Fundamentals and Advances in Multiple-Hypothesis Tracking

Stefano Coraluppi
600 West Cummings Park Suite 6500
Woburn, MA 01801
UNITED STATES

stefano.coraluppi@ieee.org

ABSTRACT

This manuscript describes the mathematical foundations of multiple-hypothesis tracking (MHT), a leading paradigm for multi-target tracking (MTT). We address aspects of track management, hypothesis pruning and aggregation, and the merits and limitations of centralized, distributed, and asynchronous processing for challenging multi-sensor surveillance applications. We extend the MHT formalism to the redundant-measurement setting. Finally, we derive a useful expression to assist parameter selection for MHT track management.

1.0 MULTI-TARGET TRACKING

This section introduces the *multi-target tracking* (MTT) problem and clarifies what is meant by *association-based* approaches to MTT, including the *multiple-hypothesis tracking* (MHT) approach. We address some concerns with association-based approaches that have been discussed in the literature.

1.1 Elements of Classical Estimation

Let X be a random variable of interest, and let Z be an observed random variable. We seek an *optimal estimate* $\hat{X}(Z)$ of X, based on its statistical prior as well as the observation Z. The *optimal estimator* $\hat{X}(\cdot)$ is the functional mapping from an observation Z to the corresponding estimate $\hat{X}$.

To speak of an optimal estimator requires that we specify the cost function with respect to which the estimator is best. Bayesian estimation prescribes that we consider the *Bayes risk* $J(\tilde{X})$; correspondingly, the optimal estimator is that $\tilde{X}(Z)$ that minimizes $J(\tilde{X})$. A simple manipulation shows that the optimal estimator minimizes the expected cost given any observation Z.

$$J(\tilde{X}) = E_X\left[E_Z\left[c\left(X, \tilde{X}(Z)\right)|X\right]\right] = E_{X,Z}\left[c\left(X, \tilde{X}(Z)\right)\right] = E_Z\left[E_X\left[c\left(X, \tilde{X}(Z)\right)|Z\right]\right]. \tag{1}$$

It is well known that optimal estimator $\hat{X}(\cdot)$ takes an appealing form for a number of choices of the cost $c(\cdot,\cdot)$. For simplicity, we identify these in the scalar-variable case. For $c(x,y) = (x-y)^2$, it can be shown that $\hat{X}(Z) = E[X|Z]$: the *minimum mean squared error* (MMSE) estimate is the conditional mean. For $c(x,y) = |x-y|$, it can be shown that $\hat{X}(Z) = M[X|Z]$: the *minimum mean absolute error* (MMAE) estimate is the conditional median. For $c(x,y) = 1, x-y > \delta$ (with $\delta \to 0$) and $c(x,y) = 0$ otherwise, it can be shown that $\hat{X}(Z) = \arg\max f(X|Z)$: the *minimum error* estimate, also known as the *maximum a posteriori* (MAP) estimate, is the conditional mode.

1.2 Multi-Target Estimation

In multi-target tracking, the variable of interest over a sequence of times $t^k = (t_1, \dots, t_k)$ is a set of trajectories that we denote by X^k. Each trajectory in this set has a time of birth, an evolution in target state space, and (possibly) a time of death. Hence, we are interest to identify the time evolution of an unknown (and time-varying) number of objects. We observe a sequence of sets of measurements Z^k. The usual simplifying assumption in the MTT problem formulation is that each target at each sensor measurement time gives rise to at most one measurement. However, it is not known which measurement originates from which object, and there are as well false measurements that are not target originated.

A difficulty in addressing the MTT problem is that it is not obvious how to define the Bayes risk for which we seek the optimal estimator. One could sidestep the issue and simply consider the posterior probability distribution $p(X^k|Z^k)$ and seek one of the optimal estimators noted above, e.g. the MAP estimator. Aside from the computational complexity associated with attempting such an operation, there is a conceptual difficulty in performing MAP estimation in this setting. This issue is discussed effectively in [1, pp. 494-500]. Essentially, it is problematic to compare values of the posterior probability distribution for choices of X^k that correspond to sets of objects with disparate cardinalities or temporal support. MAP estimation cannot meaningfully be performed.

One approach to resolving this conceptual difficult is to consider new types of estimators, as Mahler does with the *Marginal Multitarget* (MaM) estimator and *Joint Multitarget* (JoM) estimator. The interested reader is referred to [1] for details. Here, we note simply that the MaM estimator is conceptually similar to the MHT paradigm, which we will arrive to shortly.

1.3 Association-Based MTT

Another approach to resolving the conceptual difficulty noted above is explicitly to consider an explanation for the data, i.e. to specify which measurements are to be rejected as false and how target-originated measurements are to be associated. Let us denote by q^k one such global hypothesis or explanation. This leads to a probabilistic conditioning approach and the following expression for the multi-target posterior probability distribution $p(X^k|Z^k)$.

$$p(X^k|Z^k) = \sum_{q^k} p(X^k|Z^k, q^k) p(q^k|Z^k). \tag{2}$$

Unfortunately, it is not clear how eqn. (2) can be exploited to address the MTT problem. Further, it is worth noting that the space of global hypotheses is enormous. In fact, considering the possibility that not all targets will be detected over the time interval t^k, the set of global hypotheses is infinite.

The MHT paradigm resolves the conceptual and practical difficulties noted above. The former difficulty is addressed by focusing exclusively on $p(q^k|Z^k)$, and seeking the MAP estimate for q^k without facing the conceptual difficulty posed by continuous-valued spaces. The latter difficulty is generally addressed by neglecting undetected targets, and resolving hypotheses over small time horizons to bound computational complexity.

Thus, the MHT approach may be characterized as seeking the best explanation of the data, and then conditioning on this explanation to determine the continuous-space trajectories of interest. The latter task amounts to solving a set of nonlinear filtering problems, for which in the linear Gaussian case both MMSE and MAP estimators are given by the recursive Kalman filter; in the linear non-Gaussian case, the Kalman filter remains optimal among all linear estimators.

$$\hat{q}^k = \arg\max_{q^k} p\left(q^k | Z^k\right), \tag{3}$$

$$\hat{X}^k = \arg\max_{X^k} p\left(X^k | Z^k, \hat{q}^k\right). \tag{4}$$

1.4 Association-Based Sub-Optimality

It is worth emphasizing that association-based approaches like MHT do not directly optimize a criterion based on the multi-target posterior probability distribution $p\left(X^k|Z^k\right)$. Thus, while MHT seeks the MAP association hypothesis $\hat{q}^k$ – and this surely is a reasonable thing to do – there is no guarantee that selecting $\hat{q}^k$ will lead to optimal performance with respect to arbitrary MTT performance criteria based on $p\left(X^k|Z^k\right)$.

An interesting result in this respect is discussed in [3], where communication-constrained data association is considered. The results in [3] demonstrate that *statistical nearest neighbor* (SNN) association, i.e. the MAP solution to the single-target single-scan track maintenance problem, does not minimize the track localization error in sufficiently-high clutter environments. One can do better by following the data-association strategy detailed in the paper.

The implication of this research for our purposes is clear. In the general MTT setting, there is no guarantee that the choice of $\hat{q}^k$ leads to optimal MTT performance. Nonetheless, in practice (and not surprisingly) we find that selection of the MAP association hypothesis is generally a good thing to do.

1.5 Objections to Association-Based MTT

In [2, pp. 8-9], Vo presents a lucid discussion of potential conceptual difficulties in association-based MTT. His focus is on the problematics associated with applying Bayes rule to manipulate $p\left(q^k|Z^k\right)$. In particular, Bayes rule prescribes the following:

$$p\left(q^k|Z^k\right) = \frac{p\left(Z^k|q^k\right)p\left(q^k\right)}{p\left(Z^k\right)}. \tag{5}$$

Vo rightly observes that the use of Bayes rule in this setting, while seemingly benign, does raise some concerns. Since q^k depends on Z^k as it prescribes how to explain the data, is $p\left(q^k\right)$ a valid prior? Likewise, is $p\left(Z^k|q^k\right)$ a valid likelihood function?

The use of Bayes rule in this setting requires a conditioning argument for it to be valid. To our knowledge, this point has been overlooked in the MHT community to date. In particular, we must proceed as follows, where $\left|Z^k\right|$ is the sequence of measurement set cardinality for the time sequence t^k.

$$p\left(q^k|Z^k\right) = p\left(q^k|Z^k, \left|Z^k\right|\right) = \frac{p\left(Z^k|q^k, \left|Z^k\right|\right)p\left(q^k|\left|Z^k\right|\right)}{p\left(Z^k|\left|Z^k\right|\right)}. \tag{6}$$

Note that q^k is conditionally independent of Z^k given $|Z^k|$, hence $p(q^k||Z^k|)$ is now a valid prior and $p(Z^k|q^k,|Z^k|)$ is now a valid likelihood function.

Referring to the numerator in (5), Vo notes further that it is not clear whether $p(Z^k|q^k)p(q^k)$ is the joint density $p(Z^k,q^k)$. Indeed, the marginal $p(q^k)$ that result by integration of $p(Z^k,q^k)$ over Z^k according to (7) must not depend on the data, contracting the fact that q^k does depend on Z^k. In (7), we denote by Z the dummy integration variable. Note that (7) neglects the fact that the integration can only meaningfully be performed over those measurement sets that are consistent with q^k.

$$p(q^k) = \int_Z p(Z|q^k)p(q^k)dZ. \tag{7}$$

Once more, we can resolve this difficulty by considering instead the joint density conditioned on a given $|Z^k|$. Referring to the numerator in (6), we have:

$$p(q^k||Z^k|) = \int_{Z,|Z|=|Z^k|} p(Z|q^k,|Z^k|)p(q^k||Z^k|)dZ \tag{8}$$

Finally, the normalizing constant must again be understood as integration over all global hypotheses consistent with a given measurement cardinality. This resolves the concern raised by Vo as to whether the normalizing constant even exists. Note that q^k is discrete-valued, hence (9) is simply a sum.

$$p(Z^k||Z^k|) = \int_{q^k \text{ consistent with } |Z^k|} p(Z^k|q^k,|Z^k|)p(q^k||Z^k|)dq^k. \tag{9}$$

In [2], Vo provides a valuable contribution to the MTT literature by raising the concerns noted here. These concerns lead to a more rigorous derivation of the MHT recursion, as described below.

2.0 MULTIPLE-HYPOTHESIS TRACKING

This section introduces the fundamental recursion that is utilized in MHT; to our knowledge, the exact form considered here has not appeared in the literature, and relies on insight from the discussion in Section 1.5. The resulting track-oriented MHT recursion is well known. Additionally, we touch upon additional aspects including track management, hypothesis aggregation, and processing architectures.

2.1 Recursive Formulation of Global Hypotheses

Computational and real-time constraints require that we adopt a recursive formulation of (6). Thus, we proceed as follows.

$$p(q^k|Z^k) = \frac{p(q^k|Z^k)}{p(|Z_k||Z^k)} = p(q^k|Z^k,|Z_k|) = \frac{p(Z_k|Z^{k-1},q^k,|Z_k|)p(q^k|Z^{k-1},|Z_k|)}{p(Z_k|Z^{k-1},|Z_k|)}. \tag{10}$$

We consider in turn each of the factors in (10). Noting that $|Z_k|$ is known given q^k, the first numerator factor in (10) may be manipulated as follows:

$$p\left(Z_k|Z^{k-1},q^k,|Z_k|\right)=p\left(Z_k|Z^{k-1},q^k,|Z_k|\right)p\left(|Z_k||Z^{k-1},q^k\right)$$
$$=p\left(Z_k,|Z_k||Z^{k-1},q^k\right)=p\left(Z_k|Z^{k-1},q^k\right). \tag{11}$$

The second numerator factor in (10) may be manipulated as follows:

$$p\left(q^k|Z^{k-1},|Z_k|\right)=p\left(q_k|Z^{k-1},|Z_k|,q^{k-1}\right)p\left(q^{k-1}|Z^{k-1},|Z_k|\right)$$
$$=\frac{p\left(q_k|Z^{k-1},|Z_k|,q^{k-1}\right)p\left(q^{k-1}|Z^{k-1}\right)}{p\left(|Z_k||Z^{k-1},\right)}. \tag{12}$$

The denominator in (10) may be manipulated as follows:

$$p\left(Z_k|Z^{k-1},|Z_k|\right)=\frac{p\left(Z_k|Z^{k-1}\right)}{p\left(|Z_k||Z^{k-1}\right)}. \tag{13}$$

Combining (11-13) according to (10) yields the following:

$$p\left(q^k|Z^k\right)=\frac{p\left(Z_k|Z^{k-1},q^k\right)p\left(q_k|Z^{k-1},|Z_k|,q^{k-1}\right)p\left(q^{k-1}|Z^{k-1}\right)}{p\left(Z_k|Z^{k-1}\right)}. \tag{14}$$

This is the global hypothesis recursion that expresses $p\left(q^k|Z^k\right)$ as a function of $p\left(q^{k-1}|Z^{k-1}\right)$ and the current scan of data Z_k. The noteworthy numerator factor is $p\left(q_k|Z^{k-1},|Z_k|,q^{k-1}\right)$: in the literature, this factor is erroneously identified as $p\left(q_k|Z^{k-1},q^{k-1}\right)$ [4].

2.2 Track-Oriented MHT

Though useful, the recursion (14) is generally intractable in the sense that the space of global hypotheses is quite large. Fortunately, under some simplifying assumptions, namely Poisson-distributed number of target births and number of false alarms at each scan, the posterior probability of a global hypothesis $p\left(q^k|Z^k\right)$ may be expressed as a product over local (or *track*) hypotheses associated with q^k. This fundamental contribution to the MHT literature is derived in [4].

The Poisson assumptions above are quite reasonable in many settings. Indeed, consider a continuous-time birth-death process with exponentially-distributed target inter-arrival (birth) times with parameter λ_b, and exponentially distributed target lifetime with parameter λ_χ. Discrete-time statistics may be readily obtained, leading to a Poisson distributed number of births with mean $\mu_b(t)$ and death probability $p_\chi(t)$ over an interval of duration t. The expressions are given in equations (15-16).

$$\mu_b(t)=\frac{\lambda_b}{\lambda_\chi}\left(1-e^{-\lambda_\chi t}\right), \tag{15}$$

$$p_\chi(t)=1-e^{-\lambda_\chi t}\,. \tag{16}$$

For simplicity, in the following we will omit the time interval t and use the birth rate and death probability μ_b and p_χ, respectively. (Time arguments should be noted explicitly when the sensor revisit interval is time-varying.)

Similarly, the Poisson false alarm assumption (with mean Λ) is a reasonable one as it matches clutter statistics in many application domains. It results as a limiting case of the Binomial distribution with a large number of detection cells N and vanishingly small false detection probability p_F, with $p_F \cdot N \to \Lambda$. We assume that at every scan, each target is detected with probability p_d.

Let τ be the number of targets in global hypothesis q^{k-1} at time t^{k-1}, $r = |Z_k|$ be the number of measurements in the current scan at time t^k, and b, χ, and d are the number of target births, deaths, and measurement updates in global hypothesis q^k at time t^k, respectively. Note that the classical MHT only considers global hypotheses for which targets are detected at birth.

We now express the global hypothesis recursion (14) explicitly. It can be shown that the factor $p(q_k|Z^{k-1}, |Z_k|, q^{k-1})$ may be written as follows:

$$p(q_k|Z^{k-1}, |Z_k|, q^{k-1}) = \left\{\frac{\exp(-p_d\mu_b-\Lambda)\Lambda^r}{r!}\right\} p_\chi^\chi \left((1-p_\chi)(1-p_d)\right)^{\tau-\chi-d} \left(\frac{(1-p_\chi)p_d}{\Lambda}\right)^d \left(\frac{p_d\mu_b}{\Lambda}\right)^b. \quad (17)$$

The factor $p(Z_k|Z^{k-1}, q^k)$ in (14) accounts for the probability of observing a set of measurements given a global hypothesis. It is simply a product over filter residual scores; hence, it may be written as follows, where, under q_k, J_d is the set of track update measurements, J_{fa} is the set of false alarms, J_b is the set of target birth measurements, and $f(\cdot)$ is the filter score.

$$p(Z_k|Z^{k-1}, q^k) = \prod_{j\in J_d} f(z_j|Z^{k-1}, q^k) \prod_{j\in J_b} f(z_j|Z^{k-1}, q^k) \prod_{j\in J_{fa}} f(z_j|Z^{k-1}, q^k). \quad (18)$$

Equations (17-18) may be combined into (14), resulting in the following track-oriented MHT recursion. Equation (19) is of fundamental importance in that it factors global hypothesis scores into track scores. This allows the recursive determination of $\hat{q}^k$ as the solution to an integer programming problem, without requiring explicit enumeration of global hypotheses. Note that $f_{fa}(\cdot)$ is the false alarm distribution.

$$p(q^k|Z^k) = p_\chi^\chi \left((1-p_\chi)(1-p_d)\right)^{\tau-\chi-d} \prod_{j\in J_d} \frac{(1-p_\chi)p_d f(z_j|Z^{k-1}, q^k)}{\Lambda f_{fa}(z_j)} \prod_{j\in J_b} \frac{p_d\mu_b f(z_j|Z^{k-1}, q^k)}{\Lambda f_{fa}(z_j)}$$

$$\cdot \frac{\left\{\frac{\exp(-p_d\mu_b-\Lambda)\Lambda^r}{r!}\right\} \prod_{j\in Z_k} f(z_j|Z^{k-1}, q^k)}{p(Z_k|Z^{k-1})} p(q^{k-1}|Z^{k-1}). \quad (19)$$

2.3 Hypothesis Pruning and Track Management

Though useful, the recursion (19) still is insufficient for viable MHT processing. Indeed, in principle one must form all track hypotheses over a temporal batch of data followed by solution to an optimization problem that results in $\hat{q}^k$. This incurs unacceptable computational expense and solution latency for all but very small surveillance problems.

Hypothesis pruning allows tractable computational expense. Effective pruning schemes exist, based on reduction to a single global hypothesis with a bounded temporal delay; this enables both reduced computations and real-time processing. A straightforward solution to the integer programming problem is via *linear programming* (LP) relaxation; this was studied independently in [5-6].

Optimal processing in principle requires full hypothesis formation (with no hypothesis pruning) as well as track extraction as a single processing step. As mentioned above, hypothesis pruning is necessarily required. Further, track extraction is generally performed only for resolved hypotheses. Thus, hypothesis resolution and track management are generally decoupled in most MHT implementations. Correspondingly, hypothesis resolution based on LP relaxation employs equality constraints; that is, all sensor measurements are accounted for in all global hypotheses [5].

While suboptimal, the use of distinct hypothesis resolution and track extraction functions offers processing advantages. In particular, we note that confirmed tracks may be favored in data-association processing; see [7-8] for an analysis of advantages resulting from feedback processing from track extraction to data association functions.

2.4 Hypothesis Aggregation

Generally, we do not distinguish between the data-association hypothesis $\hat{q}^k$ and global hypotheses that are consistent with it. In general there any many data-indistinguishable global hypotheses associated with the same data association hypothesis. In particular, multiple target birth and death times are possible, and there may as well be targets with no associated detections. The merits of considering a larger hypothesis space and the ability to do so without incurring additional computation expense are discussed in [9].

Hypothesis aggregation for data-indistinguishable hypotheses takes at least two forms. One involves aggregating over all target birth and death times to result in a single (aggregated) track hypothesis for a single associated-measurement sequence [10]. The other involves aggregation over indistinguishable sensor measurements, as will occur in cardinality-estimation applications [11]. Aggregation over similar (but not data-indistinguishable) hypotheses may also be performed with appreciable benefits [12].

2.5 Centralized, Distributed, and Asynchronous Processing

While centralized fusion provides excellent performance in many settings, effective exploitation of multi-sensor data with good performance and robustness characteristics often requires advanced processing architectures [13]. Fading detection statistics and sensor registration errors are best handled in distributed architectures. In addition to improved robustness characteristics, multi-stage data association provides an effective means to handling disparate sensor update rates and to exploit same-sensor association performance [14-15].

While the success of distributed processing solutions over centralized processing may give pause to those familiar with detection and estimation theory and the optimality results associated with centralized solutions, we must recognize that the MTT problem is exceedingly complex. Hence, the choice is between suboptimal centralized solutions and suboptimal distributed solutions. Hence, distributed processing must be viewed as a flexible approach to suboptimal but effective surveillance solutions.

Similarly, surprising results have been shown recently regarding the value of asynchronous processing in forensic settings to content with disparate data sources where the low-rate sensor is highly informative [16]. In such settings, the purposeful use of out-of-sequence processing enables effective MHT solutions that are impossible to achieve in time-sequential processing.

3.0 REDUNDANT-MEASUREMENT MHT

This section extends the MHT paradigm to allow for multiple measurements per target from a single sensor scan. In this context we discuss practical processing considerations, the nature of the track-repulsion effect, and the benefits of processing feedback.

3.1 Redundant-Measurement MTT

MTT with redundant measurements poses a significant challenge. For simplicity, most paradigms adapt a Bernoulli measurement model. There are some exceptions, e.g. the *probabilistic MHT* (PMHT) and its *non-generative* sensor model [17]. A complementary difficulty – merged measurements due to more than one target – also is not considered in most MTT treatments.

Redundant measurements induced by multipath phenomena or multiple emissions have been addressed in an MHT setting; see [18-19] and references therein. However, while these papers are of interest, they do not address the challenging problem considered here, where all redundant measurements are characterized by the *same* measurement equation. A recent treatment of redundant measurement in the context of *probability hypothesis density* (PHD) research is discussed in [20-22]. Both merged and redundant measurements are addressed using a *Markov Chain Monte Carlo* (MCMC) approach in [23], and in [24-25] with the *probabilistic data association filter* (PDAF).

3.2 The MHT Recursion

We are interested to establish a generalization to (17) that relaxes the Bernoulli measurement cardinality assumption. Thus, target birth and target update hypotheses now allow for an arbitrary number of measurements. We denote by $p(\cdot)$ the measurement-cardinality distribution; this replaces the Bernoulli measurement cardinality that is specified by the single parameter p_d. We consider the auxiliary (cardinality) variable ψ_k that specifies the number of births for each measurement-cluster cardinality, the number of target deaths, and the number of targets with measurement update for each measurement-cluster cardinality. In particular, τ is the number of targets at time t^{k-1}, b_i is the number of target births with a corresponding measurement cluster of size i; χ is the number of target deaths; and d_i is the number of targets with measurement update with a cluster of size i (note that this includes the target missed detection case, for which $i=0$). We consider n false alarms, with each a singleton cluster; i.e. $n_1 = n$.

Conditioning on ψ_k, $p\left(q_k|Z^{k-1},|Z_k|,q^{k-1}\right)$ can be expressed as follows:

$$p\left(q_k|Z^{k-1},|Z_k|,q^{k-1}\right) = p\left(\psi_k|Z^{k-1},|Z_k|,q^{k-1}\right)p\left(q_k|Z^{k-1},|Z_k|,q^{k-1},\psi_k\right), \tag{20}$$

$$p\left(\psi_k|Z^{k-1},|Z_k|,q^{k-1}\right) = \left\{\binom{\tau}{\chi}p_\chi^\chi\left(1-p_\chi\right)^{\tau-\chi}\right\}\left\{\frac{(\tau-\chi)!}{\prod_i d_i!}\prod_i p(i)^{d_i}\right\}\left\{\frac{\Lambda^{n_1}e^{-\Lambda}}{n_1!}\right\}\left\{\prod_i \frac{e^{-p(i)\mu_b}p(i)^{b_i}\mu_b^{b_i}}{b_i!}\right\}, \tag{21}$$

$$p\left(q_k|Z^{k-1},|Z_k|,q^{k-1},\psi_k\right) = \frac{1}{\left\{\binom{\tau}{\chi}\left(\frac{(\tau-\chi)!}{\prod_i d_i!}\right)\left(\prod_i\frac{(d_i+b_i+n_i)!}{(b_i+n_i)!}\right)\left(\frac{r!}{\prod_i(d_i+b_i+n_i)!}\right)\prod_i\binom{b_i+n_i}{b_i}\right\}}. \tag{22}$$

$$n_1 + \sum_i ib_i + \sum_i id_i = r, \tag{23}$$

$$\sum_i d_i = \tau - \chi. \tag{24}$$

Notes:

- The factors in (21) represent, respectively: the probability of χ deaths among τ tracks; the probability of measurement cardinalities according to d_i for surviving tracks, using the multinomial distribution; the probability of n_1 false alarms; and the probability of birth cardinalities according to b_i, using Poisson sifting.
- The denominator factors in (22) represent, respectively: the number of ways to select track terminations; the number of ways of selecting tracks for specific cardinality updates; the number of ways of assigning measurement clusters to tracks; the number of ways of assigning measurements to clusters; and the number of ways of selecting birth clusters.
- The total number of returns is r according to (23); the total number of update clusters is $\tau - \chi$ according to (24.).
- All products and summations are over $i = 0, \ldots$ Undetected births can be accounted for via generalized birth statistics [10]. Thus w.l.o.g. we have $b_0 = 0$.

Simplifying (20-24), the resulting generalization to (17) is given below. We note that an earlier derivation is given in [10], but with an error due to the mistaken assumption that ordering does not matter in assigning measurements to clusters.

$$p\left(q_k | Z^{k-1}, q^{k-1}\right) = \left\{\frac{\Lambda^r e^{-\Lambda} e^{-\mu_b}}{r!}\right\} p_\chi^\chi \prod_i \left(\frac{p(i)}{\Lambda^i}\left(1 - p_\chi\right)\right)^{d_i} \prod_i \left(\frac{p(i)}{\Lambda^i}\mu_b\right)^{b_i}. \tag{25}$$

This result is significant in that hypothesis factorization is achieved in a more general setting than previously established, so that track-oriented MHT may be adopted in the redundant-measurement setting. Further, the local track scores take an intuitive form, directly generalizing what is in (17) to the redundant-measurement case.

A special case of interest that is applicable in several domains results from a Poisson measurement-cardinality assumption, whereby we have:

$$p(i) = \frac{\lambda^i}{i!} e^{-\lambda}. \tag{26}$$

Under this special case, (25) takes the following form:

$$p\left(q_k | Z^{k-1}, q^{k-1}\right) = \left\{\frac{\Lambda^r e^{-\Lambda} e^{-\mu_b}}{r!}\right\} p_\chi^\chi \prod_i \left(\frac{\lambda^i}{i!\Lambda^i} e^{-\lambda}\left(1 - p_\chi\right)\right)^{d_i} \prod_i \left(\frac{\lambda^i}{i!\Lambda^i} e^{-\lambda}\mu_b\right)^{b_i}. \tag{27}$$

3.3 Two-Stage Processing

The principled derivation leading to (25) is important. However, in practical settings, we must recognize that redundant-measurement MHT – or MTT of any sort – poses a significant computational challenge. A first simplification may be achieved by decoupling the measurement-clustering and measurement-to-track association stages. Alternatively, one could consider target tracking (under the usual Bernoulli measurement-cardinality model) followed by track clustering. With either approach, it is worth noting that, even for a single-sensor problem, multi-stage MHT processing can provide benefits over single-stage (centralized) processing. An excellent, practical illustration in passive bistatic radar tracking with digital audio broadcast data may be found in [26].

Even with upfront measurement clustering, the number of ways to cluster N measurements is given by the Bell number B_N, which grows roughly as $O(N^N)$. Thus, modified, suboptimal (heuristic) measurement-clustering schemes are necessary in large-scale settings.

In the simplified setting of 1D measurement data, optimal measurement clustering is achievable, and does lead to improved detection statistics under the *Optimal Sub-pattern Assignment* (OSPA) metric [27]. As a simple illustration, Figure 3-1 illustrates on the left one realization of target existence (Poisson number of targets, Gaussian 1D positional prior), measurement generation (under the Poisson measurement-cardinality assumption), and optimal clustering. In this 1D setting, it can be shown that it is sufficient to consider 2^N ways to cluster the measurements. On the right, the benefit of measurement clustering with respect to the OSPA metric is shown, based on 100 Monte Carlo realizations of measurement data.

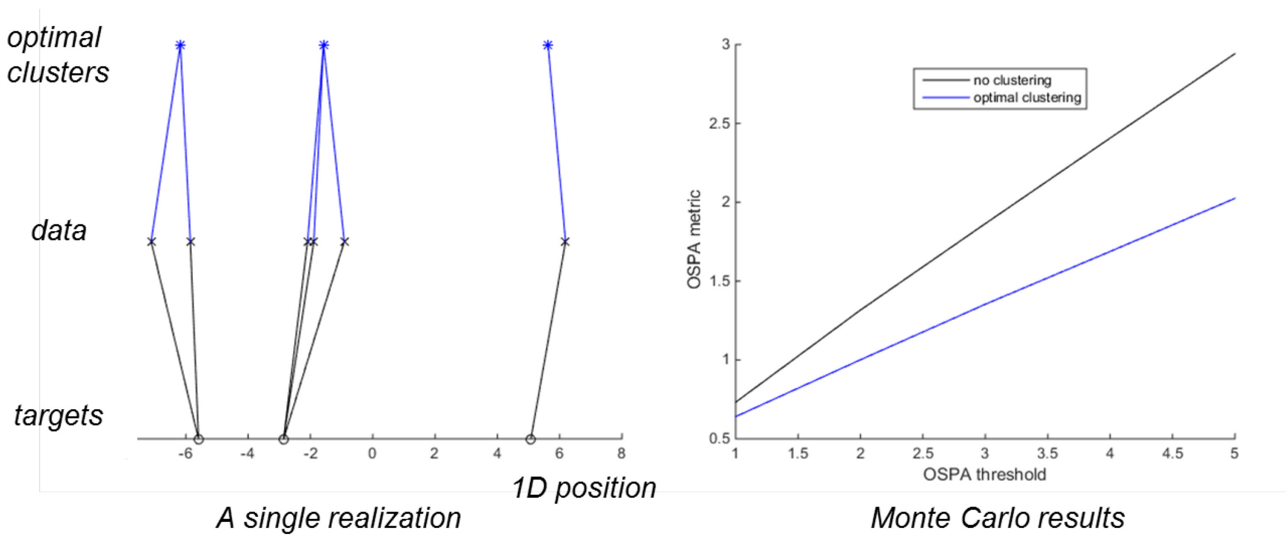

Figure 3-1: Optimal measurement clustering of 1D measurement data.

Ultimately, we expect that the sub-optimality associated with decoupled measurement clustering and target tracking will be mitigated by feedback that informs the clustering process. That is, rather than considering a diffuse prior for both the number of targets and their location when performing measurement clustering, improved clustering may be achieved with additional knowledge of the number and location of targets. This is illustrated in Figure 3-2.

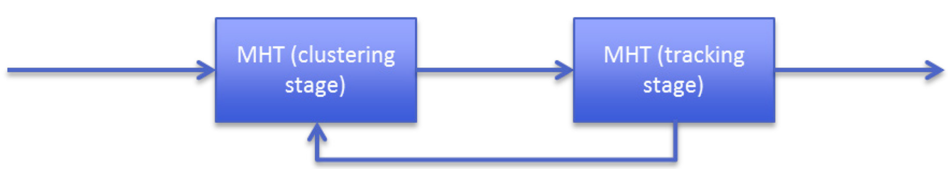

Figure 3-2: A practical two-stage approach to redundant-measurement MHT.

An interesting limiting case of the redundant-measurement tracking problem remains to be considered: the cardinality-estimation problem with indistinguishable measurements [11]. Under the Poisson assumption for both the number of targets and the number of measurements per target, a compound Poisson distribution results for the overall number of measurements per scan. Recent analysis of cardinality-estimation performance (under the simpler Bernoulli measurement-cardinality assumption) for a naïve Kalman filtering solution and for the *Cardinalized Probability Hypothesis Density* (CPHD) filter may be found in [28].

3.4 Track Repulsion

The *track-coalescence effect* degrades the performance of probabilistic data association trackers in dense-target scenarios [29]; attempts at sub-optimal Bayesian processing to combat this effect have been reported [30]. It has been observed that an opposite effect exists with trackers that utilize hard data association, which we denote as the *track-repulsion effect* [31]; the effect can be mitigated with multi-stage multi-hypothesis tracking [32-33]. A recently introduced algorithm, known as *Set JPDA* (SJPDA), successfully overcomes both coalescence and repulsion effects but at the cost of track labeling [34].

We now explore the question of whether track repulsion is greater in the presence of redundant measurements. On the one hand, an increased number of measurements provides additional information that ought to reduce the effect. On the other hand, the nature of the optimal clustering solution may increase the effect, since measurement outliers from one target might erroneously be clustered with measurements from a nearby target.

We consider a simple static problem. Two target defined by 1D positional states are displaced by a given distance, and give rise to m measurements (we consider up to six). As a function of target displacement and number of measurements, we can compute the mean OSPA (MOSPA) error that results from optimal clustering, whereby the m smallest measured values are clustered and the m largest values are clustered. The resulting mean measurement-clustering error is shown in Figure 3-3, based on a 1m measurement standard deviation error.

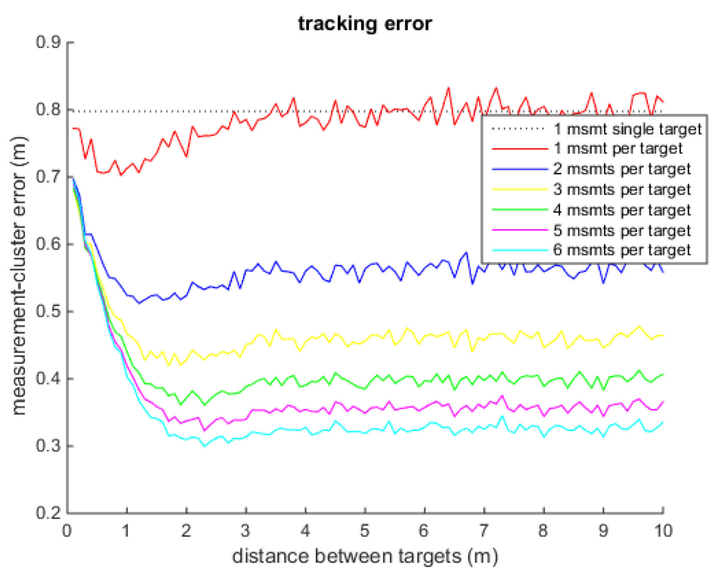

Figure 3-3: Average localization error as a function of target displacement and measurement cardinality, based on 1000 Monte Carlo realizations of measurement data.

There are several points to note. First, for distant targets, localization error improves with increased measurement cardinality. This effect is obvious, since more information is gained regarding target state and measurement-association errors are minimal. Indeed, the distant-target MOSPA error matches the single-target average localization error (black dotted line). For targets that are closer, the MOSPA error is smaller, since the metric exploits optimal cluster-to-target mapping. Hence, in this sense there is a localization benefit observed in closely-spaced target settings, regardless of measurement cardinality. On the other hand and quite interestingly, for very close targets, the MOSPA error is seen to grow. This illustrates the track repulsion effect. Indeed, when the two targets are co-located, there is a nontrivial displacement error for both the left cluster (smaller measurement values) and right cluster (larger measurement values) with respect to the true target locations.

While the track repulsion effect may not appear overly pronounced in Figure 3-3, in the sense that the total tracking error remains below that of the distant-target case, it is useful to consider the multi-scan case where a sequence of measurement sets exist. For the static target case, global nearest neighbour data association (and, in fact, MHT processing) dictates that all left clusters observed at multiple scans be fused, and that all right clusters be fused. The resulting steady-state tracking error removes the random error component and identifies the residual bias error. This is illustrated in Figure 3-4, where measurement clusters are fused over 100 scans. There is virtually no steady-state bias error for distant targets, and the effect is most pronounced for nearby targets.

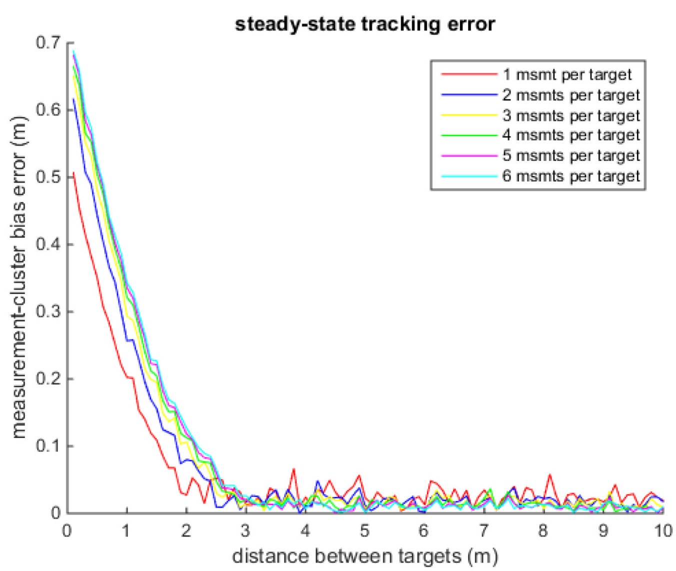

Figure 3-4: Steady-state localization error as a function of target displacement and measurement cardinality, based on 1000 Monte Carlo realizations of measurement data.

Perhaps most interestingly, this bias error is more pronounced with increasing measurement cardinality: more measurements on the targets exacerbate the track-repulsion effect. This is shown in the close-up view of Figure 3-5. The effect can be explained readily. The measurement-aggregation step in the non-unity m-measurement cardinality cases considers a larger hypothesis space than in the unity-cardinality scan case. The grouping of all smaller measurements *increases* the probability of erroneous measurement assignment over what would occur if the measurement scan were treated as a sequence of m unity-cardinality scans.

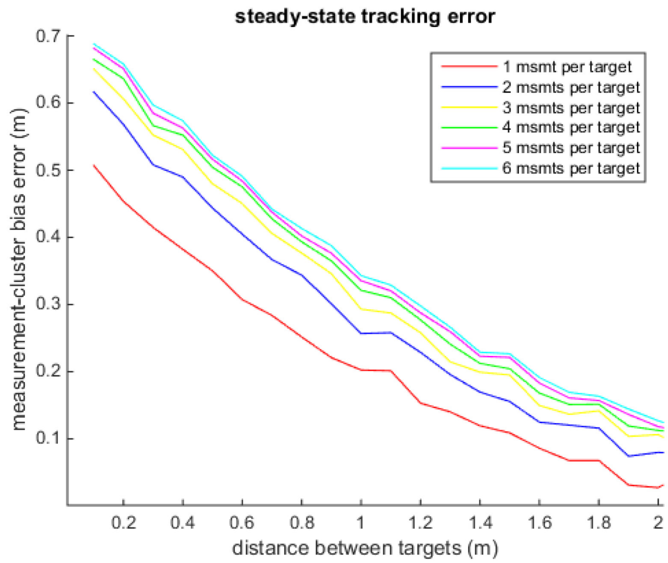

Figure 3-5: Steady-state localization error.
(Close-up view of previous figure, for small target displacement).

The statements above are best understood with an illustrative example. At the top of Figure 3-6, we see five scans of unity-cardinality returns for two targets. The returns shown in red will be associated and fused; these are shown in the middle portion of Figure 3-6. Conversely, if the same measurements originate from a single scan with five returns per target, the aggregation decision shown in the bottom portion of Figure 3-6 will result. For closely-spaced targets, the latter association approach will tend to increase the track-repulsion effect.

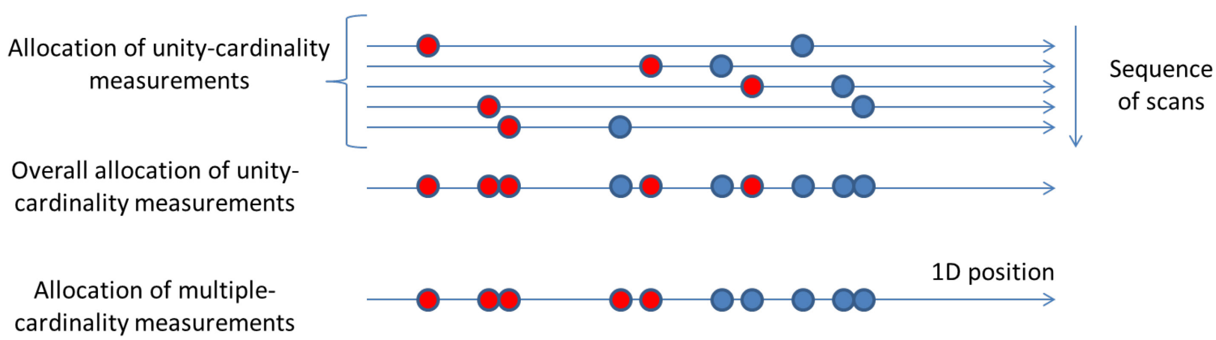

Figure 3-6: Increased track repulsion due to clustering of redundant-return measurements (bottom) relative to processing a sequence of unity-return measurement scans (top and middle). Red measurements are associated.

4.0 TRACKER PERFORMANCE MODELING

Performance models and bounds for MTT systems have been documented in the literature [13]. Typically, such analyses rely on simplifying approximations to yield tractable analysis. In this section, we present a temporal sliding-window generalization to the well-known detection statistics associated with *M-of-N* fusion logic [35]. This generalization is useful to estimate logic-based track-confirmation and track-duration statistics as a function of track-management parameters. For simplicity, we assume track confirmation on *N* consecutive detections and track termination on K consecutive missed detections. We assume target existence during the surveillance interval of T scans.

4.1 Problem Notation

N = number of (consecutive) detections to confirm

K = number of allowed (consecutive) misses

p = detection probability

T = scenario duration (number of frames)

μ = expected track duration (beyond confirmation)

τ_μ = expected track duration (total)

ε = expected track formation time (before confirmation)

τ_ε = expected track formation time (total)

Q = expected target completeness

λ = expected fragmentation rate

F = expected number of track fragments per target

4.2 Model-Based Performance

A useful identity is given by the following:

$$f(x,n) = \sum_{i=0}^{n} i \cdot x^{i-1} = \frac{n \cdot x^{n+1} - (n+1)x^n + 1}{(1-x)^2}. \tag{28}$$

Equation (28) may be established easily as shown below:

$$F(x,n) = \int_0^x f(\xi,n)d\xi = \sum_{i=0}^{n} \int_0^x i \cdot \xi^{i-1} d\xi = \sum_{i=0}^{n} x^i = \frac{1-x^{n+1}}{1-x}$$

$$f(x,n) = \frac{dF(x,n)}{dx} = \frac{-(n+1)x^n}{1-x} + \frac{1-x^{n+1}}{(1-x)^2} = \frac{(n+1)x^{n+1} - (n+1)x^n + 1 - x^{n+1}}{(1-x)^2}$$

$$= \frac{n \cdot x^{n+1} - (n+1)x^n + 1}{(1-x)^2}$$

We now estimate expected track duration (beyond confirmation) as follows:

$$\mu = \sum_{i=0}^{K} (1-p)^i p(i+1+\mu) + (1-p)^{K+1}K = p\mu \sum_{i=0}^{K} (1-p)^i + p \sum_{i=0}^{K} (1-p)^i (i+1) + (1-p)^{K+1}K$$

$$= p\mu \frac{1-(1-p)^{K+1}}{1-(1-p)} + p\sum_{j=1}^{K+1}(1-p)^{j-1}j + (1-p)^{K+1}K$$

$$= \mu(1-(1-p)^{K+1}) + p\sum_{j=0}^{K+1}(1-p)^{j-1}j + (1-p)^{K+1}K$$

$$= \mu(1-(1-p)^{K+1}) + \frac{(K+1)\cdot(1-p)^{K+2}-(K+2)(1-p)^{K+1}+1}{p} + (1-p)^{K+1}K$$

Hence, we have:

$$\mu = \frac{(K+1)\cdot(1-p)^{K+2}-(K+2)(1-p)^{K+1}+1}{p(1-p)^{K+1}} + K. \tag{29}$$

Special cases associated with (29) are the following:

1. $K = 0$: $\mu = \frac{p}{1-p}$.
2. $p \to 1$: $\mu \to \infty$.
3. $p \to 0$: $\mu \to \frac{(K+1)\left(1-(K+2)p+\frac{(K+2)(K+1)}{2}p^2\right)-(K+2)\left(1-(K+1)p+\frac{(K+1)K}{2}p^2\right)}{p} + K = \frac{(K+1)(K+2)}{2}p + K$.

Expected track duration (total) is given by:

$$\tau_\mu = 1 + \mu. \tag{30}$$

Expected track formation time (before confirmation) can be derived in the same manner as μ, with $p \leftrightarrow 1-p$, $N-1 \leftrightarrow K$. This results in the following:

$$\varepsilon = \sum_{i=0}^{N-1} p^i(1-p)(i+1+\tau_c) + p^N(N-1),$$

$$\varepsilon = \frac{N\cdot p^{N+1}-(N+1)p^N+1}{(1-p)p^N} + N - 1. \tag{31}$$

Expected track formation time (total) is given by:

$$\tau_\varepsilon = 1 + \varepsilon. \tag{32}$$

Expected target completeness may be inferred from track duration and track formation times:

$$Q = \frac{\tau_\mu}{\tau_\mu+\tau_\varepsilon}. \tag{33}$$

Finally, the expected track fragmentation rate and expected number of track fragments per target may be estimated as well:

$$\lambda = \frac{1}{\tau_\mu+\tau_\varepsilon}, \tag{34}$$

$$F = \lambda \cdot T. \tag{35}$$

4.3 An Example

As a simple illustration of the performance model derived above, consider a sensor with detection probability $p = 0.933$, a scenario duration of $\mathrm{T} = 900$ scans, and track-management logic with $N = 3$ and $K = 0$. The following model-based track statistics can be computed: $\tau_{\mu} = 15.0$, $\tau_{\varepsilon} = 3.449$, $Q = 0.8130$, and $\mathrm{F} = 48.78$.

The model-based analysis discussed here assists system analysis. An equally-important task is system design, whereby desired performance statistics are defined and tracker parameters are to be selected. Some discussion of MTT parameter selection may be found in [36].

5.0 CONCLUSIONS

This manuscript provides an accessible introduction to association-based multi-target tracking and, in particular, to the multiple-hypothesis tracking paradigm. We discuss aspects of track and hypothesis management and advanced processing architectures. Next, we present in greater detail a recent extension to multiple-hypothesis tracking that accounts for multiple returns per target per sensor scan. Finally, some compact system-analysis expressions are introduced.

The scope of multi-target tracking includes a wealth of technical challenges and a broad set of applications. It is hoped that this manuscript will stimulate the interest of engineers and scientists, provide some useful pointers to learn more, and encourage contributions to advancing the science of information fusion.

6.0 REFERENCES

[1] R. Mahler, *Statistical Multisource-Multitarget Information Fusion*, Artech House, 2007.

[2] B.-T. Vo, Random Finite Sits in Multi-Object Filtering, *PhD Dissertation*, University of Western Australia, October 2008.

[3] 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.

[4] T. Kurien, “Issues in the Design of Practical Multitarget Tracking Algorithms,” Chapter 3, in *Multitarget-Multisensor Tracking: Advanced Applications*, Y. Bar-Shalom (Ed.), Artech House, 1990.

[5] S. Coraluppi, C. Carthel, M. Luettgen, and S. Lynch, All-Source Track and Identity Fusion, in *Proceedings of the MSS 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] 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, USA, April 2006.

[8] S. Coraluppi and C. Carthel, Modified Scoring in Multiple-Hypothesis Tracking, *ISIF Journal of Advances in Information Fusion*, vol. 7(2), pp. 153-164, December 2012.

[9] S. Coraluppi and C. Carthel, If a Tree Falls in the Woods, it Does Make a Sound: Multiple-Hypothesis Tracking with Undetected Target Births, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50(3), pp. 2379-2388, July 2014.

[10] S. Coraluppi and C. Carthel, Generalizations to the Track-Oriented MHT Recursion, in *Proceedings of the 18th International Conference on Information Fusion*, Washington DC, USA, July 2015.

[11] 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.

[12] S. Coraluppi and C. Carthel, An Equivalence-Class Approach to Multiple-Hypothesis Tracking, in *Proceedings of the IEEE Aerospace Conference*, Big Sky MT, USA, March 2012.

[13] Y. Bar-Shalom, P. Willett, and X. Tian, *Tracking and Data Fusion: A Handbook of Algorithms*, YBS Publishing, 2011.

[14] C.-Y. Chong, G. Castanon, N. Cooprider, S. Mori, R. Ravichandran, and R. Macior, Efficient Multiple Hypothesis Tracking by Track Segment Graph, in *Proceedings of the 12th International Conference on Information Fusion*, Seattle WA, USA, July 2009.

[15] S. Coraluppi and C. Carthel, Multi-Stage Multiple-Hypothesis Tracking, *ISIF Journal of Advances in Information Fusion*, vol. 6(1), pp. 57-67, June 2011.

[16] S. Coraluppi, C. Carthel, W. Kreamer, and A. Willsky, MCMC and MHT Approaches to Multi-INT Surveillance, in *Proceedings of the 18th International Conference on Information Fusion*, Washington DC, USA, July 2015.

[17] R. Streit and T. Luginbuhl, Probabilistic Multiple-Hypothesis Tracking, Technical Report 10428, NUWC-NPT, February 1995.

[18] S. Coraluppi and C. Carthel, A Hierarchical MHT Approach to ESM-Radar Fusion, in *Proceedings of the 15th International Conference on Information Fusion*, Singapore, July 2012.

[19] T. Sathyan, T.-J. Chin, S. Arulampalam, and D. Suter, A Multiple Hypothesis Tracker for Multitarget Tracking with Multiple Simultaneous Measurements, *IEEE Journal of Selected Topics in Signal Processing*, vol. 7(3), June 2013.

[20] D. Clark and R. Mahler, Generalized PHD Filters via a General Chain Rule, in *Proceedings of the 15th International Conference on Information Fusion*, Singapore, July 2012.

[21] C. Degen, F. Govaers, and W. Koch, Emitter Localization under Multipath Propagation Using SMC-Intensity Filters, in *Proceedings of the 16th International Conference on Information Fusion*, Istanbul, Turkey, July 2013.

[22] C. Degen, F. Govaers, and W. Koch, Tracking Targets with Multiple Measurements per Scan, in *Proceedings of the 17th International Conference on Information Fusion*, Salamanca, Spain, July 2014.

[23] Q. Yu and G. Medioni, Multiple-Target Tracking by Spatiotemporal Monte Carlo Markov Chain Data Association, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31(12), December 2009.

[24] T. Kirubarajan, Y. Bar-Shalom, and K. Pattipati, Multiassignment for Tracking a Large Number of Overlapping Objects, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 37(1), Jan. 2001.

[25] B. Habtemariam, R. Tharmarasa, T. Thayaparan, M. Mallick, and T. Kirubarajan, A Multiple-Detection Joint Probabilistic Data Association Filter, *IEEE Journal of Selected Topics in Signal Processing*, vol. 7(3), June 2013.

[26] M. Daun and W. Koch, Multistatic Target Tracking for Non-Cooperative Illumination by DAB/DVB-T, in *Proceeding of the IEEE Radar Conference*, Rome, Italy, May 2008.

[27] 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.

[28] S. Coraluppi and C. Carthel, Counting Targets with MHT and CPHD, in *Proceedings of the IEEE Aerospace Conference*, Big Sky MT, USA, March 2014.

[29] R. Fitzgerald, Track Biases and Coalescence with Probabilistic Data Association, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 21(6), November 1985.

[30] H. Blom and A. Bloem, Probabilistic Data Association Avoiding Track Coalescence, *IEEE Transactions on Automatic Control*, vol. 45(2), February 2000.

[31] 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.

[32] 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.

[33] 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.

[34] L. Svennson, D. Svennson, 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.

[35] A. Reibman and L. Nolte, Optimal Detection and Performance of Distributed Sensor Systems, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 23(1), January 1987.

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