Topics in Multiple-Hypothesis Tracking

Stefano Coraluppi and Craig Carthel
Compunetix Inc.
2420 Mosside Boulevard,
Monroeville PA 15146
USA

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

ABSTRACT

This manuscript discusses some recent advances in multi-target tracking. First, we describe a target kinematic motion model that has a number of appealing characteristics and that, to our knowledge, is not in use within the data fusion community. We describe recent advances in multiple-hypothesis tracking, both in the traditional setting where measurements are informative with respect to target state, and in the limiting case of cardinality-only information. We touch upon recent work in fusion performance modelling and propose a generalization to the recently-introduced MOSPA metric.

1 INTRODUCTION

Effective multi-sensor surveillance requires a high-performance data-fusion capability. Within the class of contact-based, scan-based fusion methodologies (those amenable to network-centric real-time surveillance requirements), track-oriented *multiple-hypothesis tracking* (MHT) is widely acknowledged as the most powerful currently known methodology. It performs better than soft-data association approaches (such as PDA and JPDA), it allows for target maneuvers (unlike conventional ML-based techniques), it has mature track-initiation (unlike most PMHT implementations), and it provides explicit track-label information (unlike SJPDA, PHD, and CPHD). Like all automatic-tracking paradigms, it is suboptimal: in the case of MHT, the sub-optimality stems from the need to manage hypothesis growth and to reduce hypotheses through track pruning or merging techniques. Despite advances in techniques for hypothesis management, complexity remains the key challenge to effective MHT in large-scale scenarios.

Some MHT researchers have proposed advanced hypothesis-management techniques to contend with hypothesis tree explosion. This is generally attempted either locally (based on track-hypothesis scores) or globally (through so-called *k-best global* approaches). These and related techniques seek to maintain hypothesis richness with as small a set of track hypotheses as possible. Unfortunately, these techniques have been found to provide only modest performance improvements. For instance, with *k*-best global techniques, it is found that the top-scoring global hypotheses are quite similar, so that *k* must be quite large to provide sufficient richness. Further, managing global hypotheses is an extremely expensive approach to MHT, indeed it is what motivated development of the track-oriented MHT approach by Kurien and others in the 1980s [1]. Recently, some group tracking ideas have been exploited to address in part the hypothesis explosion in MHT [2].

Single-stage (centralized) MHT is challenged due to its inability to maintain a rich enough set of track hypotheses over long time intervals. Loosely stated, there is an inability "to see the forest behind the hypothesis trees". In our ongoing research, we have discovered that hierarchical fusion provides a powerful methodology for advanced hypothesis management. We have applied multi-stage MHT to address a number of challenges in ground, undersea, and maritime surveillance [3-5], as well as to address fundamental limitations in tracking performance [6].

In this manuscript, we explore some further research directions in multiple-target tracking. Some of the contributions discussed here are described at greater length in a number of recent publications. The reader is referred to these papers for further details on each topic.

This manuscript is organized as follows. In Section 2, we discuss the Mixed Ornstein-Uhlenbeck process [7]. In Section 3, we present a simple tracker performance model that is useful to study the impact of the quality and number of sensors on fusion processing. In Section 4, we discuss a modified MHT approach that is effective at suppressing the damaging impact of spurious tracks on target-originated ones [8]. In Section 5, we discuss the target cardinality problem, a limiting form of the general multi-target tracking problem applicable to large-scale surveillance challenges [9]. Section 6 proposes a generalization to the *Mean Optimal Subpattern Assignment* (MOSPA) metric that was first proposed in [10]. Conclusions are in Section 7.

2 THE MIXED ORNSTEIN-UHLENBECK PROCESS

Kinematic modeling is an important aspect of tracking system design. Significant attention has been given to this topic; the reader is referred to [11-14] and the references therein, as well as the more recent literature, much of which is documented in [15]. While nearly constant velocity (NCV) motion modeling is often adequate, much attention has been devoted in recent years to more complex, multiple-model approaches. Here, we are concerned with the lack of stationary inherent in the NCV-based models, including unbounded position and velocity under long-term prediction.

We introduce the Mixed Ornstein-Uhlenbeck (MOU) process. This model is a generalization of a number of existing models, each with its strengths and limitations:

- The NCV has a well defined velocity, but both velocity and position grow unbounded over time;
- The Ornstein-Uhlenbeck process [14] has bounded position, but no velocity is defined;
- The Integrated Ornstein-Uhlenbeck process [14] has a well defined, bounded velocity, but the position grows unbounded over time.

The MOU process exhibits a well defined velocity and bounded velocity and position over time. As such, it lends itself to steady-state analysis. The initial target state can be defined in a natural way based on the steady-state characteristics of the MOU process, leading to a stationary stochastic process. Interestingly, there are constraints on the allowable steady-state position and velocity; further, there are two choices for the process drift terms leading to identical position and velocity dynamics.

2.1 The basic form

We introduce the following *Mixed Ornstein-Uhlenbeck* (MOU) process that includes drift terms in both position and velocity components. For simplicity, we consider the one-dimensional case. (The extension to multiple dimensions with uncorrelated process noises is immediate.) The process noise is a zero-mean white Gaussian process.

$$\dot{X}(t) = FX(t) + \begin{bmatrix} 0 \\ 1 \end{bmatrix} w(t),\ X(\cdot) = \begin{bmatrix} x(\cdot) \\ \dot{x}(\cdot) \end{bmatrix},\ F = \begin{bmatrix} -\gamma_1 & 1 \\ 0 & -\gamma_2 \end{bmatrix},\ w(t) \sim N(0, q \cdot \delta(0)),\ \gamma_1, \gamma_2 > 0. \tag{2.1}$$

The corresponding discrete-time dynamics are given by (2.2-2.7), where $\Delta t_k = t_{k+1} - t_k$.

$$X_{k+1} = A_k X_k + w_k,\ w_k \sim N(0, Q_k), \tag{2.2}$$

$$A_k = \begin{bmatrix} \exp(-\gamma_1 \Delta t_k) & \dfrac{\exp(-\gamma_2 \Delta t_k) - \exp(-\gamma_1 \Delta t_k)}{\gamma_1 - \gamma_2} \\ 0 & \exp(-\gamma_2 \Delta t_k) \end{bmatrix}, \tag{2.3}$$

$$Q_k = \begin{bmatrix} Q_k^{11} & Q_k^{12} \\ Q_k^{12} & Q_k^{22} \end{bmatrix}, \tag{2.4}$$

$$Q_k^{11} = \frac{q}{(\gamma_1-\gamma_2)^2}\left\{\left(\frac{1-\exp(-2\gamma_1\Delta t_k)}{2\gamma_1}\right)+\left(\frac{1-\exp(-2\gamma_2\Delta t_k)}{2\gamma_2}\right)-2\left(\frac{1-\exp(-(\gamma_1+\gamma_2)\Delta t_k)}{\gamma_1+\gamma_2}\right)\right\}, \tag{2.5}$$

$$Q_k^{12} = \frac{q}{(\gamma_1-\gamma_2)}\left\{\left(\frac{1-\exp(-2\gamma_2\Delta t_k)}{2\gamma_2}\right)-\left(\frac{1-\exp(-(\gamma_1+\gamma_2)\Delta t_k)}{\gamma_1+\gamma_2}\right)\right\}, \tag{2.6}$$

$$Q_k^{22} = q\left(\frac{1-\exp(-2\gamma_2\Delta t_k)}{2\gamma_2}\right). \tag{2.7}$$

It can be shown (via Taylor expansions of the exponential terms) that in the limit $\gamma_1,\gamma_2 \to 0$ we revert to the standard *nearly constant velocity* (NCV) motion model. Similarly, in the limit $\Delta t_k \to 0$ the model converges to the NCV form. The MOU model is characterized by bounded uncertainties in position and velocity:

$$\Delta t_k \to \infty \Rightarrow A_k \to 0,\ Q_k \to \overline{Q} = q\begin{bmatrix} \frac{1}{2\gamma_1\gamma_2(\gamma_1+\gamma_2)} & \frac{1}{2\gamma_2(\gamma_1+\gamma_2)} \\ \frac{1}{2\gamma_2(\gamma_1+\gamma_2)} & \frac{1}{2\gamma_2} \end{bmatrix}, \tag{2.8}$$

$$\Delta t_k \to \infty \Rightarrow X_{k+1} \sim N(0,\overline{Q}). \tag{2.9}$$

2.2 The canonical form

We wish to express the MOU process in canonical form, with the following state variables:

$$Y_k = \begin{bmatrix} x_k \\ \dot{x}_k \end{bmatrix} = \mathrm{T}X_k, \mathrm{T} = \begin{bmatrix} 1 & 0 \\ -\gamma_1 & 1 \end{bmatrix}. \tag{2.10}$$

Accordingly, the state transition matrices and noise covariance matrix are now given by the following:

$$\tilde{A}_k = \mathrm{T}A_kT^{-1} = \begin{bmatrix} \tilde{A}_k^{11} & \tilde{A}_k^{12} \\ \tilde{A}_k^{21} & \tilde{A}_k^{22} \end{bmatrix}, \tag{2.11}$$

$$\tilde{A}_k^{11} = \frac{-\gamma_2\exp(-\gamma_1\Delta t_k)+\gamma_1\exp(-\gamma_2\Delta t_k)}{\gamma_1-\gamma_2}, \tag{2.12}$$

$$\tilde{A}_k^{12} = \frac{-\exp(-\gamma_1\Delta t_k)+\exp(-\gamma_2\Delta t_k)}{\gamma_1-\gamma_2}, \tag{2.13}$$

$$\tilde{A}_k^{21} = \frac{\gamma_1\gamma_2\exp(-\gamma_1\Delta t_k)-\gamma_1\gamma_2\exp(-\gamma_2\Delta t_k)}{\gamma_1-\gamma_2}, \tag{2.14}$$

$$\tilde{A}_k^{22} = \frac{\gamma_1\exp(-\gamma_1\Delta t_k)-\gamma_2\exp(-\gamma_2\Delta t_k)}{\gamma_1-\gamma_2}, \tag{2.15}$$

$$\tilde{Q}_k = TQ_kT' = \begin{bmatrix} \tilde{Q}_k^{11} & \tilde{Q}_k^{12} \\ \tilde{Q}_k^{12} & \tilde{Q}_k^{22} \end{bmatrix}, \tag{2.16}$$

$$\tilde{Q}_k^{11} = \frac{q}{(\gamma_1-\gamma_2)^2}\left\{\left(\frac{1-\exp(-2\gamma_1\Delta t_k)}{2\gamma_1}\right)+\left(\frac{1-\exp(-2\gamma_2\Delta t_k)}{2\gamma_2}\right)-2\left(\frac{1-\exp(-(\gamma_1+\gamma_2)\Delta t_k)}{\gamma_1+\gamma_2}\right)\right\}, \tag{2.17}$$

$$\tilde{Q}_k^{12} = \frac{q}{2(\gamma_1-\gamma_2)^2}\{\exp(-2\gamma_1\Delta t_k)+\exp(-2\gamma_2\Delta t_k)-2\exp(-(\gamma_1+\gamma_2)\Delta t_k)\}, \tag{2.18}$$

$$\tilde{Q}_k^{22} = \frac{q}{(\gamma_1-\gamma_2)^2}\left(\frac{\gamma_1}{2}(1-\exp(-2\gamma_1\Delta t_k)) + \frac{\gamma_2}{2}(1-\exp(-2\gamma_2\Delta t_k)) - \frac{2\gamma_1\gamma_2}{\gamma_1+\gamma_2}(1-\exp(-(\gamma_1+\gamma_2)\Delta t_k))\right). \quad (2.19)$$

The steady-state form is now given by (2.20-2.21).

$$\Delta t_k \to \infty \Rightarrow Y_{k+1} \sim N(0,\overline{Q}), \quad (2.20)$$

$$\overline{Q} = q\begin{bmatrix} \frac{1}{2\gamma_1\gamma_2(\gamma_1+\gamma_2)} & 0 \\ 0 & \frac{1}{2(\gamma_1+\gamma_2)} \end{bmatrix}. \quad (2.21)$$

The continuous-time state transition matrix consistent with this change of variables is the following:

$$\tilde{F}_k = \mathrm{T}FT^{-1} = \begin{bmatrix} 0 & 1 \\ -\gamma_1\gamma_2 & -(\gamma_1+\gamma_2) \end{bmatrix}. \quad (2.22)$$

This leads to the so-called *controller canonical form* [16], where the drift terms are given by:

$$\hat{\gamma}_1 = \gamma_1\gamma_2, \quad (2.23)$$

$$\hat{\gamma}_2 = \gamma_1 + \gamma_2. \quad (2.24)$$

We can express the steady-state covariance matrix in terms of these variables.

$$\overline{Q} = q\begin{bmatrix} \frac{1}{2\hat{\gamma}_1\hat{\gamma}_2} & 0 \\ 0 & \frac{1}{2\hat{\gamma}_2} \end{bmatrix}. \quad (2.25)$$

The solution to (2.23-2.24) is given by the following.

$$\gamma_1,\gamma_2 = \frac{\hat{\gamma}_2 \pm \sqrt{\hat{\gamma}_2^2 - 4\hat{\gamma}_1}}{2}. \quad (2.26)$$

Given requirements for steady-state position and velocity standard deviations σ_p and σ_v, the drift terms are to be fixed as follows.

$$\hat{\gamma}_1 = \frac{\sigma_v^2}{\sigma_p^2}, \quad (2.27)$$

$$\hat{\gamma}_2 = \frac{q}{2\sigma_v^2}. \quad (2.28)$$

From (2.26), the condition for non-oscillatory behavior is given by the following.

$$\frac{q^2}{4\sigma_v^4} - 4\frac{\sigma_v^2}{\sigma_p^2} \ge 0 \Rightarrow \sigma_p^2 \ge \frac{4\sigma_v^3}{q}. \quad (2.28)$$

2.3 Multiple targets

In the multi-target setting, we consider a continuous-time Poisson birth process with mean λ_b and death rate λ_χ. Birth and death processes can be discretized as follows.

$$\lambda_{b,k+1} = \int_{t_k}^{t_{k+1}} \lambda_b \cdot \exp(-\lambda_\chi(t_{k+1}-t))dt = \frac{\lambda_b}{\lambda_\chi}(1-\exp(-\lambda_\chi\Delta t_k)), \quad (2.29)$$

$$p_{\chi,k+1} = \int_{t_k}^{t_{k+1}} \lambda_\chi \exp(-\lambda_\chi t)dt = 1-\exp(-\lambda_\chi\Delta t_k). \quad (2.30)$$

In the limit of large revisit time, the number of births converges to the steady-state distribution on the number of targets:

$$\Delta t_k \to \infty \Rightarrow \lambda_{b,k+1} \to \frac{\lambda_b}{\lambda_\chi} = \bar{\lambda}_b . \qquad (2.31)$$

In the limit of small death rate, the number of births is a direct accumulation of births over the time interval:

$$\lambda_\chi \to 0 \Rightarrow \lambda_{b,k+1} \to \lambda_b \Delta t_k . \qquad (2.32)$$

Multi-target stationarity is insured by setting the birth rate in the first scan according to the Poisson distribution with steady-state parameter $\bar{\lambda}_b$. Additionally, the initial target state for every birth (for all scans) is chosen according to the steady-state distribution $N\left(0,\bar{Q}\right)$.

2.4 Discussion

The MOU process in basic form and in canonical form is illustrated in Figures 2.1-2.2. It is interesting to note from that the steady-state position and velocity are uncorrelated. Also, it is interesting to observe that there is a pair of solutions to (2.26) leading to the same drift terms in controller canonical form, given by (2.23-2.24).

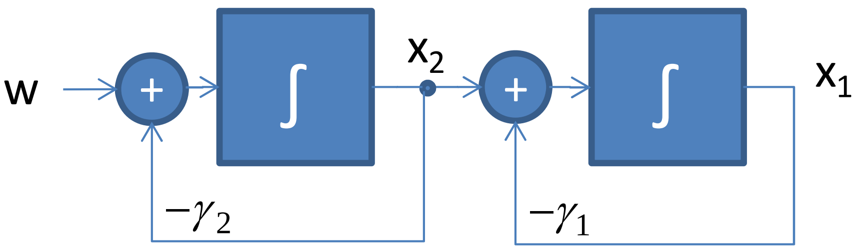

Figure 2.1: The MOU process.

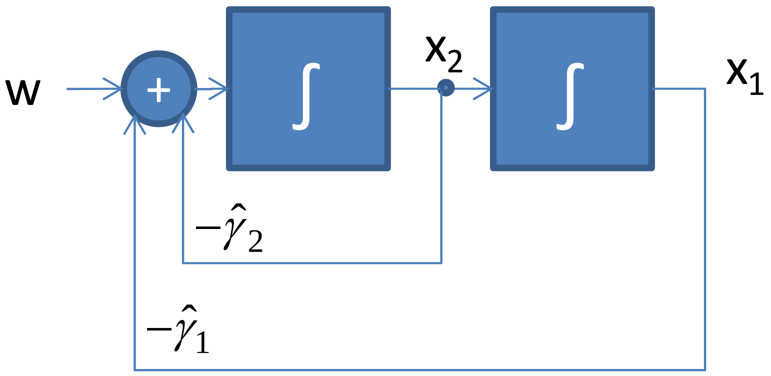

Figure 2.1: The MOU process in controller canonical form.

In figures 2.3-2.4, we illustrate some realizations of the MOU process.

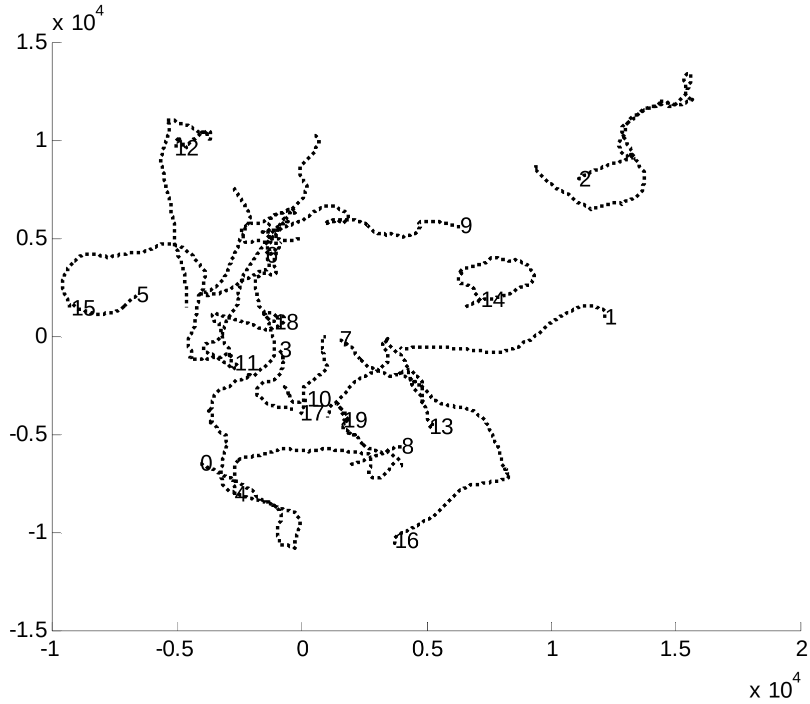

Figure 2.3: A 2D stationary MOU process realization of 1000 scans; only the positional state is plotted. Note the higher target density near the positional origin. For this example, the mean number of targets at any time is 10.

Figure 2.4: Another realization, initialized with one target and birth and death rates set to zero.

In summary, the MOU process exhibits a well defined velocity and bounded velocity and position over time. The model lends itself to steady-state analysis. The initial target state can be defined in a natural way based on the steady-state characteristics of the MOU process, leading to a stationary stochastic process. Interestingly, there are constraints on the allowable steady-state position and velocity; further, there are two choices for the process drift terms leading to identical position and velocity dynamics.
The characteristics of the MOU process make it a good candidate for long-duration target simulations in which the density of targets and their kinematic characteristics remain stationary and well-behaved. The model can be applied directly in recursive filtering as well, by using the state transition matrix and process noise covariance matrix derived here in Kalman filtering or its nonlinear and multiple-model extensions.

Multi-target stationary MOU processes are obtained by setting the initial distribution on number of targets according to steady-state statistics. Similarly, the initial state for all target births is to be set according to the steady-state state covariance.

3 TRACKER MODELING

It is useful to adopt a simple model that captures salient features of fusion processing. Our model relies on the following parameters:

- Target signal-to-noise ratio (s);
- Number of sensors (S);
- Number of sensor resolution cells (C);
- Sensor detection threshold (τ);
- Sensor revisit time [sec] (T_s);
- Track confirmation time [sec] (T_c);
- Track termination time [sec] (T_k);

The sensor detection characteristics are given by (3.1-3.2).

$$p_D = \exp\left(-\frac{\tau}{1+s}\right), \tag{3.1}$$

$$p_F = \exp(-\tau). \tag{3.2}$$

The track detection characteristics are given by (3.3-3.4), where M is the number of contacts required for track confirmation.

$$N = \left\lfloor \frac{ST_c}{T_s} \right\rfloor, \tag{3.3}$$

$$p_{DT} = \sum_{i=M}^{N} \binom{N}{i} p_D^i (1-p_D)^{N-i}, \tag{3.4}$$

$$p_{FT} = \sum_{i=M}^{N} \binom{N}{i} p_F^i (1-p_F)^{N-i}, \tag{3.5}$$

$$\lambda_{FT}(M,N) = \frac{Cp_{FT}(M,N)}{T_c} \text{ [sec}^{-1}\text{]}. \tag{3.6}$$

An illustration of this fusion model is in Figure 3.1, where the track confirmation threshold M is varied, and each curve corresponds to a different surveillance network size.

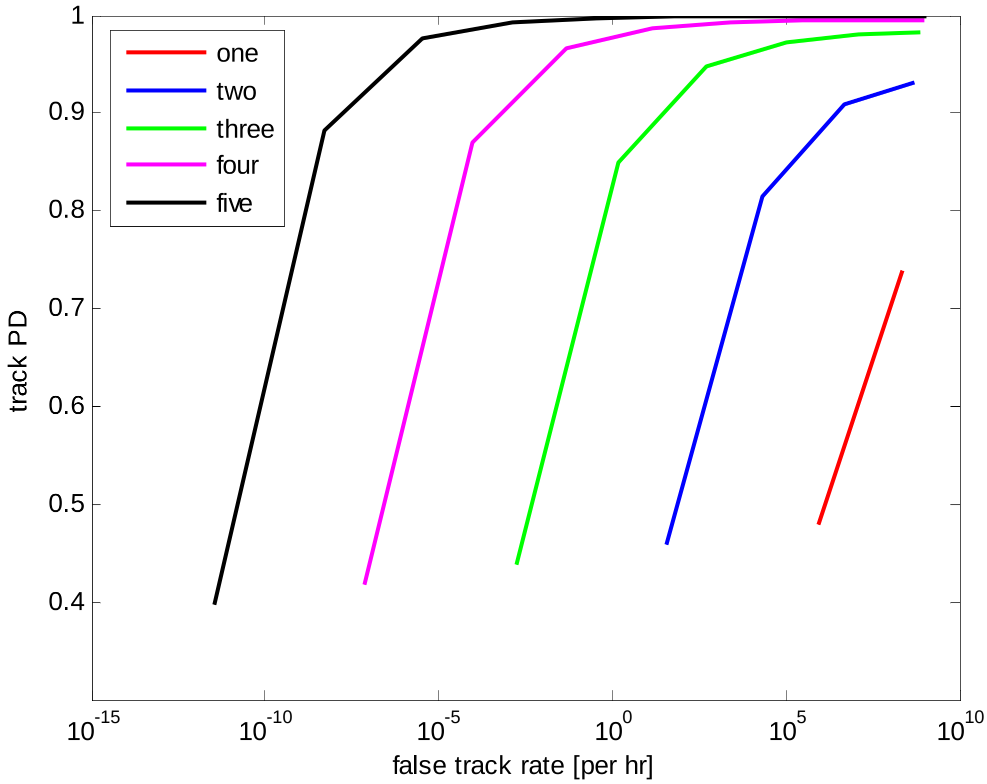

Figure 3.1: Model-based fusion performance.

More discussion of fusion performance modelling is in [12, 17-18] and references therein.

4 MODIFIED MHT FOR IMPROVED PERFORMANCE

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 [1]. 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 $Z^k = (Z_1,...,Z_k)$, 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 $(t_1,...,t_k)$. 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.

$$p(X^k \mid Z^k) = \sum_{q^k} p(X^k, q^k \mid Z^k) = \sum_{q^k} p(X^k \mid Z^k, q^k) p(q^k \mid Z^k). \tag{4.1}$$

A pure MMSE approach would yield the following:

$$\hat{X}_{MMSE}(Z^k) = E[X^k \mid Z^k] = \sum_{q^k} E[X^k \mid Z^k, q^k] p(q^k \mid Z^k). \tag{4.2}$$

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.

$$\hat{X}(Z^k) = \hat{X}_{MMSE}(Z^k, \hat{q}^k), \tag{4.3}$$

$$\hat{q}^k = \hat{q}_{MAP}(Z^k) = \arg\max p(q^k \mid Z^k). \tag{4.4}$$

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 functional optimization without the need for explicit enumeration of global hypotheses. We do so through repeated use of Bayes' rule. Note that, for notational simplicity, we use $p(\cdot)$ for both probability density and probability mass functions. Also, the normalizing constant c_k does not impact MAP estimation.

$$p(q^k \mid Z^k) = \frac{p(Z_k \mid Z^{k-1}, q^k) p(q^k \mid Z^{k-1})}{c_k} = \frac{p(Z_k \mid Z^{k-1}, q^k) p(q_k \mid Z^{k-1}, q^{k-1}) p(q^{k-1} \mid Z^{k-1})}{c_k}, \tag{4.5}$$

$$c_k = p(Z_k \mid Z^{k-1}) = \sum_{q^k} p(Z_k \mid Z^{k-1}, q^k) p(q^k \mid Z^{k-1}). \tag{4.6}$$

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 target are detected with probability p_d. The recursive expression (5) involves two factors that we consider in turn, with the discrete state probability one first. It will be useful to introduce the aggregate variable ψ_k that accounts for the number of detections d for the τ existing tracks, the number of track deaths χ, the number of new tracks b, and the number of false returns $r-d-b$, where r is the number of contacts in the current scan.

$$p(q_k \mid Z^{k-1}, q^{k-1}) = p(\psi_k \mid Z^{k-1}, q^{k-1}) p(q_k \mid Z^{k-1}, q^{k-1}, \psi_k), \tag{4.7}$$

$$p\left(\psi_k \mid Z^{k-1}, q^{k-1}\right) = \left\{\binom{\tau}{\chi} p_\chi^\chi (1-p_\chi)^{\tau-\chi}\right\} \cdot \left\{\binom{\tau-\chi}{d} p_d^d (1-p_d)^{\tau-\chi-d}\right\} \cdot \left\{\frac{\exp(-\lambda_b) p_d^b \lambda_b^b}{b!}\right\} \cdot \left\{\frac{\exp(-\lambda_{fa}) \lambda_{fa}^{r-d-b}}{(r-d-b)!}\right\}, \tag{4.8}$$

$$p\left(q_k \mid Z^{k-1}, q^{k-1}, \psi_k\right) = \frac{1}{\binom{\tau}{\chi}\binom{\tau-\chi}{d}\left(\frac{r!}{(r-d)!}\right)\binom{r-d}{b}}. \tag{4.9}$$

Substituting (4.8-4.9) into (4.7) and simplifying yields the following.

$$p\left(q_k \mid Z^{k-1}, q^{k-1}\right) = \left\{\frac{\exp(-\lambda_b - \lambda_{fa}) \lambda_{fa}^r}{r!}\right\} p_\chi^\chi \left((1-p_\chi)(1-p_d)\right)^{\tau-\chi-d} \left(\frac{(1-p_\chi) p_d}{\lambda_{fa}}\right)^d \left(\frac{p_d \lambda_b}{\lambda_{fa}}\right)^b. \tag{4.10}$$

The first factor in (4.5) is given below, where $Z_k = \{z_j, 1 \le j \le r\}$, $|J_d| + |J_b| + |J_{fa}| = r$, 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). For example, in the linear Gaussian case, $f_d\left(z_j \mid Z^{k-1}, q^k\right)$ 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, $f_b\left(z_j \mid Z^{k-1}, q^k\right)$ is generally the value of the uniform density function over measurement space. Similarly, $f_{fa}\left(z_j \mid Z^{k-1}, q^k\right)$ 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.

$$p\left(Z_k \mid Z^{k-1}, q^k\right) = \prod_{j \in J_d} f_d\left(z_j \mid Z^{k-1}, q^k\right) \cdot \prod_{j \in J_b} f_b\left(z_j \mid Z^{k-1}, q^k\right) \prod_{j \in J_{fa}} f_{fa}\left(z_j \mid Z^{k-1}, q^k\right). \tag{4.11}$$

Substituting (4.10-4.11) into (4.5) and simplifying results in (4.12-4.13). This expression is the key enabler of track-oriented MHT. In particular, it provides a recursive expression for $p\left(q^k \mid Z^k\right)$ that consists of a number of factors that relate to its constituent local track hypotheses.

$$p\left(q^k \mid Z^k\right) = p_\chi^\chi \left((1-p_\chi)(1-p_d)\right)^{\tau-\chi-d} \cdot \prod_{j \in J_d}\left[\frac{(1-p_\chi) p_d f_d\left(z_j \mid Z^{k-1}, q^k\right)}{\lambda_{fa} f_{fa}\left(z_j \mid Z^{k-1}, q^k\right)}\right] \cdot \prod_{j \in J_b}\left[\frac{p_d \lambda_b f_b\left(z_j \mid Z^{k-1}, q^k\right)}{\lambda_{fa} f_{fa}\left(z_j \mid Z^{k-1}, q^k\right)}\right] \frac{p\left(q^{k-1} \mid Z^{k-1}\right)}{\bar{c}_k}, \tag{4.12}$$

$$\bar{c}_k = \frac{c_k}{\left\{\frac{\exp(-\lambda_b - \lambda_{fa})}{r!} \lambda_{fa}^r\right\} \prod_{j \in J_d \cup J_b \cup J_{fa}} f_{fa}\left(z_j \mid Z^{k-1}, q^k\right)}. \tag{4.13}$$

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 $p_d \lambda_b$. 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. 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 detection 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 (4.4), generally by a relaxation approach to an integer programming problem, followed by pruning of all

local hypotheses that differ from $\hat{q}^k$ in the first scan. This 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:

$$\hat{X}_{MAP}(y) = \arg\max f(y \mid X) f(X). \tag{4.14}$$

$$\hat{X}_{ML}(y) = \arg\max f(y \mid X). \tag{4.15}$$

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. 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, batch track extraction would incur significant processing delay.

Given the need for post-association track confirmation and termination logic, a reasonable simplification 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.

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 (4.12) that will prove useful. In particular, a confirmation reward factor $\xi \geq 1$ is applied to track updates for *confirmed track hypotheses* J_c only.

$$\begin{aligned} p(q^k \mid Z^k) = p_\chi^\chi \left((1-p_\chi)(1-p_d)\right)^{\tau-\chi-d} \cdot \prod_{j \in J_c} \left[\frac{(1-p_\chi)\xi p_d f_d(z_j \mid Z^{k-1}, q^k)}{\lambda_{fa} f_{fa}(z_j \mid Z^{k-1}, q^k)} \right] \\ \cdot \prod_{j \in J_d \setminus J_c} \left[\frac{(1-p_\chi) p_d f_d(z_j \mid Z^{k-1}, q^k)}{\lambda_{fa} f_{fa}(z_j \mid Z^{k-1}, q^k)} \right] \cdot \prod_{j \in J_b} \left[\frac{p_d \lambda_b f_b(z_j \mid Z^{k-1}, q^k)}{\lambda_{fa} f_{fa}(z_j \mid Z^{k-1}, q^k)} \right] \frac{p(q^{k-1} \mid Z^{k-1})}{\bar{c}_k}. \end{aligned} \tag{4.16}$$

This mechanism introduces a means to give precedence to confirmed tracks in the data association process. The scheme is shown to provide improved tracking performance results, as described in [8]. Some theoretical justification for its use comes from considering a limiting form of the problem [8].

5 CARDINALITY TRACKING

5.1 The unique-hypothesis approach

In the case of cardinality-only information, the key recursive track-oriented MHT equation is the following:

$$p\left(q^k \mid Z^k\right) = \frac{p_\chi^\chi \left(1-p_\chi\right)^{\tau-\chi} \left(1-p_d\right)^{\tau-\chi-d} p_d^{b+d} \lambda_b^b}{\lambda_{fa}^{d+b}} \frac{p\left(q^{k-1} \mid Z^{k-1}\right)}{\bar{c}_k}, \tag{5.1}$$

where for all scans λ_b is the mean for Poisson-distributed births, λ_{fa} is the mean for Poisson-distributed false returns, p_d is the object detection probability, p_χ is the object termination probability, $q^k = \left(q_1,...,q_k\right)$ is the discrete state that accounts for target births and terminations as well as all data associations, and $\bar{c}_k$ is a factor that is independent of q^k. In the standard track-oriented MHT formulation, we seek the likeliest q^k given measurement data, and then condition on q^k to estimate target trajectories.

$$\hat{q}^k = q_{MAP}\left(Z^k\right) = \arg\max\, p\left(q^k \mid Z^k\right), \tag{5.2}$$

$$\hat{X}\left(Z^k\right) = \hat{X}_{MMSE}\left(Z^k, \hat{q}^k\right). \tag{5.3}$$

In cardinality tracking, each target trajectory is characterized by a track initiation time and a track termination time.

The use of (5.1) for cardinality tracking is problematic. A first difficulty is a computational one: since all measurements gate with all tracks, there is a large number of track hypotheses.

A partial reduction in the complexity of the optimization problem (5.1) can be achieved by recognizing that measurements are interchangeable. Assume we have *N* resolved tracks, n-scan delayed hypothesis resolution, and a sequence of sets of measurements of cardinality $\left\{|Z_1|,...,|Z_{n-scan+1}|\right\}$. We recast the hypothesis generation logic such that a single measurement exists in each scan, but it must be used multiple times. Accordingly, hypothesis branching is more limited than in standard MHT, though no measurement gating is applied. It is easy to see that the number of track hypotheses n_i at level *i* is bounded by the following, with the inequality due to track termination of some branches (inability to reach track confirmation or repeated missed detections). Note that $n_0 = N$.

$$n_i \le 2n_{i-1} + |Z_i|, i = 1,...n - scan + 1. \tag{5.4}$$

The solution to eqn. (5.4) is given by:

$$n_i \le 2^i N + \sum_{j=1}^{i} 2^{i-j} |Z_j|, i = 1,...n - scan + 1. \tag{5.5}$$

Having generated $n_{n-scan+1}$ track hypotheses characterized by a vector of track scores *c*, we have the following optimization problem with $\xi_{n-scan+1}$ =*N*+*n*-*scan*+1 constraints.

$$J = c'x, \tag{5.6}$$

$$Ax = b, \tag{5.7}$$

where the matrix A is a $\xi_{n-scan+1}$-by-$n_{n-scan+1}$ matrix with structure defined by the form of the set of track hypothesis trees, x is a vector of ones and zeros, and the vector b is given by N ones (resolved track constraints) and n-$scan$+1 entries given by $\{|Z_1|,\ldots,|Z_{n-scan+1}|\}$ (measurement constraints). We seek to minimize (5.6) while satisfying (5.7).
The reduced-complexity optimization problem formulation does not address a more serious concern with cardinality tracking: a large number of comparably-scoring global hypotheses in large-scale surveillance applications, leading to the null solution as the optimal choice under (5.2). Indeed, for large surveillance problems, the posterior probability $p(q^k \mid Z^k)$ will be very small for all non-trivial choice of q^k (i.e. for all but the null solution).

Further, unlike conventional tracking, measurements are much less informative and we lack kinematic filter residuals that lead to relatively large hypothesis scores for some association decisions and relatively small scores for others. This leads to a need for a small sensor revisit time, further exacerbating these difficulties.

The difficulties associated with identifying a single global hypothesis point to the fundamental limitation in tracking solutions (like MHT) that do not provide the rich solution that optimal Bayesian tracking would provide: the full posterior probability distribution $p(q^k \mid Z^k)$. Nonetheless, it would be beneficial to indentify a set of global hypotheses that are indistinguishable (due to measurement equivalence) and that provide significant probability mass. This motivates the development in the next section.

5.2 The hypothesis-class approach

We wish to identify a most probably *set* of discrete states $\{q^k\}$, where all member of the set are equivalent under measurement re-ordering. To do so, we must revisit the derivation of track-oriented MHT equations and introduce suitable modifications.

We are interested in a *recursive* and *computationally efficient* expression for $p(\{q^k\} \mid Z^k)$ that lends itself to functional optimization without the need for explicit enumeration of global hypotheses. We do so through repeated use of Bayes' rule. Note that, for notational simplicity, we use $p(\cdot)$ for both probability density and probability mass functions.

$$p(\{q^k\} \mid Z^k) = \frac{p(Z_k \mid Z^{k-1},\{q^k\})\,p(\{q^k\} \mid Z^{k-1})}{c_k} = \frac{p(Z_k \mid Z^{k-1},\{q^k\})\,p(\{q_k\} \mid Z^{k-1},\{q^{k-1}\})\,p(\{q^{k-1}\} \mid Z^{k-1})}{c_k}, \quad (5.8)$$

$$c_k = p(Z_k \mid Z^{k-1}) = \sum_{\{q^k\}} p(Z_k \mid Z^{k-1},\{q^k\})\,p(\{q^k\} \mid Z^{k-1}). \quad (5.9)$$

The recursive expression (5.8) involves two factors that we consider in turn, with the discrete state probability one first. It will be useful to introduce the aggregate variable ψ_k that accounts for the number of detections d for the τ existing tracks, the number of track deaths χ, the number of new tracks b, and the number of false returns $r-d-b$, where r is the number of contacts in the current scan.

$$p(\{q_k\} \mid Z^{k-1},\{q^{k-1}\}) = p(\psi_k \mid Z^{k-1},\{q^{k-1}\})\,p(\{q_k\} \mid Z^{k-1},\{q^{k-1}\},\psi_k), \quad (5.10)$$

$$p(\psi_k \mid Z^{k-1},\{q^{k-1}\}) = \left\{\binom{\tau}{\chi} p_\chi^{\chi}(1-p_\chi)^{\tau-\chi}\right\}\cdot\left\{\binom{\tau-\chi}{d} p_d^{d}(1-p_d)^{\tau-\chi-d}\right\} \cdot\left\{\frac{\exp(-p_d\lambda_b)p_d^b\lambda_b^b}{b!}\right\}\cdot\left\{\frac{\exp(-\lambda_{fa})\lambda_{fa}^{r-d-b}}{(r-d-b)!}\right\}, \quad (5.11)$$

$$p\left(\{q_k\} \mid Z^{k-1}, \{q^{k-1}\}, \psi_k\right) = \frac{b!}{\binom{\tau}{\chi}\binom{\tau-\chi}{d}}. \tag{5.12}$$

The key difference with respect to standard track-oriented MHT is in (5.12), since we must not account for differences in which measurements are taken to be track updates, which are taken to be track births, and how measurements are assigned to tracks. Substituting (5.11-5.12) into (5.10) and simplifying yields the following.

$$\begin{aligned} p\left(\{q_k\} \mid Z^{k-1}, \{q^{k-1}\}\right) = & \left\{\frac{\exp\left(-p_d\lambda_b - \lambda_{fa}\right)\lambda_{fa}^r}{r!}\right\} p_\chi^\chi \left((1-p_\chi)(1-p_d)\right)^{\tau-\chi-d} \\ & \cdot \left(\prod_{i=1}^{d}(r+1-i)\right)\left(\frac{(1-p_\chi)p_d}{\lambda_{fa}}\right)^d \cdot \left(\prod_{i=1}^{b}(r-d+1-i)\right)\left(\frac{p_d\lambda_b}{\lambda_{fa}}\right)^b. \end{aligned} \tag{13}$$

Substituting (5.8) into (5.13) and with further simplification due to the lack of filter residual terms leads to (5.14).

$$\begin{aligned} p\left(\{q^k\} \mid Z^k\right) = & \frac{p_\chi^\chi (1-p_\chi)^{\tau-\chi}(1-p_d)^{\tau-\chi-d}\, p_d^{b+d}\lambda_b^b}{\lambda_{fa}^{d+b}} \frac{p\left(\{q^{k-1}\} \mid Z^{k-1}\right)}{\bar{c}_k} \\ & \cdot \left(\prod_{i=1}^{d}(r+1-i)\right)\left(\prod_{i=1}^{b}(r-d+1-i)\right), \end{aligned} \tag{5.14}$$

$$\bar{c}_k = \frac{c_k}{\left\{\frac{\exp\left(-p_d\lambda_b - \lambda_{fa}\right)}{r!}\lambda_{fa}^r\right\}}. \tag{5.15}$$

Note that the multiplicative weights in (5.14) that are not present in (5.1) obviate in a natural way the need for the modified MHT formalism that addresses the greedy-target problem in the general tracking setting. Cardinality tracking expressed via equivalence classes leads both to computational efficiency (there are much fewer equivalence classes over global hypothesis than there are global hypotheses) and to a well-posed formulation whereby the MAP equivalence class is of interest. Further, the weights in (14) imply *structural results* regarding the form of the optimal equivalence class and optimal track-extraction rules [9]. Accordingly, optimal track extraction is straightforward.

6 GENERALIZED MOSPA

We propose here a generalization to the Mean Optimal Subpattern Assignment (MOSPA) metric that was first proposed in [10]. An existing generalization that partly addresses track continuity is given in [19]. However, that generalization is inadequate as it utilizes a problematic global mapping of tracks to truth. Further, in general one may wish to penalize missed detections and false alarms differently and more heavily than the cutoff distance in the MOSPA metric. Finally, it is useful to allow different norms in defining the distance metric and the corresponding OSPA metric need. The Generalized MOSPA (G-MOSPA) introduced here addresses these concerns. Further, the G-MOSPA metric can be used to define fusion gain in perhaps a more meaningful way than attempted in [20].

6.1 The MOSPA metric

For two sets X and Y with cardinalities $m \le n$ and with Π the set of all possible associations among the sets, the OSPA metric is given by (6.1), where the cutoff distance is given by (6.2) which in turn relies on the Euclidean metric over R^N (6.3):

$$D_{p,c}(X,Y)=\left[\frac{1}{n}\left(\min_{\pi\in\Pi}\sum_{i=1}^{m}d_c\left(x_i,y_{\pi(i)}\right)^p+(n-m)\cdot c^p\right)\right]^{1/p}, \quad (6.1)$$

$$d_c(x,y)=\min\{c,d(x,y)\}, \quad (6.2)$$

$$d(x,y)=\left(\sum_{l=1}^{N}\left|x(l)-y(l)\right|^p\right)^{1/p}. \quad (6.3)$$

The MOSPA metric is the mean OSPA metric over a sequence of pairs of sets.

6.2 The first generalization: multiple norms

Replace (6.1-6.3) by (6.4-6.6). This simply allows greater flexibility in the choice of Euclidean norm. For example, we expect that an operationally-useful choice will be *q*=2 (i.e. the usual Mahalanobis distance) and *p*=1. Indeed for *m*=*n*, this results simply in the mean Mahalanobis distance over all targets.

$$D_{p,c,q}(X,Y)=\left[\frac{1}{n}\left(\min_{\pi\in\Pi}\sum_{i=1}^{m}d_{c,q}\left(x_i,y_{\pi(i)}\right)^p+(n-m)\cdot c^p\right)\right]^{1/p}, \quad (6.4)$$

$$d_{c,q}(x,y)=\min\{c,d_q(x,y)\}, \quad (6.5)$$

$$d_q(x,y)=\left(\sum_{l=1}^{N}\left|x(l)-y(l)\right|^q\right)^{1/q}. \quad (6.6)$$

This still appears to be a metric in the mathematical sense.

6.3 The second generalization: more flexible costs

Replace (6.3-6.4) by (6.7). This considers a vector *c* of thresholds. Accordingly, it decouples the threshold c_1 that specifies how close a track estimate must be to the true target location to be considered a valid detection; the thresholds c_2 that prescribes a cost (in meters) to a missed detection; and the threshold c_3 that prescribes a cost (in meters) to a false alarm. Note that in many applications (think of ballistic missile defense) a missed detection is more costly than a false alarm.

$$D_{p,c,q}(X,Y)=\left[\frac{1}{n}\left(\min_{\pi\in\Pi}\sum_{i=1}^{m_0(c_1)}d_q\left(x_i,y_{\pi(i)}\right)^p+\left(m-m_0(c_1)\right)\cdot c_2^p+\left(n-m_0(c_1)\right)\cdot c_3^p\right)\right]^{1/p}. \quad (6.7)$$

Note that $m_0(c_1)$ is the number of associations that do not exceed the distance threshold c_1; further, for notational simplicity we assume these associations involve the first m_0 true target locations.

Note that the metric given by (6.6-6.7) is not a metric in the mathematical sense (symmetry is violated).

6.4 The third generalization: accounting for track-labelling performance

First, the optimal assignment of tracks to truth $\hat{\pi}(i)$ is performed for each pair of sets of points in the sequence. The optimal assignment is done with respect to (6.7). Subsequently, we assign a (global) label to each track, corresponding to the *mode* association, i.e. the target with which it is associated the most. This allows us to consider the following labeled-set distance, where $\alpha\geq 0$ is the weight given to track-truth mislabeling and $\delta[\cdot,\cdot]$ is the Kronecker delta function.

$$d_l(s,t)=\left(1-\delta[s,t]\right)\alpha\,. \quad (6.8)$$

Accordingly, (6.7) is now replace by (6.9).

$$D_{p,c,q,\alpha}(X,Y)=\left[\frac{1}{n}\left(\sum_{i=1}^{m_0(c_1)}\left(d_q\left(x_i,y_{\hat{\pi}(i)}\right)^p+d_l\left(s_i,t_{\hat{\pi}(i)}\right)^p\right)+\left(m-m_0(c_1)\right)\cdot c_2^p+\left(n-m_0(c_1)\right)\cdot c_3^p\right)\right]^{1/p}. \quad (6.9)$$

While the expression (6.9) accounts for mislabeling costs, it does not account for labeling-diversity costs. That is, for a tracker to perform well with respect to (6.9), there is no incentive not to assign each track estimate a unique label. Accordingly, we must introduce a penalty term for labeling diversity that is analogous to parametric penalization as done in information theoretic settings e.g. the Akaike Information Criterion [21].

The G-MOSPA metric is evaluated over two sequences of sets $X=\{X_{1,},\ldots,X_N\}$ and $Y=\{Y_{1,},\ldots,Y_N\}$ according to (5.3), where L is the number of distinct track labels in Y and where $\beta\geq 0$ is the weight given to label diversity.

$$MOSPA_{d,c,q,\alpha,\beta}(X,Y)=\frac{1}{N}\sum_{i=1}^{N}D_{d,c,q,\alpha}(X_i,Y_i)+\beta\cdot\frac{L}{N}. \quad (6.10)$$

This third generalization of the MOSPA metric shares some similarities with the approach taken in [19]. However, the key difference is the methodology to perform the global mapping to tracks to truth. We believe the approach taken in [19] is problematic; for example, it does not scale to lengthy scenarios where most targets will have multiple tracks associated with them over the course of the scenario.

6.5 Discussion

While G-MOSPA is not a metric in the mathematical sense, it is still useful as a scalar metric that quantifies overall tracking performance. Further, it allows for more flexible and operationally-relevant choices for the Euclidean norm and penalties for missed detections and false alarms, and reflects track-labelling performance without resorting to a problematic global mapping of tracks to truth. The ratio of input and output G-MOSPA can be interpreted as fusion gain.

Use of G-MOSPA provides an interesting unifying perspective on unlabeled and labelled tracking. The former seeks to perform well with respect to the G-MOSPA metric without considering labels i.e. with $\alpha=\beta=0$, while the latter adopts $\alpha,\beta\geq 0$. Naturally, an optimal labelled tracker will not perform as well as an optimal unlabeled one when compared with respect to G-MOSPA with $\alpha=\beta=0$.

It will be of interest to apply G-MOSPA and fusion gain computations to benchmark datasets to examine their properties. In particular, it will be of interest to see whether these quantities have desirable properties: small changes in fusion gain for scenarios of comparable complexity and for a given tracking solution (allowing for sensor management based of detection-level G-MOSPA); increased fusion gain for more sophisticated tracking solutions and a given scenario complexity; lower fusion gain for more complex scenarios and a given tracking solution.

7 CONCLUSIONS

Data fusion is an essential component of surveillance technology for security, defense, and other domains. In this manuscript, we discuss a number of recent research efforts. Our ongoing research is principally focused on extending the state-of-the-art in track-oriented multiple-hypothesis tracking by exploring a hierarchical processing paradigm that allows for efficient analysis of competing hypotheses in large-scale, high-clutter, dense-target surveillance settings.

8 REFERENCES

[1] T. Kurien, Issues in the Design of Practical Multitarget Tracking Algorithms, in *Multitarget-Multisensor Tracking: Advanced Applications*, Artech House, 1990.

[2] C.-Y. Chong et al, Efficient Multiple Hypothesis Tracking by Track Segment Graph, in *Proceedings of the 12th International Conference on Information Fusion*, Seattle WA, USA, July 2009.

[3] S. Coraluppi and C. Carthel, Recursive Track Fusion for Multi-Sensor Surveillance, *Information Fusion*, vol. 5(1), March 2004.

[4] S. Coraluppi and C. Carthel, Distributed Tracking in Multistatic Sonar, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41(3), July 2005.

[5] C. Carthel, S. Coraluppi, K. Bryan, and G. Arcieri, Wide-Area Feature-Aided Tracking with Intermittent Multi-Sensor Data, in *Proceedings of the SPIE Conference on Signal and Data Processing of Small Targets*, Orlando FL, USA, April 2010.

[6] 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.

[7] S. Coraluppi and C. Carthel, The Mixed Ornstein-Uhlenback Process, in *Proceedings of the 14th International Conference on Information Fusion*, Chicago IL, USA, July 2011.

[8] S. Coraluppi and C. Carthel, Addressing the Greedy-Target Problem in Multiple-Hypothesis Tracking, in *Proceedings of the IEEE Aerospace Conference*, Big Sky MT, USA, March 2011.

[9] 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.

[10] 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.

[11] A. Gelb (ed.), Applied Optimal Estimation, The MIT Press, 1974.

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

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

[14] L. Stone, C. Barlow, and T. Corwin, Bayesian Multiple Target Tracking, Artech House, 1999.

[15] *IEEE Transactions on Aerospace and Electronic Systems*.

[16] C.-T. Chen, *Linear System Theory and Design*, Hold, Rinehart, and Winston, 1970.

[17] S. Blackman, *Multi-Target Tracking with Radar Application*, Artech House, 1986.

[18] S. Coraluppi, M. Guerriero, P. Willett, and C. Carthel, Fuse-before-Track in Large Sensor Networks, *ISIF Journal of Advances in Information Fusion*, vol. 5(1), June 2010.

[19] B. Ristic, B.-N. Vo, and D. Clark, Performance evaluation of multi-target tracking using the OSPA metric, in *Proceedings of the 13th International Conference on Information Fusion*, Edinburgh, Scotland, July 2010.

[20] 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.

[21] H. Akaike, A New Look at the Statistical Model Identification, *IEEE Transactions on Automatic Control*, AC-19:716-723, 1974.