

Cardinality Tracking and the Equivalence-Class MHT

Stefano Coraluppi and Craig Carthel

Compunetix Inc.

2420 Mosside Boulevard, Monroeville PA 15146 USA

E-mail: stefano.coraluppi@compunetix.com, craig.carthel@compunetix.com

ABSTRACT

This manuscript first studies cardinality tracking, a special case of the general multi-target tracking problem for which measurements do not provide any target state information. That is, each scan only provides information as to how many targets are present. We address the problem with a modified form of the multiple-hypothesis tracking formalism using sets of hypotheses. Optimal track extraction may be performed due to the special structure of the problem. Hypothesis aggregation turns out to be essential in determining reasonable solutions to the cardinality tracking problem. Next, we consider the multi-target filtering problem for closely spaced targets, for which most available solutions incur track coalescence or track repulsion. The approach used for cardinality tracking – hypothesis aggregation – provides motivation for an approach to multi-target filtering that we call equivalence-class MHT. The idea is to choose a solution that, while suboptimal, is similar to many other solutions that, collectively, maximize the global posterior probability. That is, rather than seeking to identify a MAP solution, we seek a MAP set of solutions. We demonstrate improved performance of this approach over classical PDAF and MHT solutions in two-target closely-spaced scenarios.

1 INTRODUCTION

Track-oriented multiple hypothesis tracking (MHT) is well-established as a paradigm for multi-sensor multi-target tracking. The fundamental approach includes many variants. Hypothesis-oriented MHT was first proposed by Reid [1]. The initial integer-programming formulation of the problem is due to Morefield [2]. The hybrid-state decomposition that allows for computationally efficient track-oriented MHT is due to Kurien [3]. An efficient solution to the optimization problem required for *n-scan* hypothesis pruning via Lagrangian relaxation is due to Poore and Rijavec [4]. Linear-programming based relaxation approaches to the optimization problem were proposed independently by Coraluppi et al [5] and by Storms and Spieksma [6].

In [7], we proposed a modified MHT approach that is more effective at suppressing spurious tracks. The approach is motivated by optimality results that one can prove in the limit of large measurement errors. In this limit, we are faced with a cardinality tracking problem in which we seek only to identify how many targets are present. Unfortunately, cardinality tracking highlights a weakness in the MHT formalism, in particular its selection of a *maximum a posteriori* (MAP) tracking solution. This solution may not perform well with respect to metrics of interest.

As an example, consider a biased coin with $p(\text{heads})=0.6$. Toss the coin 100 times. Assume there is no measured data. What is the *maximum a posteriori* estimate for the number of heads? The answer, 100 heads, provides a very poor estimate of the actual number of heads in the realization (close to 60).

Indeed, in cardinality tracking, given a sufficiently high false alarm rate the *null* solution (no tracks) is the

MAP solution. To circumvent this difficulty, in Section 4 we introduce an approach whereby all indistinguishable global hypotheses are treated together. As such, the revised MHT formalism identifies the likeliest *class* of hypotheses. We call this the *cardinality tracker*. A similar paradigm turns out to be extremely useful in addressing multi-target filtering problems for closely-spaced targets. We describe the problem in Section 5 and propose an *equivalence-class MHT* (ECMHT) solution in Section 6. Performance results are promising; the ECMHT appears to suffer neither track coalescence nor track repulsion, and provides more accurate tracking estimates than the well-known classical approaches. Concluding remarks are in Section 7.

2 THE MULTI-TARGET TRACKING PROBLEM

We start by stating the general multi-target tracking (MTT) problem of interest. We are given a continuous-time Poisson birth and death process with birth and death parameters λ and λ_χ , respectively. We consider a discrete-time sequence (t_0, t_1, t_2, \dots) with inter-sampling times defined by $\Delta t_k = t_{k+1} - t_k$, $k \geq 0$. Correspondingly, we have the following discrete-time birth and death statistics; note that these do not account for birth-death events between sampling times, as these remain invisible to the sampling process.

$$\lambda_b(k) = \frac{\lambda}{\lambda_\chi} (1 - \exp(-\lambda_\chi \Delta t_{k-1})), \quad k \geq 1, \tag{2.1}$$

$$p_\chi(k) = 1 - \exp(-\lambda_\chi \Delta t_{k-1}), \quad k \geq 1. \tag{2.2}$$

We assume there are no targets at time t_0 . When the inter-sampling time is constant we may omit the time index and write λ_b and p_χ .

For each target, we assume the initial state follows the probability density function $f_b(\cdot)$ defined over target state space. We assume that a time-invariant continuous-time target motion model and discrete-time measurement model at times $t^k = (t_1, t_2, \dots)$ are known. While the linear-Gaussian model is common, it is not required. Let $\hat{f}_b(\cdot)$ denote the target birth distribution in measurement space. Further, let $f(y_k | Y^{k-1})$ denote the probability density for the sensor measurement on a target at time t_k given a preceding sequence of measurements. Let p_d denote the target detection probability at all scan times, and let λ_{fa} be the mean of the Poisson-distributed number of false alarms at each scan time; these are distributed in measurement space according to $f_{fa}(\cdot)$.

Let X^k denote a multi-target realization based on the birth-death statistics and target motion evolution discussed above. That is, X^k is a set of target trajectories defined on (t_1, t_2, \dots) . The set X^k is not to be confused with X^k , the latter being the usual notation for a single-target trajectory given by $X^k = (X_1, \dots, X_k)$. Each element of X^k is a distinct target with a birth time, a state trajectory, and (possibly) a death time. Based on the sensor modeling discussed above, we observe a sequence of sets of contacts $Z^k = (Z_1, \dots, Z_k)$ at a sequence of times t^k . We denote the contacts in each set with a double subscript as follows: $Z_i = \{Z_{i1}, \dots, Z_{i|Z_i|}\}$.

Arguably, a solution to the multi-target tracking problem is given by the probability distribution

$p(\mathbf{X}^k | \mathbf{Z}^k)$ over all sets of target trajectories conditioned on the sequence of sets of measurements. This solution is problematic for at least two reasons. The first is that, even for small problems, $p(\mathbf{X}^k | \mathbf{Z}^k)$ is extremely complicated to compute. A second difficulty is that, generally, one is interested in a specific tracking solution (i.e. a deterministic set of tracks) rather than a probability distribution over tracking solutions. Indeed, most tracking paradigms documented in the literature seek an estimate $\hat{\mathbf{X}}^k$ that is close to \mathbf{X}^k . Since $\hat{\mathbf{X}}^k$ and \mathbf{X}^k are complex objects, the issue of what are appropriate tracking metrics to capture the notion of close is a topic of interest in its own right. Generally, we find that several metrics are required to quantify the quality of a tracking solution.

3 MULTIPLE-HYPOTHESIS TRACKING

In multiple-hypothesis tracking (MHT), one seeks the *maximum a posteriori* (MAP) tracking solution. That is, we seek the following:

$$\hat{\mathbf{X}}^k = \arg \max p(\mathbf{X}^k | \mathbf{Z}^k). \quad (3.1)$$

The estimate $\hat{\mathbf{X}}^k$ given by (3.1) is difficult to compute. Thus, we proceed with a hybrid-state formulation that will aid us in determining an approximation to $\hat{\mathbf{X}}^k$. We introduce the discrete state variable \mathcal{Q}^k ; this variable identifies all birth, death and measurement-association events up to time t_k . Conditioned on \mathcal{Q}^k , the general tracking problem simplifies to a set of nonlinear filtering problems. One could recover the original tracking formulation by a weighed combination over all discrete-state realizations, as stated with eqn. (3.2). Rather, we see the MAP estimate for \mathcal{Q}^k , given by eqn. (3.3). Next, we rely on the approximation given by eqn. (3.4) to replace eqn. (3.1) by the much simpler eqn. (3.5).

$$p(\mathbf{X}^k | \mathbf{Z}^k) = \sum_{\mathcal{Q}^k} p(\mathbf{X}^k | \mathbf{Z}^k, \mathcal{Q}^k) p(\mathcal{Q}^k | \mathbf{Z}^k). \quad (3.2)$$

$$q^k = \arg \max p(\mathcal{Q}^k | \mathbf{Z}^k). \quad (3.3)$$

$$p(\mathbf{X}^k | \mathbf{Z}^k) \approx p(\mathbf{X}^k | \mathbf{Z}^k, q^k). \quad (3.4)$$

$$\hat{\mathbf{X}}^k = \arg \max p(\mathbf{X}^k | \mathbf{Z}^k, q^k). \quad (3.5)$$

To conclude, since solving (3.1) is intractable, the MHT approach can be understood as the selection of a single high-probability global hypothesis (eqn. 3.3), followed by the determination of the corresponding tracking solution (eqn. 3.5).

In order for the MHT approach to be both computationally viable and utilizable in real-time, we do not have the luxury to wait for all the data \mathbf{Z}^k to be received and for all possible discrete-state hypotheses to be considered as prescribed by eqns. (3.3-3.5). Thus, one generally introduces a number of simplifying steps. One approach is to select the global hypothesis q^k incrementally, with a fixed delay relative to the incoming sequence of measurement sets. This is often referred to as n_{scan} pruning. Also, not all possible discrete states are considered: unlikely measurement associations are not considered, track hypotheses are terminated when sufficiently degraded (e.g. after several missed detections) with no earlier track-termination hypothesis spawn. (These comments will become clearer with the example at the end of this

section.)

Even with incremental, real-time (albeit fixed-delay) determination of the global hypothesis q^k , generally it is computationally infeasible to consider explicit enumeration of global hypotheses. Fortunately, it is possible to cast the determination of the optimal global hypothesis as an optimization problem in which feasible global hypotheses are implicitly defined by constraints on the set of local or track hypotheses that compose the global hypotheses. A local (or track) hypothesis identifies the birth event, measurement associations, and (possibly) death event for a single target. This approach to MHT, which does not require explicit enumeration of global hypotheses, is often referred to as *track-oriented MHT* (TO-MHT) or *multi-dimensional assignment* (MDA).

A key enabler for TO-MHT is the following recursive expression that relates the posterior probability of global hypothesis $Q^k = (Q^{k-1}, Q_k)$, to the posterior probability of its constituent hypothesis Q^{k-1} [3] and to local track hypotheses associated with track terminations, track missed detections, track updates, and track births.

$$p(Q^k | Z^k) = p_z^\chi ((1 - p_\chi)(1 - p_d))^{\tau - \chi - d} \cdot \prod_{j \in J_d} \left[\frac{(1 - p_\chi) p_d f(z_j | Z^{k-1}, Q^k)}{\lambda_{fa} f_{fa}(z_j | Z^{k-1}, Q^k)} \right] \cdot \prod_{j \in J_b} \left[\frac{p_d \lambda_b f_b(z_j | Z^{k-1}, Q^k)}{\lambda_{fa} f_{fa}(z_j | Z^{k-1}, Q^k)} \right] \frac{p(Q^{k-1} | Z^{k-1})}{\bar{c}_k}, \quad (3.6)$$

$$c_k = f(Z_k | Z^{k-1}) = \sum_{Q^k} f(Z_k | Z^{k-1}, Q^k) p(Q^k | Z^{k-1}), \quad (3.7)$$

$$\bar{c}_k = \frac{c_k}{\left\{ \frac{\exp(-p_d \lambda_b - \lambda_{fa})}{r!} \lambda_{fa}^r \right\} \prod_{j \in J_d \cup J_b \cup J_{fa}} f_{fa}(z_j | Z^{k-1}, Q^k)}. \quad (3.8)$$

In equations (3.6-3.8), c_k is a normalizing constant independent of Q_k ; τ is the number of targets at t_{k-1} ; d is the number of detections of existing targets at t_k ; χ is the number of target terminations at t_k ; b is the number of detected target births at t_k ; $r = d + b + f$ is the number of measurements at t_k (with f being the number of false alarms); $J_d \subset Z_k$, $J_b \subset Z_k$, $J_{fa} \subset Z_k$ is the set of measurements that are from existing targets, from new targets, and false alarms, respectively; $z_j \in Z_k$ is a measurement at time t_k .

Let us consider a simple example and illustrate MHT processing at a high level. Assume $|Z_1| = 1, |Z_2| = 2, |Z_3| = 1$. There are many possible discrete states Q^k ; for each, there is a corresponding global hypothesis X^k . Figure 3.1 enumerates the full set of 40 possible discrete states (and, thus, 40 global hypotheses). In Figure 3.1, O denotes a track coast, i.e. track existence in the absence of a measurement update; \otimes denotes a track termination. Figure 3.2 illustrates the TO-MHT formulation that does not include global hypothesis enumeration; we see that there are 17 local track hypotheses. Assume that $q^3 = \{(Z_{11}, Z_{21}, O), (Z_{22}, O), (Z_{31})\}$. Figure 3.3 illustrates the result of local hypothesis pruning under sequential processing with a one-scan delay (i.e. $n_{scan} = 1$). Note that many competing global hypotheses survive. Indeed, in subsequent processing, q^3 may well be invalidated.

$$Q^3 \in \left\{ \begin{aligned} & \{(Z_{11}, Z_{21}, Z_{31}), (Z_{22}, O)\}, \{(Z_{11}, Z_{21}, Z_{31}), (Z_{22}, \otimes)\}, \\ & \{(Z_{11}, Z_{22}, Z_{31}), (Z_{21}, O)\}, \{(Z_{11}, Z_{22}, Z_{31}), (Z_{21}, \otimes)\}, \\ & \{(Z_{11}, Z_{21}, O), (Z_{22}, O), (Z_{31})\}, \{(Z_{11}, Z_{21}, O), (Z_{22}, \otimes), (Z_{31})\}, \{(Z_{11}, Z_{21}, \otimes), (Z_{22}, O), (Z_{31})\}, \{(Z_{11}, Z_{21}, \otimes), (Z_{22}, \otimes), (Z_{31})\}, \\ & \{(Z_{11}, Z_{22}, O), (Z_{21}, O), (Z_{31})\}, \{(Z_{11}, Z_{22}, O), (Z_{21}, \otimes), (Z_{31})\}, \{(Z_{11}, Z_{22}, \otimes), (Z_{21}, O), (Z_{31})\}, \{(Z_{11}, Z_{22}, \otimes), (Z_{21}, \otimes), (Z_{31})\}, \\ & \{(Z_{11}, O, Z_{31}), (Z_{21}, O), (Z_{22}, O)\}, \{(Z_{11}, O, Z_{31}), (Z_{21}, O), (Z_{22}, \otimes)\}, \{(Z_{11}, O, Z_{31}), (Z_{21}, \otimes), (Z_{22}, O)\}, \{(Z_{11}, O, Z_{31}), (Z_{21}, \otimes), (Z_{22}, \otimes)\}, \\ & \{(Z_{11}, O, O), (Z_{21}, Z_{31}), (Z_{22}, O)\}, \{(Z_{11}, O, \otimes), (Z_{21}, Z_{31}), (Z_{22}, O)\}, \{(Z_{11}, \otimes), (Z_{21}, Z_{31}), (Z_{22}, O)\}, \\ & \{(Z_{11}, O, O), (Z_{21}, Z_{31}), (Z_{22}, \otimes)\}, \{(Z_{11}, O, \otimes), (Z_{21}, Z_{31}), (Z_{22}, \otimes)\}, \{(Z_{11}, \otimes), (Z_{21}, Z_{31}), (Z_{22}, \otimes)\}, \\ & \{(Z_{11}, O, O), (Z_{22}, Z_{31}), (Z_{21}, O)\}, \{(Z_{11}, O, \otimes), (Z_{22}, Z_{31}), (Z_{21}, O)\}, \{(Z_{11}, \otimes), (Z_{22}, Z_{31}), (Z_{21}, O)\}, \\ & \{(Z_{11}, O, O), (Z_{22}, Z_{31}), (Z_{21}, \otimes)\}, \{(Z_{11}, O, \otimes), (Z_{22}, Z_{31}), (Z_{21}, \otimes)\}, \{(Z_{11}, \otimes), (Z_{22}, Z_{31}), (Z_{21}, \otimes)\}, \\ & \{(Z_{11}, O, O), (Z_{21}, O), (Z_{22}, O), (Z_{31})\}, \{(Z_{11}, O, \otimes), (Z_{21}, O), (Z_{22}, O), (Z_{31})\}, \{(Z_{11}, \otimes), (Z_{21}, O), (Z_{22}, O), (Z_{31})\}, \\ & \{(Z_{11}, O, O), (Z_{21}, O), (Z_{22}, \otimes), (Z_{31})\}, \{(Z_{11}, O, \otimes), (Z_{21}, O), (Z_{22}, \otimes), (Z_{31})\}, \{(Z_{11}, \otimes), (Z_{21}, O), (Z_{22}, \otimes), (Z_{31})\}, \\ & \{(Z_{11}, O, O), (Z_{21}, \otimes), (Z_{22}, O), (Z_{31})\}, \{(Z_{11}, O, \otimes), (Z_{21}, \otimes), (Z_{22}, O), (Z_{31})\}, \{(Z_{11}, \otimes), (Z_{21}, \otimes), (Z_{22}, O), (Z_{31})\}, \\ & \{(Z_{11}, O, O), (Z_{21}, \otimes), (Z_{22}, \otimes), (Z_{31})\}, \{(Z_{11}, O, \otimes), (Z_{21}, \otimes), (Z_{22}, \otimes), (Z_{31})\}, \{(Z_{11}, \otimes), (Z_{21}, \otimes), (Z_{22}, \otimes), (Z_{31})\}, \end{aligned} \right\}$$

Figure 3.1: Full enumeration of discrete state space (corresponding to 40 global hypotheses).

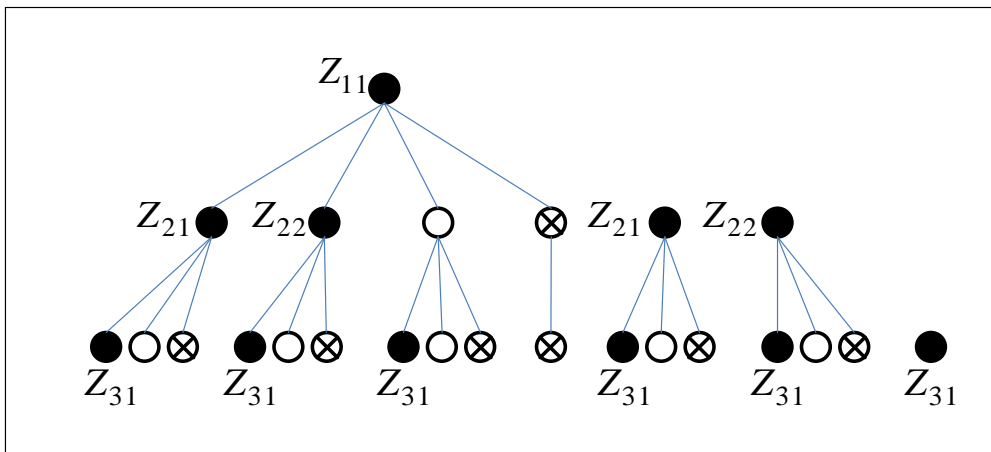


Figure 3.2: Full enumeration of track hypotheses under TO-MHT (17 local hypotheses).




Figure 3.3: TO-MHT with sequential hypothesis pruning (7 surviving local hypotheses).

This example serves to illustrate the fundamental difference between hypothesis-oriented MHT and TO-MHT. In hypothesis-oriented MHT, one enumerates global hypotheses and identifies the MAP hypothesis. In TO-MHT, one solves an optimization problem whose objective function is the global hypothesis posterior probability; however, the optimization methodology relies on local track hypotheses and appropriate constraints, and does not require enumeration of global hypotheses.

We have provided as well an illustration of local hypothesis reduction via sequential processing. As previously noted, many additional hypothesis-reduction techniques exist, including at the track-generation stage. For example, one might choose to spawn only a track coast and no track termination hypothesis on the first missed detection, and only a track termination hypothesis on the second consecutive mixed detection. Such rules must be chosen judiciously based on target and sensor modeling assumptions. Additionally, sufficiently-unlikely association hypotheses are generally discarded via measurement gating.

4 CARDINALITY TRACKING

In the case of cardinality-only information, the track-oriented MHT equation simplifies to the following:

$$p(Q^k | Z^k) = \frac{p_\chi^\chi (1 - p_\chi)^{\tau - \chi} (1 - p_d)^{\tau - \chi - d} p_d^{b+d} \lambda_b^b}{\lambda_{fa}^{d+b}} \frac{p(Q^{k-1} | Z^{k-1})}{\bar{c}_k}, \quad (4.1)$$

where for all scans λ_b is the mean for Poisson-distributed births, λ_{fa} is the mean for Poisson-distributed false returns, p_d is the object detection probability, p_χ is the object termination probability, $Q^k = (Q_1, \dots, Q_k)$ is the discrete state that accounts for target births and terminations as well as all data associations, and \bar{c}_k is a factor that is independent of Q^k . Again, we are interested in the likeliest Q^k given measurement data (that we denote by q^k), and then condition on this hypothesis q^k to estimate target trajectories. In cardinality tracking, each target trajectory is characterized solely by a track initiation time and a track termination time: there is no state trajectory.

The use of (4.1) for cardinality tracking is problematic. A first difficulty is a computational one: since all measurements gate with all tracks, there is a large number of track hypotheses. A more serious concern is a large number of comparably-scoring global hypotheses in large-scale surveillance applications, leading to the null solution (i.e. no tracks) as the optimal choice under (3.3). Indeed, for large surveillance problems, the posterior probability $p(Q^k | Z^k)$ will be very small for all non-trivial choice of Q^k (i.e. for all but the null solution). Indeed, unlike conventional tracking, measurements are much less informative and we lack kinematic filter residuals that lead to relatively large hypothesis scores for some association decisions and relatively small scores for others.

The difficulties associated with identifying a single global hypothesis point to a fundamental limitation in the MHT paradigm: the selection of a single global hypothesis. Conversely, it would be beneficial to identify a *set* of global hypotheses that are indistinguishable (due to measurement equivalence) and that provide significant probability mass. We do so next.

We wish to identify a most probably *set* of discrete states $\{Q^k\}$, where all member of the set are equivalent under measurement re-ordering. To do so, we must revisit the derivation of track-oriented MHT equations and introduce suitable modifications. Further, as with classical track-oriented MHT, we are interested in a recursive expression for $p(\{Q^k\} | Z^k)$ that lends itself to functional optimization without the need for

explicit enumeration of global hypotheses. We proceed by repeated use of Bayes' rule.

$$\begin{aligned}
 p(\{Q^k\} | Z^k) &= \frac{p(Z_k | Z^{k-1}, \{Q^k\})p(\{Q^k\} | Z^{k-1})}{c_k} \\
 &= \frac{p(Z_k | Z^{k-1}, \{Q^k\})p(\{Q_k\} | Z^{k-1}, \{Q^{k-1}\})p(\{Q^{k-1}\} | Z^{k-1})}{c_k},
 \end{aligned} \tag{4.2}$$

$$c_k = p(Z_k | Z^{k-1}) = \sum_{\{Q^k\}} p(Z_k | Z^{k-1}, \{Q^k\})p(\{Q^k\} | Z^{k-1}). \tag{4.3}$$

The recursive expression (4.2) involves two factors that we consider in turn, with the discrete state probability one first. It will be useful to introduce the aggregate variable ψ_k that accounts for the number of detections d for the τ existing tracks, the number of track deaths χ , the number of new tracks b , and the number of false returns $r - d - b$, where r is the number of contacts in the current scan.

$$p(\{Q_k\} | Z^{k-1}, \{Q^{k-1}\}) = p(\psi_k | Z^{k-1}, \{Q^{k-1}\})p(\{Q_k\} | Z^{k-1}, \{Q^{k-1}\}, \psi_k), \tag{4.4}$$

$$\begin{aligned}
 p(\psi_k | Z^{k-1}, \{Q^{k-1}\}) &= \left\{ \binom{\tau}{\chi} p_\chi^\chi (1-p_\chi)^{\tau-\chi} \right\} \cdot \left\{ \binom{\tau-\chi}{d} p_d^d (1-p_d)^{\tau-\chi-d} \right\} \\
 &\quad \cdot \left\{ \frac{\exp(-p_d \lambda_b) p_d^b \lambda_b^b}{b!} \right\} \cdot \left\{ \frac{\exp(-\lambda_{fa}) \lambda_{fa}^{r-d-b}}{(r-d-b)!} \right\},
 \end{aligned} \tag{4.5}$$

$$p(\{Q_k\} | Z^{k-1}, \{Q^{k-1}\}, \psi_k) = \frac{b!}{\binom{\tau}{\chi} \binom{\tau-\chi}{d}}. \tag{4.6}$$

The key difference with respect to standard track-oriented MHT is in (4.6), since we must not account for differences in which measurements are taken to be track updates, which are taken to be track births, and how measurements are assigned to tracks. Substituting (4.5-4.6) into (4.4) and simplifying yields the following.

$$\begin{aligned}
 p(\{Q_k\} | Z^{k-1}, \{Q^{k-1}\}) &= \left\{ \frac{\exp(-p_d \lambda_b - \lambda_{fa}) \lambda_{fa}^r}{r!} \right\} p_\chi^\chi ((1-p_\chi)(1-p_d))^{\tau-\chi-d} \\
 &\quad \cdot \left(\prod_{i=1}^d (r+1-i) \right) \left(\frac{(1-p_\chi)p_d}{\lambda_{fa}} \right)^d \cdot \left(\prod_{i=1}^b (r-d+1-i) \right) \left(\frac{p_d \lambda_b}{\lambda_{fa}} \right)^b.
 \end{aligned} \tag{4.7}$$

Substituting (4.7) into (4.2) and further simplification leads to (4.8).

$$p(\{Q^k\} | Z^k) = \frac{p_\chi^\chi (1-p_\chi)^{\tau-\chi} (1-p_d)^{\tau-\chi-d} p_d^{b+d} \lambda_b^b}{\lambda_{fa}^{d+b}} \cdot \frac{p(\{Q^{k-1}\} | Z^{k-1})}{\bar{c}_k} \left(\prod_{i=1}^d (r+1-i) \right) \left(\prod_{i=1}^b (r-d+1-i) \right), \tag{4.8}$$

$$\bar{c}_k = \frac{c_k}{\left\{ \frac{\exp(-p_d \lambda_b - \lambda_{fa})}{r!} \lambda_{fa}^r \right\}} \tag{4.9}$$

Note the multiplicative weights in (4.8) that are not present in (3.6).

Cardinality tracking expressed via equivalence classes leads both to computational efficiency (there are much fewer equivalence classes over global hypothesis than there are global hypotheses) and to a well-posed formulation whereby a MAP class of solutions is of interest. Further, the weights in (18) imply *structural results* regarding the form of the optimal equivalence class and optimal track-extraction rules.

In particular, an optimal tracking solution can be extracted from a so-called *Tetris structure*, where measurements are arranged as shown in Figure 4.1. The advantages of this structure are significant: track hypothesis trees are not required, and optimal batch extraction of tracks may be performed, as explained below.

		1																	
		1																	
		1	1																
		1	1	1										1					1
	1	1	1	1		1	1	1				1	1	1		1		1	
1	1	1	1	1	1	1	1	1	1			1	1	1	1	1		1	1

Figure 4.1: Each column corresponds to a scan of data. In each column, measurements are stacked on top of one another. Tracks are extracted by considering rows separately.

First, for each row z in the Tetris structure, we consider separately sequences of measurements divided by more than K misses, where $K = 0$ if $p_d = 1$; otherwise, K is given by the largest non-negative integer that satisfies:

$$K < \frac{\log\left(\frac{\lambda_b p_\chi}{1 - p_\chi}\right)}{\log((1 - p_\chi)(1 - p_d))} \tag{4.10}$$

For a sequence of length N that includes M measurements, the track score is given by (4.11). Note that, for tracks that start at the beginning of the data stream, the steady-state number of targets, $\frac{\lambda_b}{p_\chi}$, applies in lieu of the birth rate λ_b ; likewise, for tracks that reach the end of the data stream, the factor p_χ is not to be included in (4.11).

$$f(M, N, z, Z^N) = \lambda_b p_\chi (1 - p_\chi)^{N-1} \left(\frac{p_d}{\lambda_{fa}}\right)^M \cdot (1 - p_d)^{N-M} \prod_{\substack{i=1, \dots, N, \\ |Z_i| \neq 0}} (|Z_i - z + 1|) \tag{4.11}$$

The value K in (4.10) can be determined a priori, independent of the data realization, while the track extraction methodology depends both on row number and the input data cardinality sequence. Measurement sequences are only extracted if they contribute to the overall posterior probability of the hypothesis equivalence class. Multiple track segments are only extracted as a single track when it is

advantageous, again in terms of overall posterior probability of the hypothesis equivalence class. The extraction procedure below is applied separately to each row of data, $z \in \{1, \dots, \max\{Z_i\}\}$.

It is possible to introduce further hypothesis aggregation in cardinality tracking, yielding improved performance. Details may be found in [8]. Further, one may reformulate the cardinality tracking problem by relaxing the requirement for integer numbers of tracks and using a simple Kalman filtering approach (again, details are in [8]).

An example. Figure 4.2 illustrates a stochastic realization of ground truth, cardinality measurements, and classical MHT and cardinality tracking solutions, the latter based on the hypothesis set formulation developed here. The parameters for the realization are as follows: number of scans of data ($N = 250$), probability of death ($p_z = 0.01$), birth rate ($\lambda_b = 1$), false alarm rate per scan ($\lambda_{fa} = 10$). The classical MHT returns no tracks, thus exhibiting an extremely large track cardinality error. The cardinality tracker tends to identify lengthy tracks. As such, it tends to underestimate the true number of objects, though it vastly outperforms the classical MHT solution.

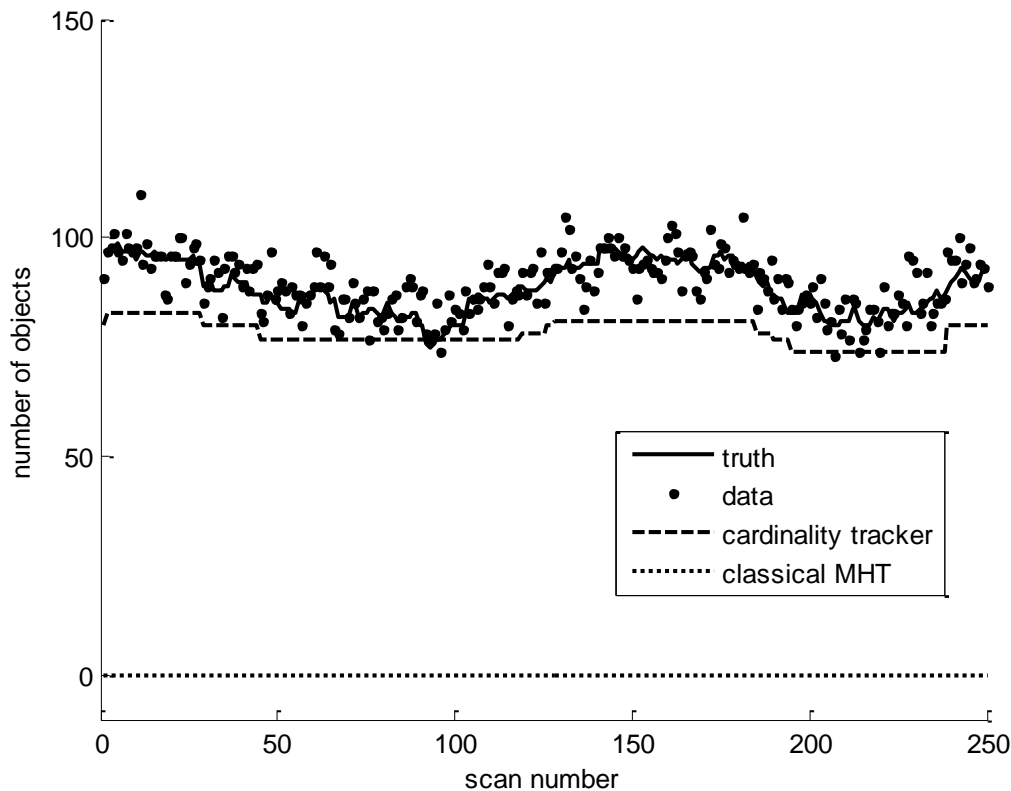


Figure 4.2: A realization of ground truth, cardinality data, and estimation sequences for cardinality tracking (with hypothesis set aggregation) and classical MHT.

5 THE MULTI-TARGET FILTERING PROBLEM

We consider now the impact of neighboring targets on tracking performance. We focus on localization accuracy for the case of two closely spaced targets. We disregard track management issues, and will consider several schemes for multi-target filtering. Here, the multi-target filtering problem assumes no

missed detections, no false alarms, and no measurement gating.

We will assume target motion according to a one-dimensional stationary *Ornstein-Uhlenbeck* (OU) motion model [9]. The noise process is independent and identically distributed. Time-discretization of the continuous-time process yields the following:

$$X_{k+1} = A_k X_k + w_k, \quad X_0 \sim N\left(0, \frac{q}{2\gamma}\right), \quad (5.1)$$

$$A_k = \exp(-\gamma\Delta t_k), \quad \gamma > 0, \quad \Delta t_k = t_{k+1} - t_k, \quad (5.2)$$

$$E[w_k] = 0, \quad E[w_k^2] = Q_k = \frac{1 - \exp(-2\gamma\Delta t_k)}{2\gamma} q. \quad (5.3)$$

The sensor measurement model is the following:

$$Z_k = C_k X_k + v_k, \quad E[v_k] = 0, \quad E[v_k^2] = R_k. \quad (5.4)$$

Overall, we are given a sequence of sets of measurements $(\{Z_{11}, Z_{12}\}, \{Z_{21}, Z_{22}\}, \{Z_{31}, Z_{32}\}, \dots)$, where for all k we have $Z_{k1} \leq Z_{k2}$. For all the filtering solutions we consider, the filter initialization and prediction equations are given by the following:

$$X_1(1|1) = Z_{11}, \quad (5.5)$$

$$X_2(1|1) = Z_{12}, \quad (5.6)$$

$$P_j(1|1) = R_1, \quad j = 1, 2, \quad (5.7)$$

$$X_j(k+1|k) = A_k X_j(k|k), \quad j = 1, 2, \quad (5.8)$$

$$P_j(k+1|k) = A_k P_j(k|k) A_k' + Q_k, \quad j = 1, 2, \quad (5.9)$$

$$S_j(k+1) = C_{k+1} P_j(k+1|k) C_{k+1}' + R_{k+1}, \quad j = 1, 2, \quad (5.10)$$

$$L_j(k+1) = P_j(k+1|k) C_{k+1}' S_j^{-1}(k+1), \quad j = 1, 2. \quad (5.11)$$

The update equations for the *probabilistic data association filter* (PDAF) [10] take the following form.

$$\beta_{ij}(k+1|k) = c_j \exp\left(-\frac{(Z_{ij} - C_k X_j(k+1|k))^2}{2S_j(k+1)}\right), \quad i = 1, 2, \quad j = 1, 2, \quad (5.12)$$

$$v_{ij}(k+1) = Z_{ij} - C_{k+1} X_j(k+1|k), \quad i = 1, 2, \quad j = 1, 2, \quad (5.13)$$

$$v_j(k+1) = \sum_i \beta_{ij} v_{ij}(k+1), \quad j = 1, 2, \quad (5.14)$$

$$X_j(k+1|k+1) = X_j(k+1|k) + L_j(k+1)v_j(k+1), j = 1, 2, \quad (5.15)$$

$$P_j(k+1|k+1) = (I - L_j(k+1)C_{k+1})P_j(k+1|k) + \tilde{P}_j(k+1), j = 1, 2, \quad (5.16)$$

$$\tilde{P}_j(k+1) = L_j(k+1) \left(\sum_i \beta_{ij} v_{ij}^2(k+1) - v_j^2(k+1) \right) L_j(k+1), j = 1, 2, \quad (5.17)$$

with the normalization term given by $c_j = 1 / \sum_i \beta_{ij}(k+1|k)$.

The update equations for the *global nearest-neighbor filter* (GNNF) [10] take the following form.

$$v_1(k+1) = Z_{(k+1)1} - X_1(k+1|k), \quad (5.18)$$

$$v_2(k+1) = Z_{(k+1)2} - X_2(k+1|k), \quad (5.19)$$

$$X_j(k+1|k+1) = X_j(k+1|k) + L_j(k+1)v_j(k+1), j = 1, 2, \quad (5.20)$$

$$P(k+1|k+1) = (I - L(k+1)C_{k+1})P(k+1|k). \quad (5.21)$$

Note that, unlike the PDAF, filter accuracy in the GNNF is data independent. Further, for this simplified problem, the GNNF and the *multiple-hypothesis tracker* (MHT) with arbitrary hypothesis tree depth are equivalent.

It has been shown that the PDAF exhibits track coalescence [11-12], while the GNNF and MHT exhibit track repulsion [13-15]. Several approaches have been investigated to mitigate the impact of track coalescence in soft data association filters, see e.g. [16]; an interesting, recently-introduced approach performs quite well at the cost of track labeling information [17]. The track repulsion effect in hard data association approaches can be mitigated through a multi-stage processing approach [14].

Assume that we have a field of k synchronized sensors. All sensors are characterized by (5.4), with no missed detections and no false alarms. We wish to determine an alternative hard data association, motivated by the equivalence-set approach described in the context of cardinality tracking. Specifically, we consider a set of global hypotheses that share the same number of measurement swaps. There are many ways to choose for which sensor to apply the swaps; this leads to (5.22), where the global hypothesis on the right hand side is chosen arbitrarily among all those with i swaps, and is denoted by Q_i^k .

$$p(\{Q_i^k\} | Z^k) \approx \binom{k}{i} p(Q_i^k | Z^k). \quad (5.22)$$

Interestingly, although the *maximum a posteriori* (MAP) solution involves no measurement swaps, the same is not true when reasoning over equivalence classes of hypotheses. We are interested to examine whether such an approach provides improved filtering accuracy in dense-target settings.

As a first analysis step, we consider a single set of (synchronous) sensor measurements. We merge two sets of k measurements into a single equivalent measurement, either according to the GNNF paradigm (no swaps), or by randomly introducing i measurement swaps before merging. Next, we examine the measurement error of the equivalent measurements, as a function of the number of swaps.

The results are shown in Figure 5.1, for a 10m displacement between targets, 30 sensors, and 10m measurement error standard deviation. Results are based on 5000 Monte Carlo realizations. Interestingly, introducing some swaps (about 4) provides a modest reduction in equivalent measurement RMS error. Further, the optimal number of swaps obtained empirically is consistent with the posterior probability of the optimal hypothesis set, determined according to (5.22) with one randomly-selected element of the set, and shown in Figure 5.2. The optimal number of swaps changes as a function of the distance between targets. In Figure 5.3, we show that, as expected, the optimal swap rate is 0.5 for identically-located targets and goes to zero for sufficiently distant targets.

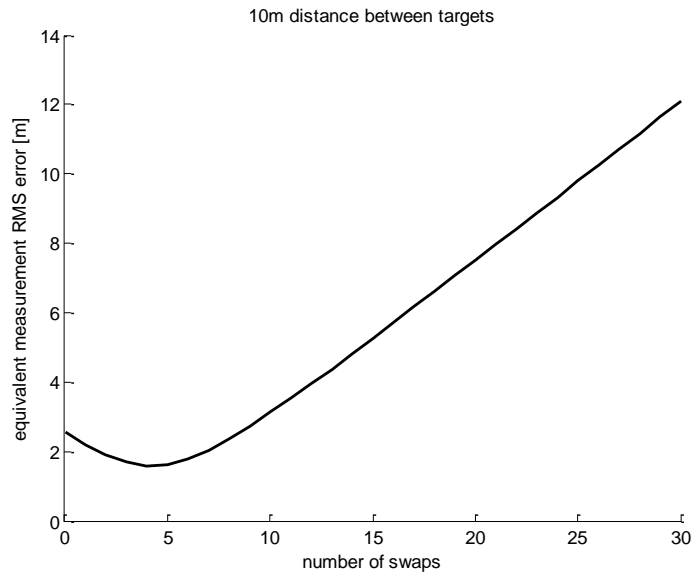


Figure 5.1: Equivalent measurement RMS error as a function of (random) measurement swaps.

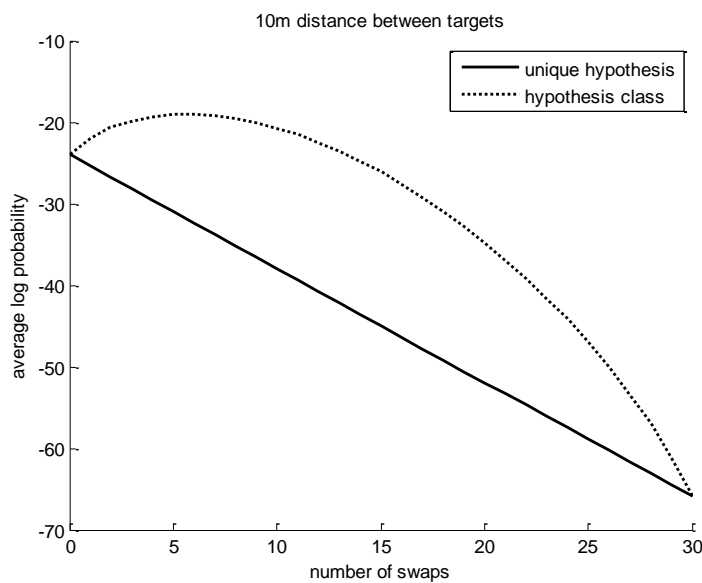


Figure 5.2: Single-hypothesis and hypothesis-class posterior probabilities.

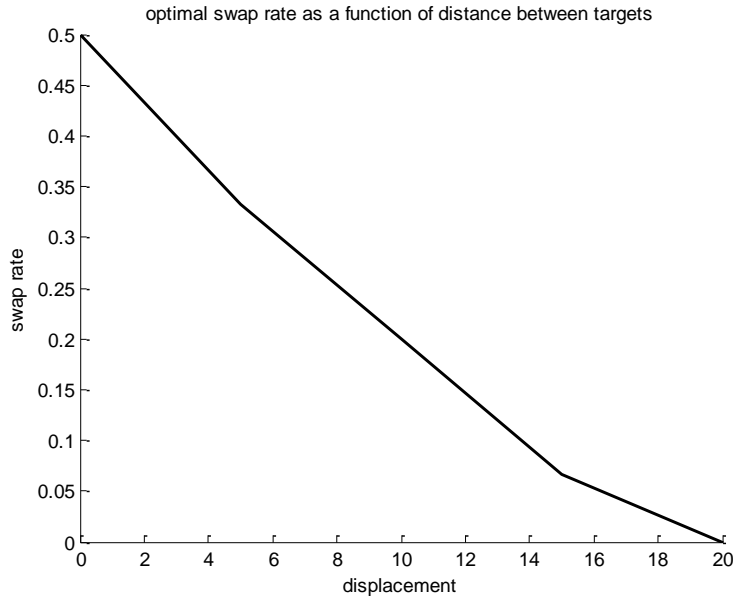


Figure 5.3: Optimal swap rate decreases as targets are more distant.

6 THE EQUIVALENCE-CLASS MHT (ECMHT)

The Equivalence-Class MHT (ECMHT) adopts the same initialization and prediction equations as the GNNF and MHT. We assume there are N synchronous sensors. We denote by $Z_{k1}(n)$ the smaller measurement from the n^{th} sensor at time t_k , and the larger by $Z_{k1}(n)$. We have the following filter initialization and update equations.

$$X_1(1|1) = \sum_n Z_{11}(n), \quad (6.1)$$

$$X_2(1|1) = \sum_n Z_{12}(n), \quad (6.2)$$

$$P_j(1|1) = \frac{R_1}{N}, \quad j=1,2, \quad (6.3)$$

$$S_j(k+1) = C_{k+1} P_j(k+1|k) C'_{k+1} + \frac{R_{k+1}}{N}. \quad (6.4)$$

Next, equation (6.5) determines the number of measurements to be swapped. Note that the each summation is understood to be zero when the upper limit of summation is smaller than the lower limit of summation. Once the number of swaps is determined according to (6.5), equivalent measurements are formed and the usual filter update equations are applied.

$$\omega(k+1) = \arg \max_s \left\{ \log \binom{N}{s} + \left(\sum_{i=1}^s \left(-\left(Z_{(k+1)1}(i) - X_1(k+1|k) \right)^2 \right) - \sum_{i=s+1}^N \left(-\left(Z_{(k+1)1}(i) - X_2(k+1|k) \right)^2 \right) \right) \right\}, \quad (6.5)$$

$$Z_{(k+1)1} = \frac{\sum_{i=1}^{\omega(k+1)} Z_{(k+1)1}(i) + \sum_{i=\omega(k+1)+1}^N Z_{(k+1)2}(i)}{N}, \quad (6.6)$$

$$Z_{(k+1)2} = \frac{\sum_{i=1}^{\omega(k+1)} Z_{(k+1)2}(i) + \sum_{i=\omega(k+1)+1}^N Z_{(k+1)1}(i)}{N}, \quad (6.7)$$

$$\nu_1(k+1) = Z_{(k+1)1} - X_1(k+1|k), \quad (6.8)$$

$$\nu_2(k+1) = Z_{(k+1)2} - X_2(k+1|k), \quad (6.9)$$

$$X_j(k+1|k+1) = X_j(k+1|k) + L_j(k+1)\nu_j(k+1), \quad j=1,2, \quad (6.10)$$

$$P(k+1|k+1) = (I - L(k+1)C_{k+1})P(k+1|k). \quad (6.11)$$

The ECMHT equations given here include some simplifying approximations. First, we estimate the score of a single global hypothesis by separately computing the filter innovation for each measurement in the (synchronous) set of measurements based on the current state covariance, rather than computing the filter innovation in a coupled, sequential fashion. Second, we compute the posterior probability of an equivalence class of hypotheses based on a single randomly-selected element of the set of similar global hypotheses, according to (5.22). One could instantiate a more precise computation of the posterior probability of each hypothesis class.

We now quantify filtering performance based on the PDAF, MHT, and ECMHT. We adopt the 1D OU process with parameters $q = 0.01 \text{m}^2 \text{sec}^{-1}$ and $\gamma = 0.005$. We consider 1000 Monte Carlo realizations of length 100, with $\Delta t = 1 \text{sec}$. There are 10 sensors, each with constant measurement covariance $R = 1 \text{m}^2$. Since track swap phenomena occur in general, localization performance will be measured according to the MOSPA metric, where in each scan we determine independently the best track-truth association based on distance [18].

Results are in Table 6.1. We find that the ECMHT performs best, followed by the MHT solution. The impact of track coalescence on the PDAF significantly hampers localization accuracy. Note that the *ideal* solution corresponds to the case where measurement origin is known; this reduces to a classical filtering problem to which the Kalman filter may be applied directly.

Table 6.1: Optimal swap rate decreases as targets are more distant.

Filtering scheme	MOSPA
ideal	0.0363m
PDAF	0.8518m
MHT	0.1376m
ECMHT	0.0586m

An illustration of one realization of ground truth and MHT, PDAF, and ECMHT solutions is in Figures 6.1-6.2. (Both figures correspond to the same truth and sensor data realization.)

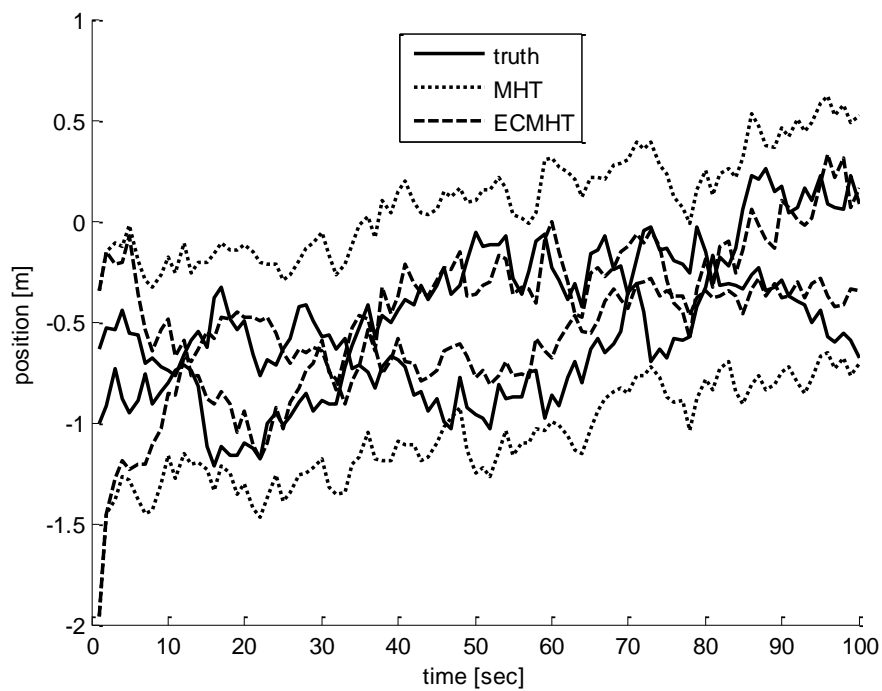


Figure 6.1: The ECMHT appears not to suffer track repulsion like the MHT.

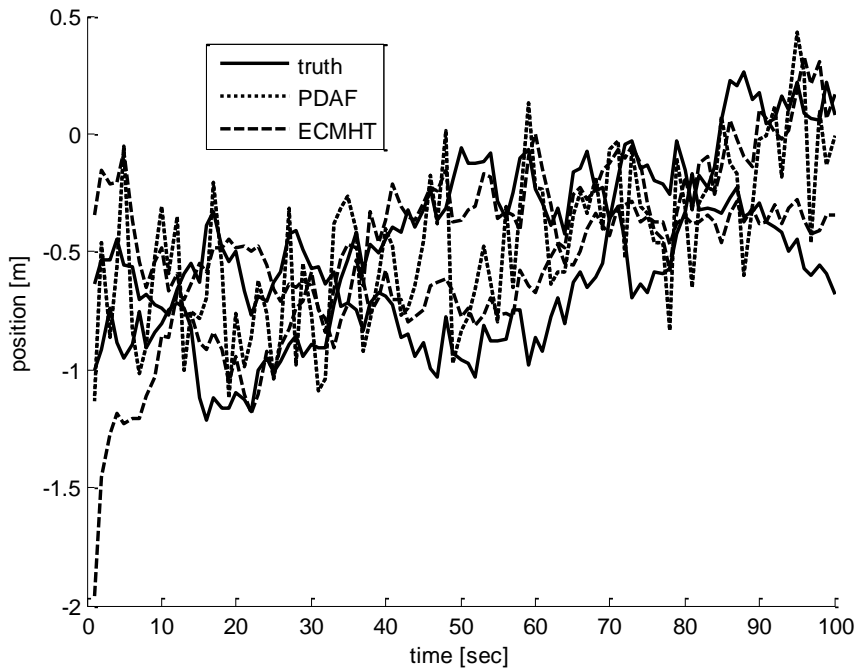


Figure 6.2: The ECMHT appears not to suffer track coalescence like the PDAF.

7 CONCLUSIONS

This manuscript introduces an interesting generalization of the MHT formalism, yielding a recursion over hypothesis classes rather than over single hypotheses. We do so for a special case of the multi-target tracking problem that we call cardinality tracking, where we measure the number of sensor measurements and no target state information is available. For this problem, a classical MHT approach does not scale: the MAP solution for non-trivial false alarm rates is the null solution. The recursion developed here allows for a significantly-improved MHT solution.

Motivated by the cardinality tracking results, we also introduce an equivalence-class MHT (ECMHT) approach to multi-target filtering for a network of synchronous sensors. The approach leads to promising performance results with respect to a recently-introduced label-free tracking metric [18]. The ECMHT appears to suffer neither track repulsion nor track coalescence. The benefits of selecting a sub-optimal solution under the MAP criterion may be related to a similar effect that has been identified recently in the context of distributed estimation with bandwidth constraints [19].

8 REFERENCES

- [1] D. Reid, An Algorithm for Tracking Multiple Targets, *IEEE Transactions on Automatic Control*, vol. 24(6), December 1979.
- [2] C. Morefield, Application of 0-1 Integer Programming to Multitarget Tracking Problems, *IEEE Transactions on Automatic Control*, vol. 22(3), June 1977.
- [3] T. Kurien, Issues in the Design of Practical Multitarget Tracking Algorithms, in *Multitarget-*

- Multisensor Tracking: Advanced Applications*, Artech House, 1990.
- [4] A. Poore and N. Rijavec, A Lagrangian Relaxation Algorithm for Multidimensional Assignment Problems Arising from Multitarget Tracking, *SIAM J. Optimization*, vol. 3(3), August 1993.
 - [5] S. Coraluppi, C. Carthel, M. Luetgen, and S. Lynch, All-Source Track and Identity Fusion, in *Proceedings of the National Symposium on Sensor and Data Fusion*, San Antonio TX, USA, June 2000.
 - [6] P. Storms and F. Spieksma, an LP-based Algorithm for the Data Association Problem in Multitarget Tracking, in *Proceedings of the 3rd International Conference on Information Fusion*, Paris, France, July 2000.
 - [7] S. Coraluppi and C. Carthel, Addressing the Greedy-Target Problem in Multiple-Hypothesis Tracking, in *Proceedings of the IEEE Aerospace Conference*, Big Sky MT, USA, March 2011.
 - [8] S. Coraluppi and C. Carthel, Aggregate Surveillance: a Cardinality Tracking Approach, in *Proceedings of the 14th International Conference on Information Fusion*, Chicago IL, USA, July 2011.
 - [9] L. Stone, C. Barlow, and T. Corwin, *Bayesian Multiple Target Tracking*, Artech House, 1999.
 - [10] Y. Bar-Shalom and X. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*, Storrs, YBS Publishing, 1995.
 - [11] R. Fitzgerald, Track Biases and Coalescence with Probabilistic Data Association, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 21(6), November 1985.
 - [12] S. Coraluppi and P. Willett, Analysis of the Track Coalescence Effect in Target Tracking with Probabilistic Data Association, in *Proceedings of the NATO Workshop on Data Fusion and Anomaly Detection for Maritime Situational Awareness*, La Spezia, Italy, September 2009.
 - [13] P. Willett, T. Luginbuhl, and E. Giannopoulos, MHT Tracking for Crossing Sonar Targets, in *Proceedings of the SPIE Conference on Signal and Data Processing of Small Targets*, August 2007.
 - [14] S. Coraluppi, C. Carthel, P. Willett, M. Dingboe, O. O'Neill, and T. Luginbuhl, The Track Repulsion Effect in Automatic Tracking, in *Proceedings of the 12th International Conference on Information Fusion*, Seattle WA, USA, July 2009.
 - [15] S. Coraluppi, C. Carthel, P. Willett, and T. Luginbuhl, Further Analysis of the Track Repulsion Effect in Automatic Tracking, in *Proceedings of the SPIE Conference on Signal and Data Processing of Small Targets*, San Diego, USA, August 2009.
 - [16] H. Blom and A. Bloem, Probabilistic Data Association Avoiding Track Coalescence, *IEEE Transactions on Automatic Control*, vol. 45(2), February 2000.
 - [17] L. Svensson, D. Svensson, and P. Willett, Set JPDA Algorithm for Tracking Unordered Sets of Targets, in *Proceedings of the 12th International Conference on Information Fusion*, Seattle WA, USA, July 2009.
 - [18] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, A consistent metric for performance evaluation of multi-

object filters, *IEEE Transactions on Signal Processing*, vol. 56(8), August 2008.

- [19] P. Braca, M. Guerriero, S. Marano, V. Matta, and P. Willett, Selective Measurement Transmission in Distributed Estimation with Data Association, in *IEEE Transactions on Signal Processing*, vol. 58(8), August 2010.