A COMPARISON OF THE JPDAF AND PMHT TRACKING ALGORITHMS

Constantino Rago and Peter Willett

University of Connecticut
Storrs, CT 06269

Roy Streit

Naval Undersea Warfare Center
Newport, RI 02841

ABSTRACT

Here we analyze the tracking characteristics of a new data-association/tracking algorithm proposed by Streit and Luginbuhl [1], the Probabilistic Multi Hypothesis Tracker (PMHT). The algorithm uses a recursive method (known amongst statisticians as the *Expectation-Maximization* or *EM* method) to compute in an optimal way the associations between measurements and targets. Until now, no comparative performance analysis has been done. In this paper, we compare the performance of this new scheme to that of a commonly used tracking algorithm, the Joint Probabilistic Data Association Filter (JPDAF).

1. INTRODUCTION

In a multi-target scenario or in a heavy-clutter environment, we are faced with multiple measurements and multiple targets. Before any estimation can be done the problem of deciding which measurements correspond to which targets, and which to clutter – the data association problem – must be dealt with. The PMHT accomplishes this in a novel way: its associations are "soft", and each represents the posterior probability (given the observations) that each measurement is associated with each corresponding target.

In this paper we first describe the operation of the PMHT, and the assumptions on which it is based. We then discuss modifications to the basic PMHT (as described in [1]) which are necessary that it work effectively with clutter and missed detections (these and other issues are explored in more detail in [2]). Finally we compare the PMHT to the JPDAF in a two-target parallel trajectory scenario. It is shown that fewer tracks are lost when the PMHT is used.

THIS WORK WAS SPONSORED BY THE OFFICE OF NAVAL RESEARCH

2. OPERATION OF THE PMHT

2.1. The Algorithm

The scenario to which the PMHT is to be applied is perhaps the purest target tracking scheme: there are assumed to be M targets, the s^{th} of which moves according to the discrete-time linear model:

$$\begin{aligned} \mathbf{x}_s(t+1) &= \mathbf{F}_s(t)\mathbf{x}_s(t) + \mathbf{G}_s(t)\mathbf{u}_s(t) + \mathbf{v}_s(t) \\ \mathbf{y}_s(t) &= \mathbf{H}_s(t)\mathbf{x}_s(t) + \mathbf{w}_s(t) \end{aligned} \quad (1)$$

for $t = 1, 2, \ldots, T$. Here, as usual, $\mathbf{x}_s(t)$ represents the trajectory of the s^{th} model at time t, and $\mathbf{y}_s(t)$ its corresponding observation; the model matrices are known, and are assumed to represent an observable and controllable system. The random sequences $\{\mathbf{v}_s(t), \mathbf{w}_s(t)\}$ are assumed white, zero-mean, Gaussian, and mutually independent, with $E\{\mathbf{v}_s(t)\mathbf{v}_s^T(t)\} = \mathbf{Q}_s(t)$ and $E\{\mathbf{w}_s(t)\mathbf{w}_s^T(t)\} = \mathbf{R}_s(t)$. The control sequences $\{\mathbf{u}_s(t)\}$ are known, and since these contribute in the form of "ownship" platform motion in a straightforward but notationally-inconvenient way, we shall take them as zero with no loss of generality.

Now, naturally in a multi-target tracking scenario the thorniest problem is of *data-association*; that is, how to determine which measurements come from which targets. For each t we define $\{k_r(t), \mathbf{z}_r(t)\}_{r=1}^{n_t}$ such that

$$\mathbf{z}_r(t) = \mathbf{y}_{k_r(t)}(t) \quad (2)$$

meaning that the r^{th} measurement at time t comes from model $k_r(t)$ – and, of course, $k_r(t)$ is unknown. As an aside, please note that we have denoted the number of observations (detections) at time t as n_t – which is not necessarily identical to M, the number of targets – but that we have at least for now insisted that each measurement comes from *some* target. We will discuss false-alarms shortly.

What remains is to provide a probabilistic structure for the measurement/target associations: our assumption is that

$$Pr(k_r(t) = s) = \pi_s \quad (3)$$

0-7803-2431-5/95 $4.00 © 1995 IEEE

and that all are *independent*. This is to be contrasted with the assumption made by the JPDAF (and MHT) [3] that

$$\{k_r(t) = j\} \Longrightarrow \{k_s(t) \neq j\} \tag{4}$$

for $r \neq s$, which introduces *dependence*. The independence of (3) is necessary that the PMHT have an appealing computational structure; however, it must be admitted that (4) is of greater practical significance. As such, in the simulations that follow, assumption (4) *will* be employed – we will show that the PMHT performs well regardless of this.

The operation of the PMHT is as follows:

1. Determine initial values for the trajectory variables $\mathbf{x}_s^1(t)$ for all targets $s = 1, 2, \ldots, M$ and all times $t = 1, 2, \ldots, T$. Set the EM iteration index $n = 1$.

2. Calculate the posterior association probabilities $w_{l,r}^n(t)$ for all targets $l = 1, 2, \ldots, M$, and for all times $t = 1, 2, \ldots, T$, and measurements $r = 1, 2, \ldots, n_t$, according to:
$$w_{l,r}^n(t) = \frac{\pi_l \mathcal{N}\{\mathbf{z}_r(t); \mathbf{H}_l(t)\mathbf{x}_l^n(t), \mathbf{R}_l(t)\}}{\sum_{p=1}^M \left[\pi_p \mathcal{N}\{\mathbf{z}_r(t); \mathbf{H}_p(t)\mathbf{x}_p^n(t), \mathbf{R}_p(t)\}\right]} \tag{5}$$
where $\mathcal{N}$ refers to a multivariate Gaussian *pdf* with the specified mean and covariance.

3. Calculate synthetic measurements $\tilde{\mathbf{z}}_s(t)$ and their associated (synthetic) measurement covariances $\tilde{\mathbf{R}}_s(t)$ for all targets $s = 1, 2, \ldots, M$ and times $t = 1, 2, \ldots, T$, according to
$$\tilde{\mathbf{z}}_s(t) \equiv \frac{\sum_{r=1}^{n_t} w_{sr}^n(t)\mathbf{z}_r(t)}{\sum_{r=1}^{n_t} w_{sr}^n(t)}$$
$$\tilde{\mathbf{R}}_s(t) \equiv \frac{\mathbf{R}_s(t)}{\sum_{r=1}^{n_t} w_{sr}^n(t)} \tag{6}$$
respectively.

4. For each target $s = 1, 2, \ldots, M$, use a Kalman smoothing algorithm to obtain the updated estimated trajectory $\mathbf{x}_s^{n+1}(t)$, according to the model of (1), but with the difference that the *synthetic* measurements and covariances $\tilde{\mathbf{z}}_s(t)$ and $\tilde{\mathbf{R}}_s(t)$ are used in place of $\mathbf{y}_s(t)$ (which naturally is not known) and $\mathbf{R}_s(t)$.

5. Increment $n = n + 1$, and return to step 2 for the next EM iteration, unless a stopping criterion is reached.

Experience shows that reasonably accurate solutions are available after 3-5 iterations, and that there is seldom any change at all after 10-20 iterations (this, of course, is problem-dependent).

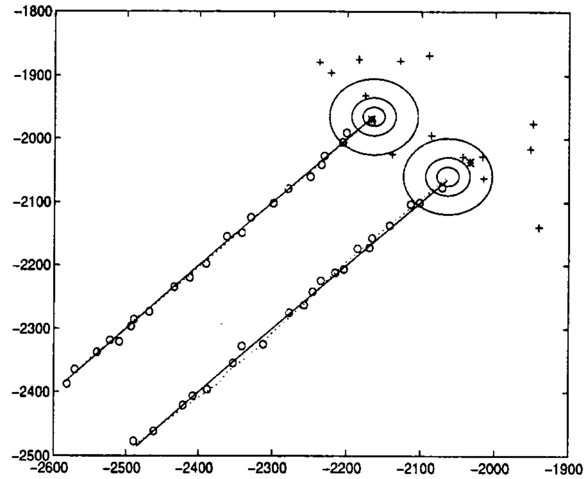

Figure 1: The situation in a typical run of the PMHT, after the sixth batch of length 5. The clutter density is $\lambda = 10^{-4}$, and the probability of detection is $P_d = 0.7$. A detection is denoted by a o, a false-alarm (only one snapshot is shown of these) by a +, and each current probabilistic centroid (i.e. $\tilde{\mathbf{z}}$) by a *. The circles surrounding the targets represent one-σ regions for homothetic sub-modes corresponding to $\kappa_1 = 1$, $\kappa_2 = 4$, and $\kappa_3 = 16$. False-alarms too distant from the targets to be of any interest are suppressed.

2.2. False Alarms

The PMHT as written makes no provision for false detections. That is, for each detection the posterior probability (w) of it being associated with each target is generated, and these w's must sum to unity. With no modification the result is, usually, a lost track.

Some provision must therefore be made. The most straightforward approach is to allow for an $(M+1)^{st}$ "dummy" target, from which all false detections are assumed to arise. The dynamics of this target are not important, and it can be assumed stationary; what is important is that it have a sufficiently large measurement covariance so as to have its associated measurements appear approximately uniformly in space over the region of interest.

2.3. Homothetic Gating

In the PMHT the concept of a gate is murky, and deserves some explanation. Equations (6) describe the synthetic measurements and covariances to be applied to a Kalman smoother, and it is easy to see that the former is a posterior-probabilistic centroid of measurements, which makes intuitive sense. The latter indicates the weighting factor for this synthetic measurement: if $\sum_{r=1}^{n_t} w_{sr}(t)$ is small, this must have arisen because no measurements at time t seem likely to have come from target s, and hence the synthetic measurement covariance $\tilde{\mathbf{R}}_s(t)$ supplied to the Kalman smoother is large, reflecting the PMHT's lack of confidence in $\tilde{\mathbf{z}}_s(t)$.

While this is sensible and intuitive, it is not a gate. As with the JPDAF and MHT algorithms, a gate is not necessary in the PMHT, but there is still gating action as implied by the posterior association probabilities (the w's for the PMHT). From equation (5) it is clear that these are determined by the modeled measurement covariances (the $\mathbf{R}$'s) and not the synthetic ones (not the $\tilde{\mathbf{R}}$'s) – and this is the problem: the $\mathbf{R}$'s are fixed. Without dynamically-sized gating (which the JPDAF and MHT employ) targets are easily lost, and performance is poor.

Our solution is to change the model from that of equation (1) to

$$\begin{aligned} \mathbf{x}_s(t+1) &= \mathbf{F}_s(t)\mathbf{x}_s(t) + \mathbf{G}_s(t)\mathbf{u}_s(t) + \mathbf{v}_s(t) \\ \mathbf{y}_{sp}(t) &= \mathbf{H}_s(t)\mathbf{x}_s(t) + \mathbf{w}_{sp}(t) \end{aligned} \tag{7}$$

where $p = 1, 2, \ldots, P$. All quantities behave as before; the difference is that for each target s there are P submodels, the p^{th} of which has measurement noise $\{\mathbf{w}_{sp}(t)\}$ and for which $E\{\mathbf{w}_{sp}(t)\mathbf{w}_{sp}^T(t)\} = \kappa_{sp}\mathbf{R}_s(t)$. Each submodel of the same s-index has the same trajectory $\{\mathbf{x}_s(t)\}$, but the measurements obtained are different, and have multiplicatively-different covariances. Intuitively each target carries along with it P concentric and fixed gates, presumably some small (small κ_{sp}) and some large (large κ_{sp}). The PMHT "prefers" to find its measurements, if possible, within the smallest gate; but if no such detection is available it turns its attention successively to its larger gates.

This "homothetic" (meaning that groups of P targets which *have the same trajectory*) PMHT operates exactly as described before, except that we have

$$\begin{aligned} \tilde{\mathbf{z}}_s(t) &\equiv \frac{\sum_{p=1}^{P}\sum_{r=1}^{n_t} w_{(Ps+p)r}^n(t)\mathbf{z}_r(t)/\kappa_{(Ps+p)r}}{\sum_{p=1}^{P}\sum_{r=1}^{n_t} w_{(Ps+p)r}^n(t)/\kappa_{(Ps+p)r}} \\ \tilde{\mathbf{R}}_s(t) &\equiv \frac{\mathbf{R}_s(t)}{\sum_{p=1}^{P}\sum_{r=1}^{n_t} w_{(Ps+p)r}^n(t)/\kappa_{(Ps+p)r}} \end{aligned} \tag{8}$$

instead of (6). As promised, a target's homothetic modes with larger covariance multipliers (κ's) produce lower gain responses. Suggested prior probabilities (i.e. π's) are:

$$\pi_s = \begin{cases} \frac{T}{P\sum_{t=1}^{T} n_t} & s = 1, 2, \ldots, PM \\ 1 - \frac{MT}{\sum_{t=1}^{T} n_t} & s = PM + 1 \end{cases} \tag{9}$$

3. RESULTS

Our model is two-dimensional and kinematic, with position measurements, and contains two targets; that is, with reference to equation (1) we have

$$\begin{aligned} \mathbf{F}_s(t) &= \begin{Bmatrix} 1 & \Delta t & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & \Delta t \\ 0 & 0 & 0 & 1 \end{Bmatrix} \\ \mathbf{H}_s(t) &= \begin{Bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{Bmatrix} \\ \mathbf{Q}_s(t) &= \begin{Bmatrix} \Delta t^3/3 & \Delta t^2/2 & 0 & 0 \\ \Delta t^2/2 & \Delta t & 0 & 0 \\ 0 & 0 & \Delta t^3/3 & \Delta t^2/2 \\ 0 & 0 & \Delta t^2/2 & \Delta t \end{Bmatrix} \frac{1}{100} \\ \mathbf{R}_s(t) &= \begin{Bmatrix} 25 & 0 & 0 & 0 \\ 0 & 25 & 0 & 0 \\ 0 & 0 & 25 & 0 \\ 0 & 0 & 0 & 25 \end{Bmatrix} \end{aligned} \tag{10}$$

for $s = 1, 2$ and all $t = 1, 2, \ldots, T$, with $T = 100$ snapshots of measurements. The known input $\mathbf{u}_s(t)$ is taken as zero. The sampling interval is $\Delta t = 3$ seconds.

In all cases the trajectories used are constant velocity and parallel. (These are the true trajectories; for tracking a positive-definite $\mathbf{Q}$ is necessary, and the small one of (10) is not inconsistent with a constant-velocity target.) The probability of detection P_d is assumed to be the same for both targets, but its value will depend on the simulation. False alarms are Poisson-distributed, with spatial density λ whose value will depend on the simulation. We specify that each target can produce no more than one detection at any given time; that is, the model used is that of (4) which we have called "realistic", and is that upon which the MHT and JPDAF are based, rather than that giving rise to the PMHT. For both schemes two-point initialization is used.

A typical snapshot and its associated tracking behavior are shown in figure 1. At the current snapshot the detection from the lower target appears to have been missed – the synthetic measurement $\tilde{\mathbf{z}}$ is found near some false alarms.

In figure 2 we show the result of varying the κ's. It appears that $\kappa_{sp} = \{1, 4, 16\}$ is a reasonable choice, at least in the cases we have examined. We suspect that such a choice is in general reasonable; however, this may be a matter of tuning to a particular application.

We compare the PMHT to the JPDAF via 100 Monte Carlo runs in our constant-velocity parallel-track situation. The main metric is percentage of tracks lost, and the results are given in tables 1 and 2. We additionally show in figure 3 the mean square estimation

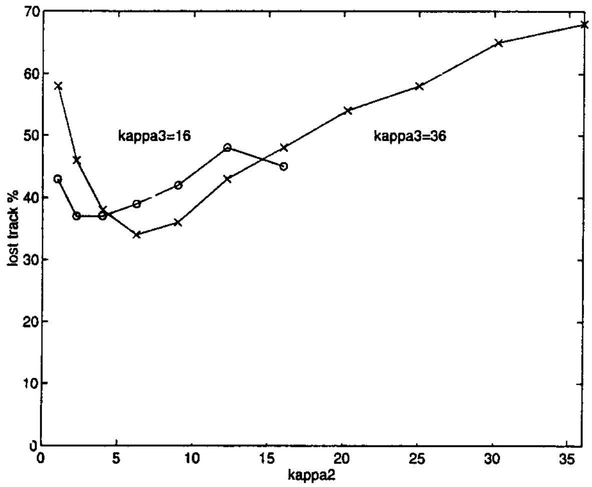

Figure 2: The percentage of lost tracks as a function of the κ's. These specify the ratios of the variances of the homothetic sub-models. Many such simulations have been run, and these results are typical: it appears to be a good choice to use $\kappa_2 = 4$, and a considerably larger κ_3 whose exact value is not important.

error (MSE) from those tracks which are not lost. Our conclusion is that for moderate P_d and λ, the improvement of the PMHT over the JPDAF is impressive.

λ	10^{-4}	$10^{-4.5}$	10^{-5}	$10^{-5.5}$	10^{-6}
$P_d = 0.4$	92	73	68	28	18
$P_d = 0.5$	74	50	21	16	11
$P_d = 0.6$	48	28	9	9	2
$P_d = 0.7$	34	11	6	8	5
$P_d = 0.8$	17	5	4	0	1
$P_d = 0.9$	12	0	1	0	1
$P_d = 1.0$	3	0	0	0	0

Table 1: The percentage of lost tracks for two parallel and constant-velocity target models, from simulation, PMHT algorithm. The average number of false-alarms per square meter is λ.

4. SUMMARY

The purpose of this paper has been two-fold: to introduce some needed modifications to the original PMHT design, and to compare the PMHT to the JPDAF. With respect to the first, we have explained the use of a high measurement variance "dummy" target to absorb false alarms; and we have elaborated the "homothetic" multi-model concept as a means to obtain variable gating behavior. As for the second, we have shown via simulation that in both terms of lost tracks and of MSE, the PMHT can offer a satisfying improvement over the JPDAF in a practical two-target scenario.

λ	10^{-4}	$10^{-4.5}$	10^{-5}	$10^{-5.5}$	10^{-6}
$P_d = 0.4$	98	82	75	64	74
$P_d = 0.5$	87	73	55	59	60
$P_d = 0.6$	69	52	36	56	39
$P_d = 0.7$	57	39	33	25	43
$P_d = 0.8$	37	23	30	26	18
$P_d = 0.9$	25	19	17	16	21
$P_d = 1.0$	0	0	0	6	5

Table 2: The percentage of lost tracks for two parallel and constant-velocity target models, from simulation, JPDAF algorithm. The average number of false-alarms per square meter is λ.

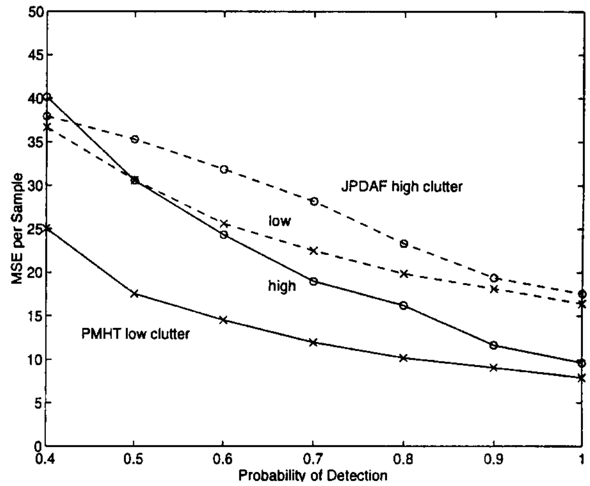

Figure 3: The mean-square estimation error (MSE) for those tracks which in tables 1 and 2 are not lost. "High" and "low" clutter correspond to the left- and right-most columns in those tables, respectively.

5. REFERENCES

[1] R.L. Streit and T.E. Luginbuhl, "Probabilistic Multi-Hypothesis Tracking", submitted to *IEEE Transactions on Automatic Control.*

[2] P. Willett, C. Rago, and R. Streit, "The Probabilistic Multi-Hypothesis Tracking Algorithm: Implementation Issues and Results", *University of Connecticut Technical Report TR 95-1*, January 1995.

[3] Y. Bar-Shalom, X.R. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*, YBS Publishing, 1995.