Modified Scoring in Multiple-Hypothesis Tracking

STEFANO CORALUPPI
CRAIG CARTHEL

Track-oriented multiple-hypothesis tracking is a powerful and widely-accepted methodology in multi-target tracking. We show that the target-death problem inherent in the probability hypothesis density filter does not arise in the MHT. However, the MHT suffers from a problem of its own: excessive competition for measurements from tentative tracks. We introduce a mechanism to mitigate this effect by favoring confirmed tracks in the association process. A heuristic justification for the technique is that it mitigates the sub-optimality associated with hypothesis pruning and sequential track extraction. Perhaps more convincingly, the modification to the MHT equations is provably optimal in the limiting case of cardinality tracking with unity detection probability. We show that modified-scoring MHT improves upon standard MHT in several benchmark studies.

Manuscript received April 1, 2011; revised August 12, 2011, January 12, 2012, and May 15, 2012; released for publication May 18, 2012.

Refereeing of this contribution was handled by William Blair.

Author's address: Compunetix Inc., 2420 Mosside Boulevard, Monroeville, PA 15146, E-mail: (stefano.coraluppi@compunetix.com, craig.carthel@compuneitx.com).

1557-6418/12/$17.00

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 [10]. The initial integer-programming formulation of the problem is due to Morefield [8]. The hybrid-state decomposition that allows for computationally-efficient track-oriented MHT is due to Kurien [7]. An efficient solution to the optimization problem required for n_{scan} hypothesis pruning via Lagrangian relaxation is due to Poore and Rijavec [9]. A linear-programming based relaxation approach to the same optimization problem was proposed independently by Coraluppi et al [3] and by Storms and Spieksma [12].

In practice, MHT implementations must limit the number of local (or track) hypotheses. This can be achieved by measurement gating, by limiting hypothesis generation, and by pruning or merging existing hypotheses. Additionally, sequential track extraction schemes are adopted in lieu of optimal (batch) track extraction [1]. These techniques, while necessary for computationally-realizable and real-time MHT processing, lead to suboptimal data association decisions and track extraction. In this paper, we show that the suboptimality can be mitigated by favoring nearly-confirmed and confirmed tracks over tentative ones in the data-association process with suitable modification to the MHT track scoring equations.

This paper is organized as follows. In Section 2, we summarize the hybrid-state derivation of the track-oriented MHT, with some modifications with respect to the original derivation [7]. In Section 3, we address briefly the *target-death problem* that arises in *probability hypothesis density* (PHD) filtering as discussed in Erdinc et al [6], and show that it does not arise in track-oriented MHT. In Section 4, we introduce the modified-scoring MHT equations and considering a limiting case of the general tracking problem that we call *cardinality tracking*. Section 5 provides simulation results that demonstrate the improved performance of modified-scoring MHT over standard MHT. Concluding remarks are in Section 6.

2. MULTIPLE-HYPOTHESIS TRACKING

A key challenge in multi-sensor multi-target tracking is *measurement origin uncertainty*. That is, unlike a classical nonlinear filtering problem, we do not know how may objects are in the surveillance region, and which measurements are to be associated. New objects may be born in any given scan, and existing objects may die.

We assume that for each sensor scan, contact-level (or detection-level) data is available, in the sense that signal processing techniques are applied to raw sensor data yielding contacts for which the detection and localization statistics are known. We are interested in a scan-based (or real-time) approach that, perhaps with some delay, yields an estimate of the number of objects and corresponding object state estimates at any time.

Several approaches to contact-level scan-based tracking exist. In this section, we employ a hybrid-state formalism to describe the track-oriented multiple-hypothesis tracking approach. Our approach follows closely the one introduced in [7]. We assume Poisson distributed births at each scan with mean λ_b, Poisson distributed false returns with mean λ_{fa}, object detection probability p_d, object death or termination probability p_χ at each scan. (We neglect the *time-dependent* nature of birth and death probabilities as would ensue from an underlying continuous-time formulation, and we neglect as well inter-scan birth and death events.)

We have a sequence of sets of contacts $Z^k = (Z_1,\ldots, Z_k)$, and we wish to estimate the state history X^k for all objects present in the surveillance region. X^k is compact notation that represents the state trajectories of targets that exist over the time sequence $(t_1,\ldots,t_k)$. Note that each target may exist for a subset of these times, with a single birth and a single death occurrence i.e. targets do not reappear. We introduce the auxiliary discrete state history q^k that represents a full interpretation of all contact data: which contacts are false, how the object-originated ones are to be associated, and when objects are born and die. There are two fundamental assumptions of note. The first is that there are no target births in the absence of a corresponding detection, i.e. we do not reason over new, undetected objects. The second is that there is *at most* one contact per object per scan.

We are interested in the probability distribution $p(X^k \mid Z^k)$ for object state histories given data. This quantity can be obtained by conditioning over all possible auxiliary states histories q^k.

$$p(X^k \mid Z^k) = \sum_{q^k} p(X^k, q^k \mid Z^k)$$
$$= \sum_{q^k} p(X^k \mid Z^k, q^k) p(q^k \mid Z^k). \qquad (1)$$

A pure MMSE approach would yield the following:

$$\hat{X}_{\mathrm{MMSE}}(Z^k) = E[X^k \mid Z^k] = \sum_{q^k} E[X^k \mid Z^k, q^k] p(q^k \mid Z^k). \qquad (2)$$

The track-oriented MHT approach is a mixed MMSE/MAP one, whereby we identify the MAP estimate for the auxiliary state history q^k, and identify the corresponding MMSE estimate for the object state history X^k conditioned on the estimate for q^k.

$$\hat{X}(Z^k) = E[X^k \mid Z^k, \hat{q}_k] \qquad (3)$$

$$\hat{q}^k = \hat{q}_{\mathrm{MAP}}(Z^k) = \arg\max p(q^k \mid Z^k). \qquad (4)$$

The MHT recursion. Each feasible q^k corresponds to a global hypothesis. (The set of global hypotheses is generally constrained via measurement gating and hypothesis generation logic.) We are interested in a *recursive* and *computationally efficient* expression for $p(q^k \mid Z^k)$ that lends itself to maximization without the need for explicit enumeration of global hypotheses. We do so through repeated use of Bayes' rule. Note that $f(\cdot)$ denotes the probability density function and $p(\cdot)$ denotes the probability mass functions. The normalizing constant c_k does not impact MAP estimation.

$$p(q^k \mid Z^k) = \frac{f(Z_k \mid Z^{k-1}, q^k) p(q^k \mid Z^{k-1})}{c_k}$$
$$= \frac{f(Z_k \mid Z^{k-1}, q^k) p(q_k \mid Z^{k-1}, q^{k-1}) p(q^{k-1} \mid Z^{k-1})}{c_k} \qquad (5)$$

$$c_k = f(Z_k \mid Z^{k-1}) = \sum_{q^k} f(Z_k \mid Z^{k-1}, q^k) p(q^k \mid Z^{k-1}). \qquad (6)$$

Recall that we assume that in each scan the number of target births is Poisson distributed with mean λ_b, the number of false returns is Poisson distributed with mean λ_{fa}, targets die with probability p_χ, and targets are detected with probability p_d. The recursive expression (5) involves two factors that we consider in turn.

Computation of $p(q_k \mid Z^{k-1}, q^{k-1})$. It will be useful to introduce the aggregate variable ψ_k (consistent with the approach in [7]) 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.

$$p(q_k \mid Z^{k-1}, q^{k-1}) = p(\psi_k \mid Z^{k-1}, q^{k-1}) p(q_k \mid Z^{k-1}, q^{k-1}, \psi_k) \qquad (7)$$

$$p(\psi_k \mid Z^{k-1}, q^{k-1}) = \left\{ \binom{\tau}{\chi} p_\chi^\chi (1-p_\chi)^{\tau-\chi} \right\} \cdot \left\{ \binom{\tau-\chi}{d} p_d^d (1-p_d)^{\tau-\chi-d} \right\} \cdot \left\{ \frac{\exp(-p_d\lambda_b) p_d^b \lambda_b^b}{b!} \right\} \cdot \left\{ \frac{\exp(-\lambda_{fa}) \lambda_{fa}^{r-d-b}}{(r-d-b)!} \right\} \qquad (8)$$

$$p(q_k \mid Z^{k-1}, q^{k-1}, \psi_k) = \frac{1}{\binom{\tau}{\chi} \binom{\tau-\chi}{d} \left(\frac{r!}{(r-d)!}\right) \binom{r-d}{b}}. \qquad (9)$$

Substituting (8–9) into (7) and simplifying yields the following.

$$p(q_k \mid Z^{k-1}, q^{k-1}) = \left\{ \frac{\exp(-p_d\lambda_b - \lambda_{fa})\lambda_{fa}^r}{r!} \right\} \cdot p_\chi^\chi((1-p_\chi)(1-p_d))^{\tau-\chi-d} \cdot \left(\frac{(1-p_\chi)p_d}{\lambda_{fa}}\right)^d \left(\frac{p_d\lambda_b}{\lambda_{fa}}\right)^b. \tag{10}$$

Computation of $f(Z_k \mid Z^{k-1}, q^k)$. This quantity is given by (11), where $Z_k = \{z_j, 1 \le j \le r\}$, J_d is the set of measurements associated with detections of existing tracks, J_b is the set of measurements associated with target births, J_{fa} is the set of measurements hypothesized as false, $|J_d| + |J_b| + |J_{fa}| = r$, and the factors on the R.H.S. are derived from filter innovations, filter initiations, and the false contact distribution (generally uniform over measurement space).

$$f(Z_k \mid Z^{k-1}, q^k) = \prod_{j\in J_d} f_d(z_j \mid Z^{k-1}, q^k) \cdot \prod_{j\in J_b} f_b(z_j \mid Z^{k-1}, q^k) \prod_{j\in J_{fa}} f_{fa}(z_j \mid Z^{k-1}, q^k). \tag{11}$$

For example, in the linear Gaussian case, $f_d(z_j \mid Z^{k-1}, q^k)$ is a Gaussian residual, i.e. it is the probability of observing z_j given a sequence of preceding measurements. If there is no prior information on the target, $f_b(z_j \mid Z^{k-1}, q^k)$ is generally the value of the uniform density function over measurement space. Similarly, $f_{fa}(z_j \mid Z^{k-1}, q^k)$ is as well usually taken to be the value of the uniform density function over measurement space, under the assumption of uniformly distributed false returns. Note that the expressions given here are general and allow for quite general target and sensor models.

Final form of the MHT recursion. Substituting (10–11) into (5) and simplifying results in (12–13). This expression is the key enabler of track-oriented MHT. In particular, it provides a recursive expression for $p(q^k \mid Z^k)$ that consists of a number of factors that relate to its constituent local track hypotheses.

$$p(q^k \mid Z^k) = p_\chi^\chi((1-p_\chi)(1-p_d))^{\tau-\chi-d} \cdot \prod_{j\in J_d}\left[\frac{(1-p_\chi)p_d f_d(z_j \mid Z^{k-1}, q^k)}{\lambda_{fa} f_{fa}(z_j \mid Z^{k-1}, q^k)}\right] \cdot \prod_{j\in J_b}\left[\frac{p_d\lambda_b f_b(z_j \mid Z^{k-1}, q^k)}{\lambda_{fa} f_{fa}(z_j \mid Z^{k-1}, q^k)}\right] \frac{p(q^{k-1} \mid Z^{k-1})}{\bar{c}_k} \tag{12}$$

$$\bar{c}_k = \frac{c_k}{\left\{\frac{\exp(-p_d\lambda_b - \lambda_{fa})}{r!}\lambda_{fa}^r\right\} \prod_{j\in J_d\cup J_b\cup J_{fa}} f_{fa}(z_j \mid Z^{k-1}, q^k)}. \tag{13}$$

Note that the constant $\bar{c}_k$ normalizes the recursion with respect to the case in which all returns in the current scan are false. That is, for the case $b = 0$ (no births) and $\tau = d = \chi = 0$ (no current tracks, and correspondingly no detections or terminations on current tracks), we have

$$p(q^k \mid Z^k) = \left\{\frac{\exp(-p_d\lambda_b - \lambda_{fa})}{r!}\lambda_{fa}^r\right\} \cdot \prod_{j\in J_d\cup J_b\cup J_{fa}} f_{fa}(z_j \mid Z^{k-1}, q^k) \frac{p(q^{k-1} \mid Z^{k-1})}{c_k}.$$

That is, the denominator in (13) is precisely the product of the probability of no detected births, i.e. $\exp(-p_d\lambda_b) \cdot p_d^0\lambda_b^0/0!$, the probability of r false alarms, i.e. $(\exp(-\lambda_{fa})/r!)\lambda_{fa}^r$, and the filter residuals associate with all measurements being false.

An implicit reduction in the set of hypotheses in (12–13) is that target births are assumed to occur only in the presence of a detection (i.e. there is no reasoning over un-detected births). Correspondingly, the factor p_d reduces the effective birth rate to $p_d\lambda_b$ (though surprisingly the factor is absent in [7]). Further, in the first scan of data, it would be appropriate to replace $p_d\lambda_b$ by $p_d\lambda_b/p_\chi$ to account properly for the steady-state expected number of targets. (More generally, target birth and death parameters should reflect sensor scan rates, as the underlying target process is defined in continuous time.) Further reduction in the set of hypotheses is generally achieved via *measurement gating* procedures [1]. Finally, for a given track hypothesis, one usually applies rule-based spawning of a missed detection or termination hypothesis, but not both (e.g. only spawn a missed detection hypothesis until a sufficiently-long sequence of missed detections is reached).

One cannot consider too large a set of scans before pruning or merging local (or track) hypotheses in some fashion. A popular mechanism to control these hypotheses is n_{scan} *pruning*. This amounts to solving (4), generally by a relaxation approach to an integer programming problem [3, 8–9, 12], followed by pruning of all local hypotheses that differ from $\hat{q}^k$ at a depth of n_{scan}. That is, all remaining global hypotheses are identical up to scan $k - n_{\text{scan}}$. Note that, if one were to set $n_{\text{scan}} = 0$, this would amount to immediate resolution of association hypotheses up to the current time. The n_{scan} pruning methodology is applied after each new scan of data is received, resulting in a fixed-delay solution to the tracking problem.

Often, n_{scan} pruning is referred to as a *maximum likelihood* (ML) approach to hypothesis management.

ML estimation is closely related to *maximum a posteriori* (MAP) estimation. In particular, we have:

$$\hat{X}_{\mathrm{MAP}}(y) = \arg\max f(y \mid X) f(X) \tag{14}$$

$$\hat{X}_{\mathrm{ML}}(y) = \arg\max f(y \mid X). \tag{15}$$

Note that ML estimation is a *non-Bayesian* approach as it does not rely on a prior distribution on X. ML estimation can be interpreted as MAP estimation with a uniform prior. In the track-oriented MHT setting, n_{scan} pruning relies on a single parent global hypothesis, thus the ML and MAP interpretations are both valid.

Once hypotheses are resolved, in principle one has a state of object histories given by $\hat{X}(Z^k)$. In practice, it is common to apply track confirmation and termination logic to all object histories [1]. A justification for this is that it provides a mechanism to remove spurious tracks induced by the sub-optimality inherent in practical MHT implementations that include limited hypothesis generation and hypothesis pruning or merging. Further, sequential track extraction allows for real-time processing which optimal (batch) track extraction would not.

Given the use of post-association track confirmation and termination logic, a reasonable simplification that is pursued in [3] is to employ equality constraints in the data-association process, which amounts to accounting for all contact data in the resolved tracks. Spurious tracks are subsequently removed in the track-extraction stage.

To summarize, at each stage of processing, track-oriented MHT maintains a set of track trees with depth n_{scan}. When a scan of measurements is received, each measurement is compared with each (local) track hypothesis, and a new level of leaf nodes is created. All track hypotheses continue as well in the absence of a measurement. Additionally, each measurement defines the root of a new track tree. Following hypothesis generation, the MAP global hypothesis is determined via a linear programming relaxation approach [3]. Correspondingly, the set of track trees is pruned so that a single global hypothesis exists at depth $n_{\mathrm{scan}} + 1$. The process then repeats for the next scan. (If one were to set n_{scan}, the procedure reverts to a standard 2D assignment solution.)

Data association is followed by track extraction. Tentative tracks are reported at the tracker output only once a suitable track-confirmation criterion is achieved. Similarly, once a track has degraded sufficiently (or once it is determined that a still-tentative track cannot achieve the confirmation criterion), the track is terminated. This information flows back to the data association module from the track extraction module, invalidating subsequent association hypotheses for the terminated track.

3. THE TARGET-DEATH PROBLEM

A useful re-interpretation of the *probability hypothesis density* (PHD) filter, known as the bin-occupancy filter, is given in [6]. This paper describes as well the target death problem that the authors had earlier identified, and which in turn has led to the *cardinalized* PHD (CPHD) filter.

Consider the single-target case with no false alarms. In the absence of a target measurement, the PHD surface follows (16). Note that the PHD surface $D_{k|k}(x)$ at each time t_k is a function of all data received up to t_k and is computed recursively. $D_{k|k}(x)$ admits the interpretation that it identifies the probability of target presence at a given state.

$$D_{k|k}(x) = (1 - p_d(x)) D_{k|k-1}(x). \tag{16}$$

While (16) may appear reasonable, it can be shown that it is inconsistent with the following simple Bayesian argument. Let Y_{k-1} be the existence state for the target at scan $k-1$, and assume that the death probability at any scan is given by p_χ, as before. The updated probability of existence after a missed detection is given by (17).

$$\begin{aligned} p(Y_k = 1 \mid |Z_k| = 0) &= p(Y_k = 1 \mid Y_{k-1} = 1, |Z_k| = 0) p(Y_{k-1} = 1) \\ &= \frac{p(Y_k = 1, |Z_k| = 0 \mid Y_{k-1} = 1) p(Y_{k-1} = 1)}{p(|Z_k| = 0 \mid Y_{k-1} = 1)} \\ &= \frac{(1 - p_\chi)(1 - p_d)}{1 - (1 - p_\chi) p_d} p(Y_{k-1} = 1). \end{aligned} \tag{17}$$

Comparing (16) and (17), we see that the PHD filter penalizes missed detections too heavily; it is claimed in [6] that the CPHD appears to follow (17).

What happens with the track-oriented MHT approach? We compare the ratio of the probability associated with the track *coast* hypothesis (track is alive in the absence of a measurement) with the probability of track *coast* or *death*. That is, in the numerator we want the case "no detection and target alive," and in the denominator we want the case "no detection (target alive or dead." Let q_i^k and q_j^k denote global hypotheses that include coast and death events, respectively, for the target of interest. From (12), we see that (18) follows immediately. Indeed, all factors in the global hypothesis probability cancel except for those associated with the (undetected) track.

$$\begin{aligned} \frac{p(q_i^k \mid Z^k)}{p(q_i^k \mid Z^k) + p(q_j^k \mid Z^k)} &= \frac{(1 - p_\chi)(1 - p_d)}{(1 - p_\chi)(1 - p_d) + p_\chi} \\ &= \frac{(1 - p_\chi)(1 - p_d)}{1 - (1 - p_\chi) p_d}. \end{aligned} \tag{18}$$

Note that this validation is quite general, and in particular it is directly applicable to the multi-target case, under the assumption that no contacts satisfy the hypothesis gating criterion for the (undetected) object of interest here. We conclude that track-oriented MHT properly handles missed detections, and no target-death problem is observed.

4. MODIFIED-SCORING MHT AND CARDINALITY TRACKING

In scan-based processing, assume that a track hypothesis is *confirmed* when it achieves an M-of-N criterion, with the start of the (tentative) track defined with the first of the relevant M measurements, and that tentative tracks that have no chance to achieve the M-of-N criterion are discarded. Also, a track hypothesis is *terminated* if K missed detections are exceeded. Note that, under multiple-hypothesis processing, a confirmed track may be pruned under a hypothesis-reduction scheme such as n_{scan} pruning. Following hypothesis resolution, a single global hypothesis exists that is composed of a number of *resolved* tracks. Next, the set of resolved tracks undergoes a *track extraction* process based on the same M, N, and K parameters. With these notions of confirmed, resolved and extracted tracks, we now introduce a modification to (12) that will prove useful. In particular, confirmation reward factors $\xi_2 > \xi_1 > 1$ are applied to track updates for *confirmed* and *nearly-confirmed* track hypotheses. The later refers to tentative tracks that reach confirmation in the current scan. The measurement sets for confirmed, nearly-confirmed, and tentative tracks are denoted by J_d, J_c, and J_t, respectively.

$$p(q^k \mid Z^k) = p_\chi^\chi((1-p_\chi)(1-p_d))^{\tau-\chi-d} \cdot \prod_{j\in J_t}\left[\frac{(1-p_\chi)p_d f_d(z_j \mid Z^{k-1},q^k)}{\lambda_{fa} f_{fa}(z_j \mid Z^{k-1},q^k)}\right] \cdot \prod_{j\in J_c}\left[\frac{(1-p_\chi)\xi_1 p_d f_d(z_j \mid Z^{k-1},q^k)}{\lambda_{fa} f_{fa}(z_j \mid Z^{k-1},q^k)}\right] \cdot \prod_{j\in J_d}\left[\frac{(1-p_\chi)\xi_2 p_d f_d(z_j \mid Z^{k-1},q^k)}{\lambda_{fa} f_{fa}(z_j \mid Z^{k-1},q^k)}\right] \cdot \prod_{j\in J_b}\left[\frac{p_d \lambda_b f_b(z_j \mid Z^{k-1},q^k)}{\lambda_{fa} f_{fa}(z_j \mid Z^{k-1},q^k)}\right]\frac{p(q^{k-1} \mid Z^{k-1})}{\bar{c}'_k}. \tag{19}$$

We refer to *standard MHT* as solution to (4) based on (12–13) with a fixed hypothesis tree depth (n_{scan}). We refer to *modified-scoring MHT* as the solution to (4) based on (13, 19) with a fixed hypothesis tree depth (n_{scan}). Note that the normalization factor $\bar{c}'_k$ in (19) differs slightly from the normalization factor in (12), since the track hypothesis scores have been modified with the confirmation reward factors.

The use of the reward factors $\xi_2 > \xi_1 > 1$ amounts to favoring confirmed and nearly-confirmed tracks in the association process. While this appears reasonable in (sub-optimal) MHT processing, we provide justification for the procedure on two grounds: (1) optimality of modified-scoring MHT in the limiting case of the tracking problem known as *cardinality tracking*; (2) simulation results for the general case. We address (1) next, while (2) is treated in Section 5.

Let us consider now the case where measurements are not informative with respect to target state: we are given only a sequence of cardinality measurements. Assume we are given birth, death, detection and false alarm statistics as well as a sequence that specifies the number of measurements received. An example might be (1, 2, 3, 3, 1...). We must decide how many targets there are as a function of time.

Note that all filter residuals in (12) are identical, leading to (20); correspondingly, (19) leads to (21) with c_1 the number of nearly-confirmed tracks and c_2 is the number of confirmed tracks. In this context, note that measurement gating is not a meaningful concept as all track updates are equivalent. Then, *cardinality tracking* involves identifying the sequence of target cardinalities $|X|^k$ given the sequence of measurement cardinalities $|Z|^k$.

$$p(q^k \mid Z^k) = \frac{p_\chi^\chi(1-p_\chi)^{\tau-\chi}(1-p_d)^{\tau-\chi-d}p_d^{b+d}\lambda_b^b}{\lambda_{fa}^{d+b}}\frac{p(q^{k-1} \mid Z^{k-1})}{\bar{c}_k} \tag{20}$$

$$p(q^k \mid Z^k) = \frac{p_\chi^\chi(1-p_\chi)^{\tau-\chi}(1-p_d)^{\tau-\chi-d}\xi_1^{c_1}\xi_2^{c_2}p_d^{b+d}\lambda_b^b}{\lambda_{fa}^{d+b}} \cdot \frac{p(q^{k-1} \mid Z^{k-1})}{\bar{c}_k}. \tag{21}$$

For purposes of the ensuing analysis, it is useful to introduce some assumptions regarding the parameters in (20) so that the form of the optimal solution to (4) yields a reasonable structure as explained below. It will be useful to represent a tracking solution X^k in a compact manner, where each track is represented as a sequence of existence states, with 1 denoting measurement update, 0 denoting existence with no measurement update i.e. a track *coast*, and x denoting non-existence. For example, $(x,x,1,1,0)$ represents a track that exists beginning with the third sensor scan, involves two measurements and one track coast and then terminates.

As a reminder, the assumptions below apply only to the cardinality-tracking problem.

Assumption 1 (preference for longer tracks).

$$\frac{p_d}{\lambda_{fa}}(1-p_\chi) > 1.$$

Consider $|Z|^k = \{1,1\}$. Assumption 1 insures that solution $\bar{X}^k = \{(1)\}$ or $\bar{X}^k = \{(x,1)\}$ has lower probability than $\tilde{X}^k = \{(1,1)\}$, i.e. $p(\tilde{X}^k \mid Z^k) > p(\bar{X}^k \mid Z^k)$.

Assumption 2 (singleton tracks discarded).

$$\lambda_b p_\chi \frac{p_d}{\lambda_{fa}} < 1.$$

Consider $|Z|^k = \{1\}$. Assumption 2 insures that solution $\bar{X}^k = \{(1)\}$ has lower probability than $\tilde{X}^k = \emptyset$, i.e. $p(\tilde{X}^k \mid Z^k) > p(\bar{X}^k \mid Z^k)$.

		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

Fig. 1. For $|Z|^k = \{1,2,6,4,3,1,2,2,2,1,0,0,2,3,2,1,2,0,3,1\}$, a Tetris solution with parameter 3 is illustrated above, for which $|X|^k = \{1,2,3,3,3,1,2,2,2,1,0,0,2,3,2,1,1,0,0,0\}$. Measurements that are part of the solution are denoted by grey cells.

…	0	1_A	1_B	0	…
…	1_C	1_A	0	…	…

Fig. 2. A violation of Tetris structure.

Assumptions 1–2 imply the following (preference for track association). $\lambda_b p_\chi < 1 - p_\chi$.

Consider $|Z|^k = \{1,1\}$. The above inequality insures that solution $\bar{X}^k = \{(1),(x,1)\}$ has lower probability than $\tilde{X}^k = \{(1,1)\}$, i.e. $p(\tilde{X}^k \mid Z^k) > p(\bar{X}^k \mid Z^k)$.

Note that Assumptions 1–2 place limits on allowable clutter rates for non-empty optimal tracking solutions; the interested reader is referred to [2] for further discussion of this issue.

We consider now a special case of cardinality tracking ($p_d = 1$) for which a number of results can be established. First, we define a *Tetris solution* to be the solution obtained with sequential track extraction maximizing track length with contiguous sequences of measurements. The Tetris solution is parameterized by a minimum track length parameter, such that tracks shorter than a specified threshold are extracted. The solution is best described by illustration: see Fig. 1.

Result 1 (structure of optimal solution). Let $p_d = 1$. An optimal solution to (4) is given by the Tetris solution with minimum track length parameter

$$k_0 = \min_i \left\{ i, \frac{\lambda_b (1-p_\chi)^{i-1} p_\chi}{\lambda_{fa}^i} \geq 1 \right\}.$$

Indeed, since tracks with length less than k_0 contribute a score less than unity to the posterior probability, a Tetris solution with parameter less than k_0 is not optimal. Similarly, a Tetris solution with parameter greater than k_0 will not include tracks that contribute a score greater than unity to the posterior probability. Thus, such a Tetris solution is not optimal either. It remains to show that a non-Tetris solution cannot outperform the Tetris solution with parameter k_0. Assume a non-Tetris optimal solution exists, and that it cannot be re-expressed as a Tetris solution by a re-ordering of entire rows (else the solution is equivalent to a Tetris one). In particular, the non-Tetris solution must contain two (possibly non-neighboring) row portions that are as shown in Fig. 2, where each cell denotes a sequence of zeros or ones of arbitrary dimension.

We now show that dropping 1_B to the lower row (i.e. partial row reordering) yields a posterior probability that is equal or higher. Indeed, if 1_A is a track of length k_0 or greater, the two posterior probabilities are the same (Fig. 3-top). If 1_A is a track of length less than k_0, the solution with the top-row 1_A replaced by zeros has larger posterior probability (Fig. 3-bottom). Thus, by a sequence of steps of this kind, we recover a Tetris structure. This shows that a non-Tetris solution cannot outperform the optimal Tetris one.

…	0	1_A	0	0	…
…	1_C	1_A	1_B	…	…

…	0	0	0	0	…
…	1_C	1_A	1_B	…	…

Fig. 3. Equivalent solution (top) or improved solution in the case of a short track 1_A (bottom).

Result 2 (optimality of modified MHT). The modified MHT solution with $M = k_0$, $N = k_0$, $K = 0$ and $n_{\text{scan}} \geq k_0 - 3$ is optimal.

Result 2 is best illustrated by example. First, assume that the target and sensor parameters are such that $k_0 = 3$ in Result 1. According to Result 2, modified MHT with $n_{\text{scan}} \geq 0$ is optimal. Consider the measurement sequence $|Z|^k = \{1,2,1\}$. With standard MHT with arbitrary n_{scan}, one obtains either the set of tentative track $X^k_{\text{tentative}} = \{(1,1,1)\}$ or $X^k_{\text{tentative}} = \{(1,1),(x,1,1)\}$. Indeed, there is no preference in terms of posterior probabilities in associating the measurement in the third scan with the longer or shorter tentative track. Correspondingly, after track extraction, one obtains either $X^k = \{(1,1,1)\}$ or $X^k = \emptyset$. The cardinality-tracking result is thus either $|X|^k = \{1,1,1\}$ or $|X|^k = \{0,0,0\}$. With modified-scoring MHT, one is guaranteed that the measurement in the third scan is associated to the tentative track of length two; indeed, this track is nearly-confirmed, and $\xi_1 > 1$ in (21) insures that the solution to (4) yields $|X|^k = \{1,1,1\}$. Thus, modified-scoring MHT achieves optimality while standard MHT is not guaranteed to do so.

Next, assume once again that $k_0 = 3$ in Result 1 and consider the measurement sequence $|Z|^k = \{1,2,2,1\}$. By similar reasoning, we see that standard MHT results in either $X^k_{\text{tentative}} = \{(1,1,1,1),(x,1,1)\}$ or $X^k_{\text{tentative}} = \{(x,1,1,1),(1,1,1)\}$. After track extraction, one thus obtains either $X^k = \{(1,1,1,1)\}$ or $X^k = \{(1,1,1),(1,1,1)\}$. The posterior probability associated with the latter solution is the same as for the solution $X^k = \{(1,1,1,1),(x,1,1)\}$. This is immediately seen to have posterior

TABLE I
Simulation Parameters for the Cardinality Tracking Problem

Parameter Description	Setting
Target birth rate	1
Target death probability	0.1
Sensor probability of detection	1
Sensor false alarm rate	0.33
Track initiation	3-of-3
Hypothesis tree depth (n_{scan})	1
Confirmation reward factors (modified-scoring MHT)	2 (confirmed tracks), 1.5 (nearly-confirmed tracks)
Track termination (maximum missed detections)	0
Number of scans in each scenario realization	100
Number of realizations	1000

probability that does not exceed that of $X^k = \{(1,1,1,1)\}$, since a track of length two does not contribute to the posterior probability (see definition of k_0 in Remark 1). With modified-scoring MHT, the measurement in the fourth scan is guaranteed to be associated to the confirmed track rather than to the nearly confirmed one, since $\xi_2 > \xi_1$ in (21) insures that the solution to (4) yields $|X|^k = \{1,1,1,1\}$.

Finally, assume that the $k_0 = 4$ in Result 1 and consider the measurement sequence $|Z|^k = \{1,2,1,1\}$. Result 2 tells us that modified-scoring MHT requires $n_{\text{scan}} \geq 1$ to insure optimality. Indeed, with $n_{\text{scan}} = 0$, the measurement in the third scan will be associated with either the shorter or longer track, since neither is nearly confirmed. Thus, either $X^k_{\text{tentative}} = \{(1,1,1,1),(x,1)\}$ or $X^k_{\text{tentative}} = \{(1,1),(x,1,1,1)\}$ results, from which we have $X^k = \{(1,1,1,1)\}$ or $X^k = \emptyset$, respectively. This in turn leads either to solution $|X|^k = \{1,1,1,1\}$ or $|X|^k = \{0,0,0,0\}$. Conversely, with $n_{\text{scan}} = 1$, we do not decide on which track is updated with the measurement in the third scan until the fourth scan is received. Accordingly, $\xi_1 > 1$ in (21) insures that the solution to (4) yields $X^k = \{(1,1,1,1)\}$ and thus $|X|^k = \{1,1,1,1\}$.

The importance of this section is that it demonstrates the superiority of modified MHT over standard MHT in a limiting case. For this case, we are able to show that modified MHT with a sufficiently large hypothesis tree depth achieves optimality in the sense of maximizing the posterior probability over all hypotheses. Modified MHT processing introduces a mechanism whereby preference is given to tracks that have achieved or will achieve track confirmation. This is an interesting result in its own right, and provides motivation for use of modified MHT in a more general setting.

We now illustrate the performance of modified-scoring MHT and standard MHT approaches to the cardinality tracking problem for a specific numerical example. The example provides experimental validation of the claims in Results 1–2. A nice aspect of evaluating cardinality tracking is that it is much easier to provide statistically significant results for which tracking parameters are matched to target and sensor characteristics. Indeed, ground truth is obtained via a Poisson birth-death process and kinematic-space realizations are absent, so that we are not limited to a small set of benchmark scenarios. The simulation parameters are captured in Table I.

The parameters in Table I satisfy Assumption 1–2. Note that, as sensor measurements are not informative with regard to target state and are only relevant to target existence, the tolerable false alarm rates are quite low compared to a general tracking problem. The tracking initiation and termination settings and the choice of n_{scan} are consistent with the requirements for Result 1–2:

$$\frac{\lambda_b(1-p_\chi)^2 p_\chi}{\lambda_{fa}^3} > 1 > \frac{\lambda_b(1-p_\chi)p_\chi}{\lambda_{fa}^2} \Rightarrow k_0 = 3,$$

$$n_{\text{scan}} \geq k_0 - 2 \Rightarrow n_{\text{scan}} \geq 1.$$

An illustration of one realization is given in Fig. 4, along with the corresponding modified-scoring MHT output. Note that we provide a compact representation of ground truth X^k, sensor measurements Z^k, and tracker output $\hat{X}^k$: we illustrate the sequence of cardinalities $|X|^k$, $|Z|^k$, and $|\hat{X}|^k$.

Statistical performance results are based on computation of the posterior probability $p(q^k \mid Z^k)$. We find as expected that the modified-scoring MHT is optimal in the posterior-probability sense. Standard MHT suffers a performance loss resulting in a (normalized) posterior probability of 0.958. (By normalized posterior probability, we mean the ratio of the probabilities associated with the standard and modified MHT solutions, respectively.)

5. SIMULATION RESULTS FOR THE GENERAL TRACKING PROBLEM

We now evaluate modified MHT and standard MHT approaches to the general tracking problem for several scenarios of interest. First, we identify the performance metrics for this analysis. Our approach to tracker performance evaluation is somewhat novel as we do not identify a global mapping of tracks to targets. Indeed, a global mapping can be problematic due to track swap phenomena, true tracks that are seduced by false contacts and become false tracks, etc. Instead, we rely on a scan-based association of tracks to targets consistent with the recently-introduced Optimal Subpattern Assignment (OSPA) metric [11].

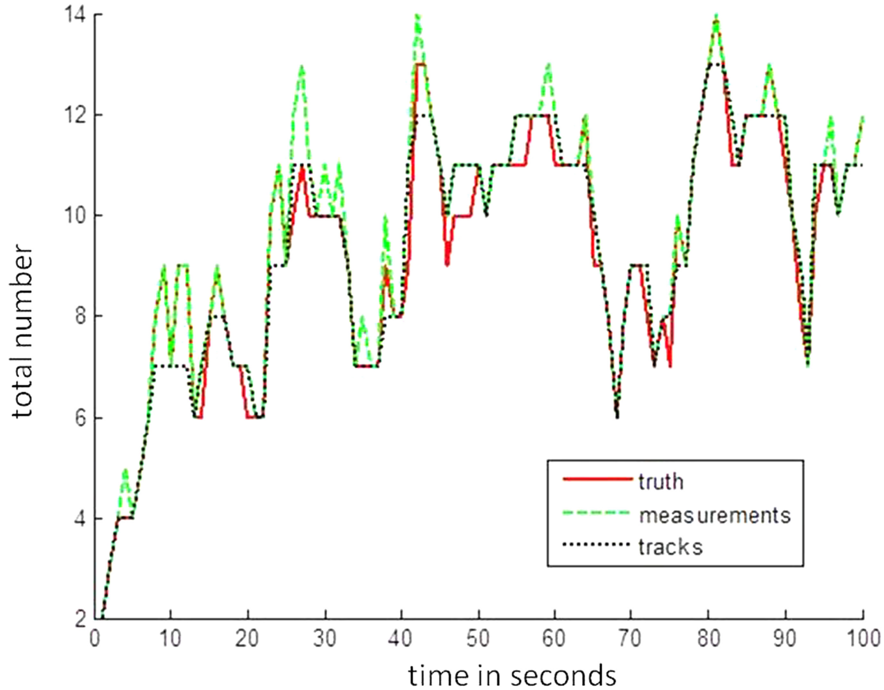

Fig. 4. One realization of truth, measurement and optimal track cardinality sequences.

For each scan time $t_i \in (t_1, \ldots, t_N)$, we have an $N_i \times M_i$ cost matrix A_i where N_i and M_i are the number of targets and tracks in existence, respectively. Given a distance threshold ξ on feasible track-truth assignments, we determine the optimal OSPA assignment between tracks and targets for each scan in a given dataset as described in [11]. For a given scan, those tracks that are assigned to targets are deemed to be *true track instances*. Correspondingly, there is a *detected target instance*. Let g_i denote the number of such (target, track) pairs for the scan at time t_i. Next, we compute the following metrics for each scenario realization:

- *Track PD*—ratio of total number of true track instances (summed over all scan times) to total number of target existence instances (summed over all targets and all scan times): $\sum_{i=1}^{N} g_i / \sum_{i=1}^{N} N_i$;
- *Track quality*—ratio of number of true track instances and total number of track instances: $\sum_{i=1}^{N} g_i / \sum_{i=1}^{N} M_i$;
- *Track purity*—ratio of number of true track instances that are as well from the mode assignment (i.e. from the most frequently associated target) and total number of track instances: $\sum_{i=1}^{N} \bar{g}_i / \sum_{i=1}^{N} M_i$; here, $\bar{g}_i \leq g_i$ is the number of truth-track assignments where truth is the *mode target* for the track, i.e. the target to which the track is associated the most.
- *Track rate*—ratio of total number of tracks to total number of targets;
- *Track localization error*—average displacement between true track instances and corresponding target location that we denote by σ_T.

Since our metrics do not rely on classifying each track as *true* or *false*, the false track statistics are understood as follows. First, the *track rate* metric answers the question: how many tracks does the system generate, relative to the true number of targets? Secondly, the *track quality* metric answers the question: for any given track at any given time, what is the probability that it is target originated? That is, track quality is the total duration of good tracks as a fraction of the overall duration of all tracks. Thus, these metrics provide an assessment of how much false track (both in number and in duration) is generated by the system, without the need for global track assessment that is often problematic when tracks are partially target-originated and partially false.

We report here on our metrics, where for each of three benchmark scenarios the metrics are averaged over multiple Monte Carlo realizations. Complete simulation parameters are identified in Table II. Illustrations of one modified MHT tracker output realization of each of three scenarios are given in Figs. 5–7; Fig. 8 illustrates the realization of the corresponding measurement data for the third scenario. Scenario 1 includes three linear-motion targets that move with identical speeds but displaced in the y dimension and with different birth and death times. Scenario 2 includes a single maneuvering target. Scenario three includes three maneuvering targets that are matched in birth and death times and in velocities and are displaced in the x dimension.

Monte Carlo performance results are given in Tables III–IV.

Encouragingly, for all scenarios we find improved performance with respect to all performance metrics of interest for modified MHT processing over standard

TABLE II
Simulation Parameters for the General Tracking Problem

Parameter Description	Setting (for scenario 1; 2; 3 respectively)
Monte Carlo realizations	100 (for each of the three scenarios)
Scenario duration	150 sec
Number of targets	3; 1; 3
Target birth (x, y) positions	$(-40,5),(-20,0),(-40,-5)$; $(-40,-5)$; $(-40,-5),(-35,-5),(-30,-5)$ m/sec
Target (x, y) velocities	$(0.5,0)$;$(0.5,0.5)$ or $(0.5,-0.5)$; $(0.5,0.5)$ or $(0.5,-0.5)$ in m/sec-all turns after 25 s
Target birth times from start	$(10,50,10)$; 10; $(10,10,10)$ sec
Target death times from start	$(140,140,100)$; 110; $(110,110,110)$ sec
Sensor footprint	2000 m^2
Sensor revisit rate	1 Hz
Sensor probability of detection	0.8
Sensor false alarm rate	5
Sensor measurement error covariance in x, y	1 m^2
Track initiation	6-of-6
Hypothesis tree depth (n_{scan})	1
Target birth rate	0.01
Target death probability	0.01
Confirmation reward factor (modified MHT)	2 (confirmed tracks), 1.5 (nearly-confirmed tracks)
Track termination (maximum missed detections)	2
Prior velocity covariance in x, y	1 m^2/s^2
Filter process noise in x, y	0.001 m^2/s^3
Data association gate	99%
Distance threshold for track-truth association	2

TABLE III
Performance Results for the Benchmark Scenarios
(Numerical results are based on 100 Monte Carlo realizations for each of the three scenarios.)

Scenario (description)	Tracker Modality	Track PD	Track Quality	Track Purity	Track Rate	Track Loc. Error
1 (3 linear)	standard MHT	0.915	0.850	0.780	1.90	0.825
1 (3 linear)	modified MHT	0.925	0.918	0.867	1.58	0.755
2 (1 maneuvering)	standard MHT	0.739	0.847	0.847	3.45	1.111
2 (1 maneuvering)	modified 3.45	1.111				
2 (1 maneuvering)	modified MHT	0.759	0.869	0.869	3.41	1.092
3 (3 maneuvering)	standard MHT	0.780	0.778	0.682	4.20	1.242
3 (3 maneuvering)	modified MHT	0.809	0.820	0.727	3.97	1.235

TABLE IV
Incremental Performance Benefit of Modified MHT, Averaged Across Scenarios, with Respect to all Metrics of Interest: Higher Track PD, Track Quality, and Track Purity; Lower Track Rate and Track Localization Error
(Numerical results accounts for all 300 Monte Carlo realizations.)

Metric	Track PD	Track Quality	Track Purity	Track Rate	Track Loc. Error
Percent change	2.43%	5.34%	8.10%	−6.09%	−3.02%

MHT. (Note that for track rate and track localization error, a reduction indicates improved performance.) Not surprisingly, since the scenarios are of increasing complexity we find consistently lower performance as we move from scenario 1 to scenario 2, and again from scenario 2 to scenario 3, as can be seen in the track quality, track rate, and track localization error. The one exception to the trend is track PD as we go from scenario 2 to scenario 3, though this can be explained: for multi-target scenarios, it is sufficient for a track instance by be close to *any* target to be deemed a detection, thus the presence of multiple nearby targets makes this easier to achieve.

Track purity is the same as track quality in the single-target scenario (scenario 2); in multi-target scenarios, track purity is lower than track quality as we require not only good track instances but from the same target as well. Indeed, track purity reflects the impact of track switching, whereby the target associated with a track may change over time. If no switching occurs, track purity equals track quality. Figure 6 illustrates what occasionally occurs, even in single-target settings: track fragmentation whereby the first track is seduced by false returns, and a second track is initiated. Note that our OSPA-based metrics correctly classify at most one track update as associated with each target at any time.

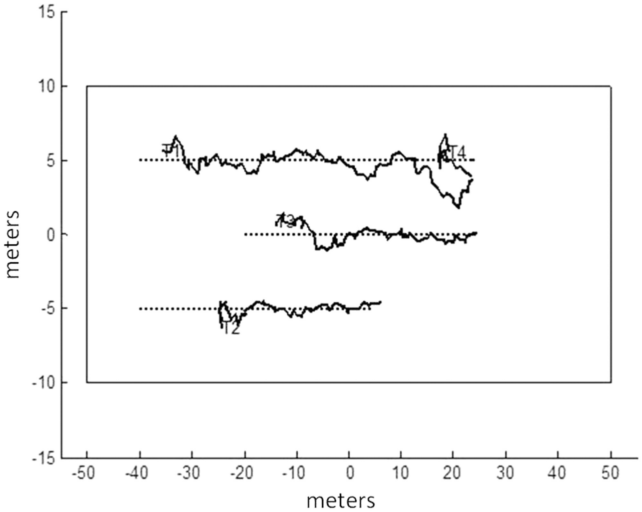

Fig. 5. An example realization of scenario 1 (target trajectories are dotted lines and modified MHT tracks are solid lines).

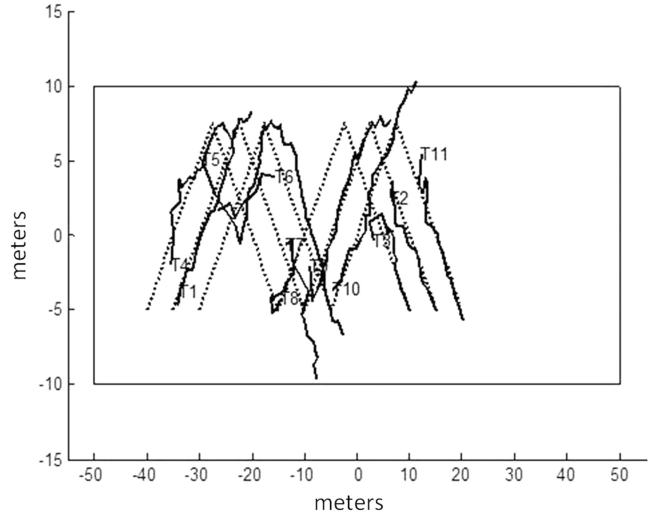

Fig. 7. An example realization of scenario 3 (target trajectories are dotted lines and modified MHT tracks are solid lines).

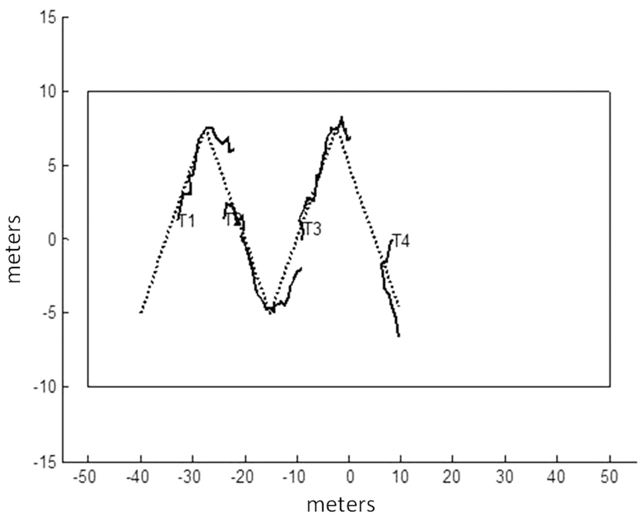

Fig. 6. An example realization of scenario 2 (target trajectories are dotted lines and modified MHT tracks are solid lines).

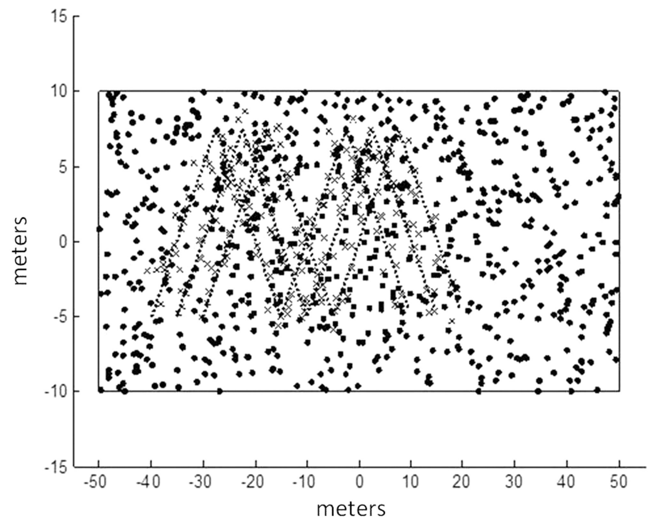

Fig. 8. Measurement data for one run of scenario 3 (crosses are target-originated returns and dots are false alarms).

TABLE V
Fusion Gain Computation (1.18 = 0.959/0.813)

Tracker Modality	IQ-scenario 1	IQ-scenario 2	IQ-scenario 3	IQ-average
Standard MHT	1.249	0.686	0.510	0.813
Modified MHT	1.610	0.729	0.532	0.959

Thus, the fragmentation and track redundancy observed here are reflected in degraded *track rate*, *track quality*, and *track purity* values.

It is helpful to capture tracker performance improvement with a single scalar metric, using the notion of *information quality* [4]; as discussed in [4], this notion can be related to the *information reduction factor* discussed in [15]. *Information quality* (IQ) is the average information content (in a Fisher information sense) of an arbitrarily-selected output track instance. With some probability, the track is associated with a true target: this is given by the *track quality* metric. Correspondingly, the information content is given by the Fisher information, which in turn is the inverse of the track error covariance matrix. Thus, IQ is the product of track quality and Fisher information:

$$\mathrm{IQ} = \frac{\sum_{i=1}^{N} g_i}{\sigma_T^2 \sum_{i=1}^{N} M_i}. \tag{22}$$

Further, we can evaluate *fusion gain* as the IQ ratio of two competing tracking solutions. For the results captured in Table III, we find that modified MHT provides a fusion gain of 1.18 (or 18%) over standard MHT. This is obtained by computing the ratio of average IQ for modified MHT and average IQ for standard MHT. Further details are in Table V. Note that, not surprisingly, for both tracker solutions the IQ metric degrades with increasing scenario difficulty.

6. CONCLUSIONS

This paper provides a compact and accessible introduction to track-oriented multiple-hypothesis tracking (MHT). It shows that track-oriented MHT does not suffer from the so-called *target-death problem* that has been observed in the probabilistic hypothesis density (PHD) filter. Unfortunately, the MHT exhibits a problem of its own, whereby unconfirmed tracks are often found to take contacts away from confirmed or nearly-confirmed tracks, degrading their quality.

We first study this problem in a simplified context with no measurement state information: this formulation reduces to the cardinality tracking problem. For this problem, and with the further assumption of unity detection probability, we are able to establish structural results for the optimal tracking solution and, remarkably, we find that the modified MHT solution with appropriately-selected track initiation and termination criteria and with sufficient hypothesis tree depth is guaranteed to achieve optimality. Simulation results are consistent with our theoretical findings.

The performance characteristics of modified-scoring MHT in the simplified cardinality-tracking context motivate its use for more general tracking problems. We study several benchmark scenarios and find improved performance for modified-scoring MHT over standard MHT. It is important to note that in all scenarios, all targets die before the scenario end; thus, there appear to be no negative side-effect to modified-scoring MHT processing whereby confirmed tracks are kept alive despite target death.

In a nutshell, modified MHT scoring is needed since we cannot perform batch track extraction from the set of track hypotheses. Indeed, batch extraction would incur computational infeasibility (unbounded n_{scan}) as well as large reporting latency. Accordingly, we must use (suboptimal) sequential track extraction. Favoring good (i.e. confirmed or nearly-confirmed) tracks over tentative ones in the extraction process can be motivated on two grounds: (1) empirically, as the scheme is found to perform better; (2) in a limiting (albeit simplistic) case—cardinality tracking, the scheme matches the performance of optimal batch extraction, provided n_{scan} is large enough (where the lower bound is quantified). While (1–2) do not *prove* that modified MHT is better than standard MHT, they do provide meaningful practical & theoretical motivation.

Our scheme is similar in its effects to the two-stage assignment scheme that has been adopted in an MHT setting [13–14]; indeed there is a need to balance track initiation and maintenance. A merit of our work, we think, is to emphasize an often-ignored aspect of making MHT work well in practice. Interestingly, giving preference to established tracks is a scheme whose applicability is not limited to the MHT approach; a recent example in the context of the *Histogram Probabilistic Multi-Hypothesis Tracker* (H-PMHT) may be found in [5].

REFERENCES

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

[2] S. Coraluppi and C. Carthel
Aggregate surveillance: A cardinality tracking approach.
In *Proceedings of the 14th International Conference on Information Fusion*, Chicago IL, July 2011.

[3] S. Coraluppi, C. Carthel, M. Luettgen, and S. Lynch
All-source track and identity fusion.
In *Proceedings of the National Symposium on Sensor and Data Fusion*, San Antonio TX, June 2000.

[4] S. Coraluppi, M. Guerriero, and C. Carthel
Fusion gain in multi-target tracking.
In *Proceedings of the 13th International Conference on Information Fusion*, Edinburgh, Scotland, July 2010.

[5] S. Davey
Histogram PMHT with particles.
In *Proceedings of the 14th International Conference on Information Fusion*, Chicago IL, July 2011.

[6] O. Erdinc, P. Willett, and Y. Bar-Shalom
The bin-occupancy filter and its connection to the PHD filters.
IEEE Transactions on Signal Processing, **57**, 11 (Nov. 2009).

[7] T. Kurien
Issues in the design of practical multitarget tracking algorithms.
In *Multitarget-Multisensor Tracking: Advanced Applications*, Y. Bar-Shalom (Ed.), Artech House, 1990.

[8] C. Morefield
Application of 0-1 integer programming to multitarget tracking problems.
IEEE Transactions on Automatic Control, **22**, 3 (June 1977).

[9] A. Poore and N. Rijavec
A Lagrangian relaxation algorithm for multidimensional assignment problems arising from multitarget tracking.
SIAM J. Optimization, **3**, 3 (Aug. 1993).

[10] D. Reid
An algorithm for tracking multiple targets.
IEEE Transactions on Automatic Control, **24**, 6 (Dec. 1979).

[11] D. Schuhmacher, B.-T. Vo, and B.-N. Vo
A consistent metric for performance evaluation of multi-object filters.
IEEE Transactions on Signal Processing, **56**, 8 (Aug. 2008).

[12] 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.

[13] A. Sinha, Z. Ding, T. Kirubarajan, and M. Farooq
Track quality based multitarget tracking algorithm.
To appear in *IEEE Transactions on Aerospace and Electronic Systems*.

[14] A. Sinha, Z. Ding, T. Kirubarajan, and M. Farooq
Track quality based multitarget tracking algorithm.
In *Proceedings of the SPIE Conference on Signal and Data Processing of Small Targets*, Orlando FL, Apr. 2006.

[15] X. Zhang, P. Willett, and Y. Bar-Shalom
Dynamic Cramer-Rao bound for target tracking in clutter.
IEEE Transactions on Aerospace and Electronic Systems, **41**, 4 (Oct. 2005).

Stefano Coraluppi received the B.S. degree in electrical engineering and mathematics from Carnegie Mellon University (1990), and M.S. and Ph.D. degrees in electrical engineering from the University of Maryland (1992, 1997).

Dr. Coraluppi has worked on the research staff at ALPHATECH Inc. (1997–2002), the NATO Undersea Research Centre (2002–2010), and Compunetix Inc. (since 2010). He has contributed to programs in ground, undersea, maritime and air surveillance for security and defense applications. In 2006, he was general cochair (with Professor Peter Willett) for the ISIF/IEEE 9th International Conference on Information Fusion in Florence, Italy. Currently, he serves on the Board of Directors of the International Society of Information Fusion (ISIF), for which he served as president in 2010. He is associate editor for target tracking and multisensor systems for the *IEEE Transactions on Aerospace and Electronic Systems* and area editor for tracking for the *ISIF Journal of Advances in Information Fusion*. His research interests include multi-target tracking, data fusion, detection and estimation theory, and optimal and stochastic control.

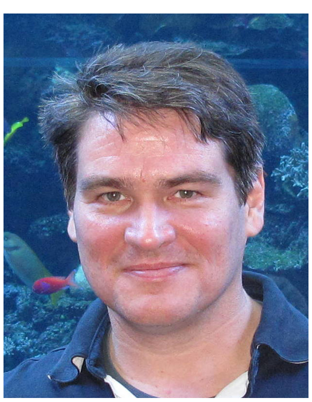

Craig Carthel received B.S. degrees in physics and mathematics in 1988, a M.S. in mathematics in 1992, and a Ph.D. in mathematics in 1995, all from the University of Houston.

He is a principal scientist at Compunetix Inc. in Monroeville, PA. He did research in numerical analysis and optimization theory at the University of Houston. From 1995 to 1997, he worked at the Institute for Industrial Mathematics at Johannes Kepler University, in Linz, Austria on parameter identification and inverse problems. From 1998 to 2002, he was a senior mathematician at ALPHATECH Inc. in Burlington, MA, where he worked on image processing, multisensor data fusion, and ground target tracking. From 2002 to 2010, he was a senior scientist in the Applied Research Department at the NATO Undersea Research Centre in La Spezia, Italy, where he worked on military operations research, simulation, optimization, and data fusion problems associated with maritime environments. In 2006, he served as the technical program chair for the 9th International Conference on Information Fusion.