns}}$

This section addresses the PMHT's problem of **Non-Adaptivity**. As we model the sensor output as a pair $\{\boldsymbol{z}_t, N_t\}$, we can split it and treat N_t separately. Some simpler calculations followed by Bayes' Rule and the product formula A.1 finally yield eqn. 2.

$$p(\mathcal{A}|\mathcal{Z}, \mathcal{X}^l, \mathcal{P}^l) = \prod_{t=1}^{T} \frac{\prod_{n=0}^{N_t} \mathcal{N}(\boldsymbol{z}_t^n; H\boldsymbol{x}_t^{ls}, S_t^{ls})\pi_t^{ns}}{\prod_{n=0}^{N_t} \sum_{s'=0}^{S} \mathcal{N}(\boldsymbol{z}_t^n; H\boldsymbol{x}_t^{ls'}, S_t^{ls'})\pi_t^{ns'}} =: \prod_{t=1}^{T} \prod_{n=0}^{N_t} w_t^{lns} \tag{2}$$

Our weights are controlled by the **innovation** covariances $S^{ls} := HP_t^{ls}H^T + R$ with measuring matrix H and measuring error R. Using these weights the PMHT works **adaptively** because it takes the quality P_t^{ls} of the current track estimation into account. If P_t^{ls} blows up soon enough, a track rescue is possible. The weights comprise two kinds of measures that evaluate the relevance of a measurement with respect to a target estimation: A distance measure $\mathcal{N}(\boldsymbol{z}_t^n; H\boldsymbol{x}_t^{ls}, S^{ls})$ and a 'visibility measure' denoted as $\pi_t^{ns} := p(a_t^n = s|N_t)$.

For $n > 0$ the latter reflects how likely it is to hit a target, not taking concrete position data into account. The weight π_t^{0s} simply is the probability of missing a target. Note that our visibility weights are posteriors depending on N_t. The original PMHT uses priors

$p(a_t^n = s)$ instead and hence is less flexible. For the calculation of π_t^{ns} we apply Bayes' Rule. Its outputs $p(a_t^n = s)$ and $p(N_t | a_t^n = s)$ are easier to handle than π_t^{ns} itself. Using binomial coefficients we also come to grips with multiple targets. By increasing N_t the weights π_t^{ns} of the real measurements ($n > 0$) converge to a uniform distribution.

Maximization of the $\mathcal{Q}$-Function

As the $\mathcal{Q}$-function can be rewritten as a sum over the targets, the maximization problem decomposes into S independent problems: one summand $\mathcal{Q}^s(\mathcal{X}; \mathcal{X}^l)$ per target. Exponentiation and successive application of the product formula A.2 yields relation 3 with evolution matrix F and process noise covariance D. $\bar{\boldsymbol{z}}_t^{ls}$ and $\bar{R}_t^{ls}$ denote synthetic measurements and corresponding error covariances respectively.

$$\exp \mathcal{Q}^s(\mathcal{X}; \mathcal{X}^l) \propto \mathcal{N}(\boldsymbol{x}_0^s; \boldsymbol{x}_{0|0}^s, P_{0|0}^s) \prod_{t=1}^{T} \mathcal{N}(\boldsymbol{x}_t^s; F\boldsymbol{x}_{t|t-1}^s, D) \prod_{n=0}^{N_t} \mathcal{N}(\bar{\boldsymbol{z}}_t^{ls}; H\boldsymbol{x}_t^s, \bar{R}_t^{ls}) \quad (3)$$

$$\text{with} \quad \bar{\boldsymbol{z}}_t^{ls} = \bar{R}_t^{ls} \sum_{n=0}^{N_t} w_t^{lns} (R_t^n)^{-1} \boldsymbol{z}_t^n \quad \text{and} \quad \bar{R}_t^{ls} = \left(\sum_{n=0}^{N_t} w_t^{lns} (R_t^n)^{-1} \right)^{-1} \quad (4)$$

At this stage the spurious measurement $n = 0$ makes an impact: Obeying the formalism we have to **renormalize** the posterior weights with respect to all measurements including the missing detection $n = 0$ and exchange w_t^{lns} by $w_t^{*\,lns} = w_t^{lns} / \sum_{n=0}^{N_t} w_t^{lns}$. Thereby an intermediate result on the way to eqn. 2 enables us to set the weight w_t^{0s} of the missing detection to π_t^{0s}. As its error covariance is $R_t^0 = \infty$ the corresponding summands in eqn. 4 vanish. So in a Cartesian system we finally obtain centroid measurements with covariances

$$\bar{R}_t^{ls} = \frac{R}{\sum_{n=1}^{N_t} w_t^{*\,lns}} \quad \text{and} \quad \sum_{n=1}^{N_t} w_t^{*\,lns} < 1 \quad (5)$$

As in the original PMHT the sum of weights in eqn. 5 can be greater than 1 it suffers from **Hospitality**. It interprets multiple measurements as one measurement of high accuracy. We have enforced the sum to be less than 1 and hence mitigated the hospitality problem.

Initializing the Iteration Process

The PMHT works on a sliding data-window. The initial states of the current window are set to the final estimates of the preceding window ($t_{\text{pre}} \in [2:T]$) and the corresponding prediction ($t_{\text{pre}} = T + 1$). If the latter states only have a slight tendency to walk off, the estimations of the current window often pursue this. Even a single false alarm near a poor **prediction** can cause wrong tendencies. However for the 'narcissistic' PMHT, the track is progressing normally: It perfectly fits the estimations to the new data situation, though it should know it better from the past.

A proper initiation can mitigate the PMHT's **Narcissism**. Consider a single iteration l before retrodiction: It consists in weight calculation, filtering and prediction (in that order) for each time scan t. We interfere between weight calculation and filtering step at each scan of the first iteration ($l = 0$). The goal is to give certainty about the already estimated states. Especially in case of missing detections the ellipsis $P^{0s}_{t|t-1}$ of the predicted estimation error comprises too much uncertainty. We propose to use the corresponding final error covariance $P^s_{t_{\rm pre}+1}$ of the preceding data-window instead. As $P^s_{t_{\rm pre}+1}$ is usually smaller, the PMHT is reminded of having already estimated the states $t = 1{:}T-1$ in the preceding iteration loop. Note that the size does not change for $t = T$.

Experimental Example

We simulated an aircraft with detection probability $P_D = 0.7$, observed by a radar: time interval $\Delta t = 5\,\mathrm{s}$, clutter density $\rho = 10^{-7.3}/m^2$, FoV-radius 25000 m and error $\sigma = (50\,\mathrm{m}, 50\,\mathrm{m})$. Figure 1 shows our PMHT on the left and the failing PDAF on the right. We chose a window length of 6 time scans and a constant number of 6 iterations. Measurements are marked as blue $+$, **final** state estimations as red $\times$. Measurements of the real target are bordered by $\square$. We show the estimation errors at the last two scans ($t = T-1, T$) of the first two iterations ($l = 0, 1$) in each case: error ellipsis before filtering in green, after filtering in red. The data situations that make the narcissistic PMHT walk off are marked as 'critical'. False alarms are plotted at all scans, only within a radius of 3000 m around the initial ($l = 0$) estimation of scan $t = T$.

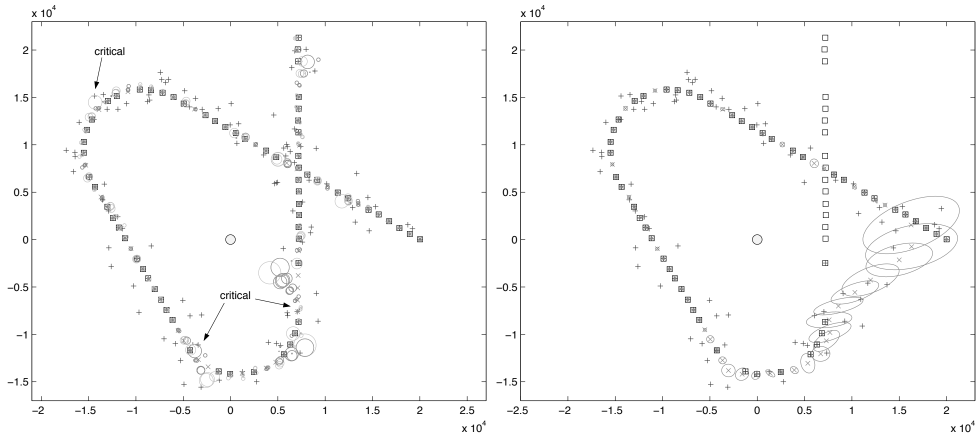

Figure 1: The new PMHT on the left and the PDAF on the right (◯ at (0,0) ≡ sensor).

3 Sequential Track Extraction

In [vK98] a sequential **Likelihood-Ratio** (LR) test has been proposed to cope with the task of track extraction and deletion. It consists in successively updating the ratio between likelihood functions $p(\boldsymbol{z}_{1:t}|H_S)$ and $p(\boldsymbol{z}_{1:t}|H_0)$ (eqn. 6). H_S denotes the hypothesis that $\boldsymbol{z}_{1:T}$ contains data from S targets and possibly clutter. H_0 represents the case that all data in $\boldsymbol{z}_{1:T}$ are false. The goal is to decide as fast as possible and as reliably as requested between

H_S and H_0. For this purpose two thresholds $A :\approx (1-P_S)/(1-P_0)$ and $B :\approx P_S/P_0$ are introduced whereas $P_S := \text{Prob}(\text{accept } H_S|H_S)$ and $P_0 := \text{Prob}(\text{accept } H_S|H_0)$ are the related statistical decision errors. If the LR at the current time scan is less than A the algorithm accepts H_0, if it is greater than B it accepts H_S, otherwise it waits for new data.

$$LR_S(t) = \frac{p(\boldsymbol{z}_{1:t}|H_S)}{p(\boldsymbol{z}_{1:t}|H_0)} = \underbrace{\frac{p(\boldsymbol{z}_t|N_t, \boldsymbol{z}_{1:t-1}, H_S)}{p(\boldsymbol{z}_t|N_t, \boldsymbol{z}_{1:t-1}, H_0)}}_{=: F_1} \cdot \underbrace{\frac{p(N_t|H_S)}{p(N_t|H_0)}}_{=: F_2} \cdot \underbrace{\frac{p(\boldsymbol{z}_{1:t-1}|H_S)}{p(\boldsymbol{z}_{1:t-1}|H_0)}}_{= LR_S(t-1)} \tag{6}$$

As the LR-test has been successfully embedded into the framework of multi-hypothesis tracking we are motivated to integrate it into the PMHT in an analogous manner. The key idea on adopting the LR-method is to replace the hypotheses H_S and H_0 in F_1 by hypotheses $H_S \wedge (P_D \gg 0)$ and $H_S \wedge (P_D \approx 0)$ respectively. We continue by including the target states $\boldsymbol{x}_t$, and the **synthetic** measurements of the final PMHT-iteration (eqn. 7).

$$p(\boldsymbol{z}_t|N_t, \boldsymbol{z}_{1:t-1}, H_S, P_D \gg 0) \propto \int_{\boldsymbol{x}_t} p(\bar{\boldsymbol{z}}_t, \boldsymbol{x}_t|N_t, \boldsymbol{z}_{1:t-1}, H_S, P_D \gg 0)\ d\boldsymbol{x}_t \tag{7}$$

As we assume well separated targets it is permitted to process each target separately and multiply the results at the end. Briefly summarized the steps that lead to a single factor of eqn. 8 are: Marginalization over the target's detection status d_t^s ($d_t^s \equiv$ detected, $\neg d_t^s \equiv$ not detected), some basic transformations and application of the product formula A.1.

$$\frac{p(\boldsymbol{z}_t|N_t, \boldsymbol{z}_{1:t-1}, H_S, P_D \gg 0)}{p(\boldsymbol{z}_t|N_t, \boldsymbol{z}_{1:t-1}, H_S, P_D \approx 0)} = \prod_{s=1}^{S} \frac{\pi_t^{ds} \mathcal{N}(\bar{\boldsymbol{z}}_t; H\boldsymbol{x}_{t|t-1}^s, \bar{S}_t^s) + \pi_t^{\neg ds} \frac{1}{|\text{FoV}|}}{\underbrace{\pi_t^{ds}}_{\approx 0} \mathcal{N}(\bar{\boldsymbol{z}}_t; H\boldsymbol{x}_{t|t-1}^s, \bar{S}_t^s) + \underbrace{\pi_t^{\neg ds}}_{\approx 1} \frac{1}{|\text{FoV}|}} \tag{8}$$

In eqn. 8 $\bar{S}_t^s := HP_{t|t-1}^s H^T + \bar{R}_t$ denotes the synthetic innovation covariance of target s and $\boldsymbol{x}_{t|t-1}^s$ the predicted target state at the last scan t. Analogously to π_t^{ns}, we define $\pi_t^{\neg ds} := p(\neg d_t^s|N_t) = p(a_t^0 = s|N_t)$ and $\pi_t^{ds} := p(d_t^s|N_t) = N_t \cdot p(a_t^n = s|N_t)$ for arbitrary $n > 0$. At this stage $p(a_t^n = s|N_t)$ is normalized with respect to the measurements.

Finally our LR-formula for the PMHT is presented in eqn. 9

$$LR_S(t) = \prod_{s=1}^{S} \left(\pi_t^{ds} \mathcal{N}(\bar{\boldsymbol{z}}_t; H\boldsymbol{x}_{t|t-1}^s, \bar{S}_t^s) \cdot |\text{FoV}| + \pi_t^{\neg ds} \right) \cdot F_2 \cdot LR_S(t-1) \tag{9}$$

As $LR_S(t)$ is a **by-product** of the PMHT we propose to treat every measurement as a potential 'track-seed' and start a separate PMHT for it. At the end of the iteration process we check $LR_S(t)$ and either make a decision or expand the PMHT-window to start a new iteration process.

Experimental Example

We generated another aircraft simulation: $\Delta t = 5\,\text{s}$, clutter density $\rho = 10^{-7.5}/m^2$, FoV-radius 30000 m, error $\sigma = (50\,\text{m}, 50\,\text{m})$. The detection probability of the aircraft is 0.8. Figure 2 shows real measurements as $+$, labeled by scan numbers. False alarms are marked as $\times$. The total length of observation is 20 scans. The target vanishes at scan 16. We let the PMHT-window grow up to 7 time scans and shifted it until the head reached scan 20. The table shows the effect of missing detections on $F_1 \cdot F_2$ at $t = 7, 10, 13, 16 - 20$.

t	$F_1 \cdot F_2$	t	$F_1 \cdot F_2$
3	123462.433	12	23651.016
4	20796.787	**13**	**0.010**
5	12013.697	14	4093.070
6	781.654	15	18419.864
7	**0.014**	**16**	**0.011**
8	138.704	**17**	**7.142**
9	18501.717	**18**	**0.302**
10	**0.010**	**19**	**0.010**
11	2884.916	**20**	**1.018**

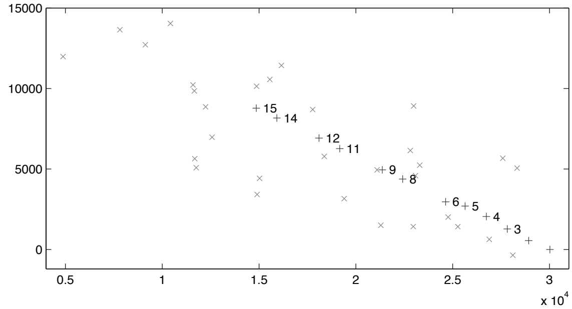

Figure 2: LR-Factors $F_1 \cdot F_2$

4 Conclusion and Future Work

In the first part we presented some means to mitigate the PMHT's main problems. What lacks is a proper statistical evaluation of our method. In the second part we sketched a new approach for track extraction by PMHT. An open question in this context is the appropriate choice of the bounds A and B. In view of the values $F_1 \cdot F_2$ it surely makes sense to work with $\log LR_S(t)$. The computational load of the approach will also be an important topic.

A Product Formulae for Gaussians

The formulae transform a product of gaussians into another product of gaussians.

$$\mathcal{N}(\mathrm{x}; \mathrm{Xy}, \mathrm{Y})\mathcal{N}(\mathrm{y}; \mathrm{z}, \mathrm{Z}) = \mathcal{N}(\mathrm{x}; \mathrm{a}, \mathrm{A})\mathcal{N}(\mathrm{y}; \mathrm{b}, \mathrm{B}) \tag{A.1}$$

with $\mathrm{a} = \mathrm{Xz}$, $\mathrm{b} = \mathrm{z}+\mathrm{W}(\mathrm{x}-\mathrm{Xz})$, $\mathrm{A} = \mathrm{XZX^T}+\mathrm{Y}$, $\mathrm{B} = \mathrm{Z}-\mathrm{WAW^T}$ and $\mathrm{W} := \mathrm{ZX^T A^{-1}}$.

$$\mathcal{N}(\mathrm{x}; \mathrm{y}, \mathrm{Y})\mathcal{N}(\mathrm{x}; \mathrm{z}, \mathrm{Z}) = \mathcal{N}(\mathrm{x}; \mathrm{a}, \mathrm{A})\mathcal{N}(\mathrm{y}; \mathrm{b}, \mathrm{B}) \tag{A.2}$$

with $\mathrm{a} = \mathrm{A}\left(\mathrm{Y}^{-1}\mathrm{y} + \mathrm{Z}^{-1}\mathrm{z}\right)$, $\mathrm{b} = \mathrm{z}$, $\mathrm{A} = \left(\mathrm{Y}^{-1} + \mathrm{Z}^{-1}\right)^{-1}$ and $\mathrm{B} = \mathrm{Y} + \mathrm{Z}$.

References

[SL95] Roy L. Streit and T. E. Luginbuhl. Probabilistic Multi-Hypothesis Tracking. Technical Report NUWC-NPT/10/428, Naval Undersea Warefare Center Division, 1995.

[vK98] Günter van Keuk. Sequential Track Extraction. *IEEE Transactions on Aerospace and Electronic Systems*, 34(4):1135–1148, 1998.

[WRS02] Peter K. Willett, Yanhua Ruan, and Roy L. Streit. PMHT: Problems and Some Solutions. *IEEE Transactions on Aerospace and Electronic Systems*, 38(3):738–754, 2002.