Simplified Distributed Tracking

Stefano Coraluppi, Noah Jenkins, and Michael Lexa
STR
600 West Cummings Park
Woburn, MA 01801 USA
stefano.coraluppi@str.us, noah.jenkins@str.us, michael.lexa@str.us

***Abstract*—The purported optimality of single-stage centralized MHT with respect to a MAP data association objective is difficult to achieve in practice. In both single-sensor and multi-sensor domains, distributed MHT offers measurable performance and robustness benefits. Some of these gains are hampered due to complex distributed tracking logic. This paper explores a simplified paradigm for distributed MHT.**

Keywords—multi-target tracking, multiple-hypothesis tracking, distributed tracking

I. INTRODUCTION

Numerous approaches to *multi-target tracking* (MTT) are considered in the literature [1-2]. MTT addresses a complex statistical estimation problem whose key challenge is data association, as the time-varying number of targets and measurement provenance (from targets or due to clutter) is unknown. A leading approach to MTT is *Multiple-Hypothesis Tracking* (MHT) [3]. Advances in MHT over the years have included the track-oriented approach [4], distributed MHT [5-6], and graph-based approximations [7]. A comprehensive overview of MHT history and advances may be found in [8]. A short summary of key aspects of MHT is in Sec. II.

While distributed MHT has provided performance and robustness gains in many operational settings, various issues remain. Sec. III identifies numerous challenges associated with current distributed MHT solutions. Sec. IV introduces a simplified distributed tracking paradigm that holds the potential to mitigate these difficulties.

Our simplified distributed tracking approach does not explicitly reason over upstream data association decisions. This can lead to performance degradation in some settings. Accordingly, in Sec. V we introduce a generalization to simplified distributed tracking that encodes upstream track label as features for downstream exploitation. Calibration is discussed in Sec. VI and further considerations in Sec. VII-VIII.

II. MULTIPLE-HYPOTHESIS TRACKING

MHT solves a *maximum a posteriori* (MAP) estimation problem whereby we seek to identify the best scoring global data association hypothesis. (Strictly speaking, there are many target existence hypotheses that correspond to each data association hypothesis, and one might seek the best existence hypothesis; see [9] for further discussion.) The optimal (global) data association hypothesis prescribes a partitioning of the data, and filtering solutions for each associated sequence of measurements yields the overall tracking solution. Real-time and computational considerations necessitate effective track management logic; see [10] for further details.

In both single-sensor and multi-target settings, optimal MHT is hampered due to complex detection and measurement phenomena including target fading effects, multipath effects, bias errors, limited state observability (notably with passive sensors), etc. Most of the MHT literature (like most of the MTT literature generally) considers a simplifying point-target assumption, whereby a target gives rise to at most one measurement per scan. MHT solution approaches with repeated and merged measurement phenomena are addressed in [11-15].

Despite efforts to extend MHT theory to address various challenges, often the most effective means to contend with complex sensor phenomena, disparate multi-sensor update rates and information content, and challenging multi-target ambiguities (due to closely spaced targets, large sensor measurement errors, Doppler ambiguity, and high sensor clutter rates) is to perform distributed (or multi-stage) MHT. This challenges the intuition afforded by the distributed detection and estimation literature, whereby one often seeks to approach centralized performance via effective distributed processing. In the MTT setting, we seek rather to outperform necessarily suboptimal centralized MHT via an effectively designed distributed MHT solution.

There is amble empirical evidence of the benefits of distributed MHT; see [8] and references therein. As a single-sensor example, a first stage of cautious tracking (low target process noise, low track coast times) followed by a second stage of more aggressive tracking (higher target process noise, larger track coast times) is an effective means to track maneuvering targets in heavy clutter (see Fig. 1). As a (passive) multi-sensor example, it is beneficial to perform single-sensor measurement-space tracking upstream, with downstream multi-sensor (Cartesian) tracking (see Fig. 2). The second stage includes crucial cross-sensor data association decisions to mitigate ghost-track formation. As a forensic example, point-of-origin tracking can be achieved with a first stage of (cautious) forward tracking, a second stage of (cautious) backwards tracking, and a third stage of (more aggressive) track stitching (see Figs. 3-4).

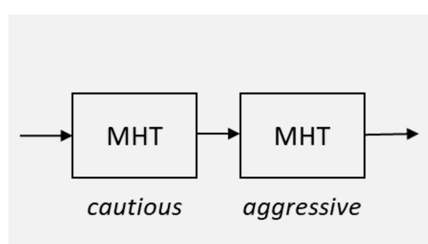

Figure 1. An example of distributed single-sensor tracking.

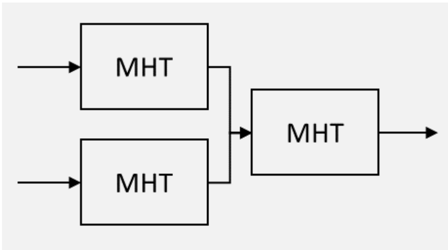

Figure 2. An example of distributed multi-sensor tracking.

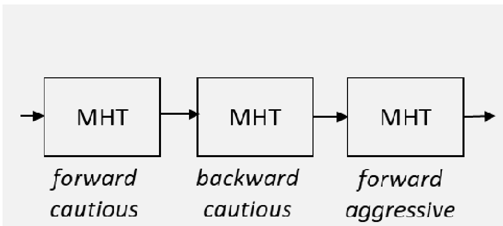

Figure 3. An example of distributed forensic tracking.

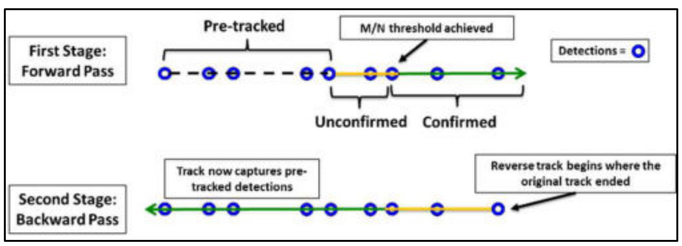

Figure 4. Illustration of first two stages of forensic tracking to include point-of-origin determination.

In most of the distributed MHT architectures that we consider, measurement information is passed to downstream stages of processing. The motivation for this is to avoid the suboptimality associated with suboptimal distributed estimation induced by correlated state estimation errors [1]. In this view, tracks are sequences of associated measurements, as this is more informative than the corresponding sequences of state estimates.

Note that the storage and transmission requirements associated with measurements (and associated measurement error covariances) is comparable to that of state estimates (and associated state covariances). The primary and drastic reduction in data flow is associated with false measurement suppression in first-stage tracking. Hence, it is a good design choice to send track-relevant measurements to downstream processing.

In some instances, an upstream MHT stage may include complex estimation processing that results in equivalent measurements, as with angle-of arrival or *time difference of arrival* (TDOA) sensors [16-18]. In such cases, for downstream processing purposes, tracks may be defined effectively as sequences of equivalent (generally, Cartesian) measurements. This approach generally neglects residual correlations in the sequence of equivalent measurement. In so doing, we recast the problem consistent with the general approach discussed above, whereby measurements rather than state estimates are utilized in subsequent processing stages.

III. Issues in Distributed MHT

Despite the proven effectiveness of distributed MHT in many single-sensor and multi-sensor settings, relying on upstream association decisions in downstream processing can be problematic. Corrective downstream logic to contend with necessarily imperfect upstream processing can be cumbersome, highly complex, and not sufficiently robust. Here, we identify through a series of vignettes several examples of difficulties that may emerge from distributed MHT.

A. Upstream tracking errors

Consider the case where the distributed architecture in Fig. 2 is used, and each of the upstream trackers produces two tracks. In one solution, the tracks IA and IB cross, while in the second solution, the tracks IIA and IIB bounce. This is illustrated in Fig. 5. The downstream MHT may associate IA & IIA as well as IB & IIB, only to be confronted with diverging input tracks with which to update fused tracks.

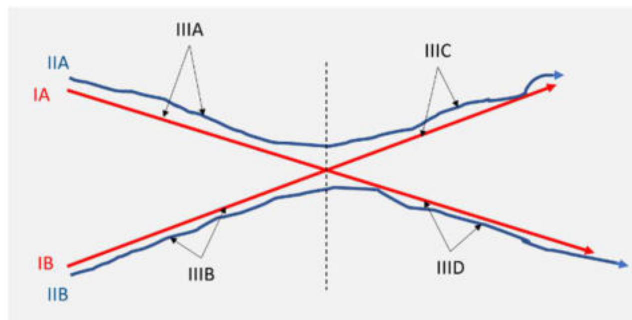

Figure 5. Fragmentation induced by upstream tracking errors.

An example of simple corrective logic is to terminate the fused track IIIA and IIIB, and thereafter to treat the input tracks as new tracks amenable to new fusion decisions. This would induce fused tracks IIIC and IIID. More complex corrective logic to avoid the fragmented fused-tracking solution may require a mix of rollback logic and suitable fragmenting of upstream track feeds.

B. Cautious upstream processing

In some settings, one may need to contend with highly disparate sensors. For instance, we may have high-rate kinematic data from one sensor, and low-rate, highly informative identity measurements from a second sensor. Single-stage MHT while enabling exploitation of identity information is difficult, as the required hypothesis tree depth is prohibitive.

A reasonable processing architecture is shown in Fig. 6, whereby a first stage of cautious processing is employed to form kinematic track fragments that are stitched in a second stage of processing with the aid of the identity measurements [7]. Due to the drastic reduction kinematic track objects after the first stage of processing, the second stage MHT may rely on a large hypothesis depth to enable correct association decisions. A notional scenario illustration is in Fig. 7.

The success of this approach relies on the ability to ensure cautious processing that fragments first-stage tracks when association ambiguity is high. This is extremely difficult to achieve as tracking solutions are based on global hypothesis scores in both hypothesis-oriented and modern track-oriented MHT solutions.

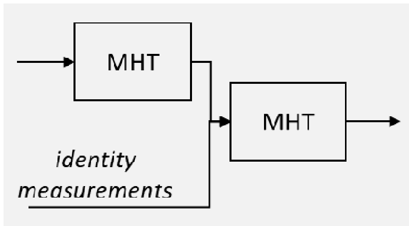

Figure 6. Multi-INT tracking architecture.

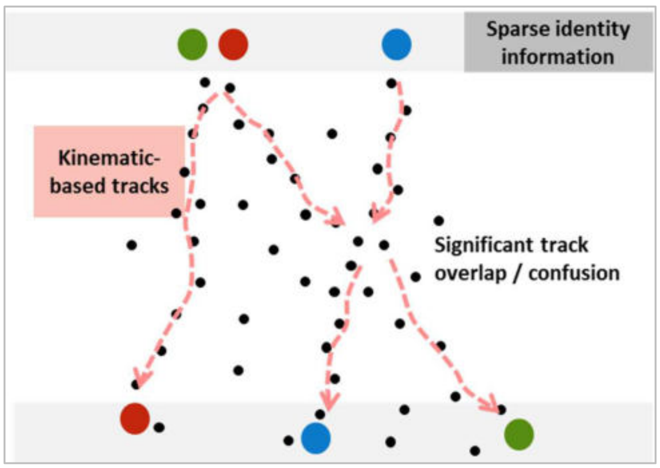

Figure 7. Multi-INT tracking problem illustration.

C. *Variable-quality sensors*

Downstream track-fusion processing generally assumes a negligible false track rate from upstream sensors. As such, downstream processing performs track association decisions, without discarding any unassociated tracks. Relaxing the *no false track* assumption while contending with unequal residual false track rates across sensors is nontrivial. In a sense, we face a generalization to the classic distributed detection problem of designing optimal fusion rules [19]. In distributed tracking, if we are to maintain upstream measurement association constraints, sequences of associated measurements are to be maintained or discarded consistently. The generalization to simple *M-of-N* type rules requires further extension to contend with associated measurement sequences.

As an illustration (Fig. 8), consider the case where system I produces fragmented tracks but with essentially no false tracks, while system II produces higher-continuity tracks but with nontrivial numbers of false tracks as well. Design of systematic and robust track-fusion logic (extensible to the case of *N* upstream systems) is lacking and appears to be highly complex.

D. *Non-fusion decisions*

Consider again the multi-sensor case where track IA and track IIA from distinct upstream trackers are not fused downstream. This decision may be due to initial poor track localization, even though the tracks may subsequently align well. Nonetheless, once it is determined not to fuse these tracks, subsequent track updates do not impact this decision. Modifying real-time track fusion logic to revisit earlier non-fusion decisions significantly complicates reasoning logic, potentially requiring we disregard upstream single-sensor associations.

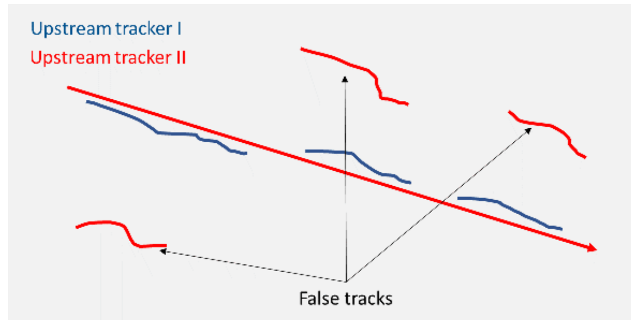

Figure 8. Track fusion with variable-quality inputs.

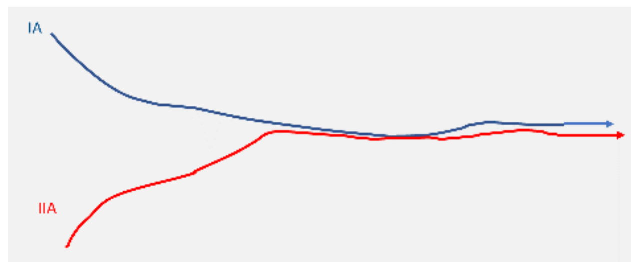

Figure 9. The case for revisiting a non-fusion decision.

E. *Upstream fragmentation*

A related instance of an upstream non-fusion decision is best explained in the single-sensor case, with the architecture in Fig. 1. Refer to Fig. 10. Consider the case where the first stage produces two track fragments. In many cases, the two tracks are temporally non-overlapping and hence amenable to second-stage fusion.

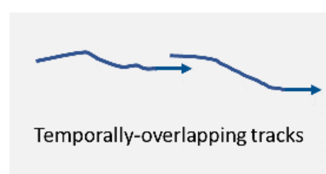

Figure 10. A challenging instance of first-stage track fragmentation.

In some cases, the fragmentation is such that the second track starts before the second is terminated. If we are fortunate, there will not be a sensor frame in which both tracks have a measurement update. If the situation occurs with both track having an update (as in the case where one track correctly makes use of a target-induced measurement while the other makes use of a spurious false alarm), we cannot fuse the tracks in downstream processing without violating the point-target assumption. This makes for a non-robust processing scheme in which the slightest data clash invalidates successful track fusion.

F. *Wide-area tracking*

In wide-area settings, an effective distributed processing solution is to partition the field of regard and to apply a bank of trackers, each devoted to a subregion. Downstream fusion combines single-sensor tracks into an overall wide-area surveillance solution.

Solution robustness for targets that are near the boundary between subregions is achieved by introducing a small overlap between subregions. Refer to Fig. 11. This, however, requires appropriate adjustments in downstream processing to avoid double counting of those measurements that contribute to more than one track.

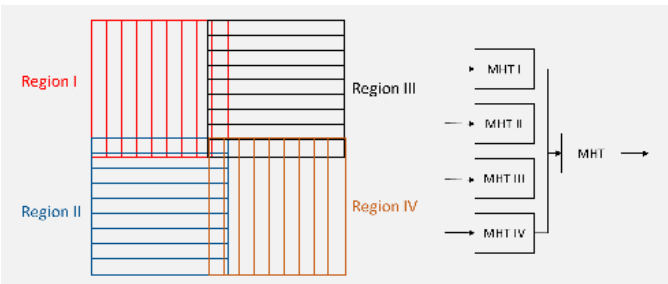

Figure 11. Wide-area distributed tracking.

G. *Distributed sensor networks*

Perhaps the most complex setting illustrating the difficulty of accounting for hard upstream data association decisions in downstream processing is in the context of distributed sensor networks with arbitrary topology. Here, we have sensors with associated processing that will perform data association decisions with local detection data as well as with associated measurement sequences received from neighbors.

In principle, maintaining pedigree (provenance) information for all detections prevents double counting, when the same data arrives to a sensor via multiple channels. On the other hand, how to reconcile conflicting fusion decisions involving the same measurement that arrives via multiple channels?

Referring to Fig. 12, Track IA may be fused with Track IIA and, separately, with Track IIIA, with IIA and IIIA incompatible for fusion. How should System IV contend with Track IA received via two channels? A similar difficult exists when a track produced by System IV flows through the network and returns to System IV, having undergone distinct (and potentially incompatible) associations along the way.

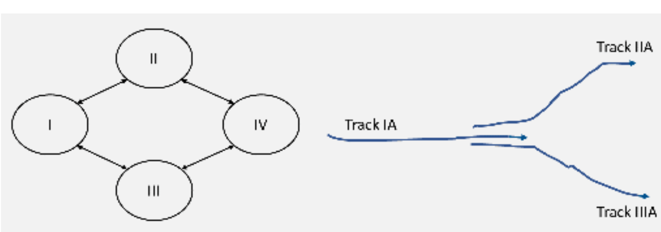

Figure 12. Tracking in a sensor network.

IV. Simplified Distributed MHT

While the distributed tracking paradigm whereby tracks are sequences of associated measurements (rather than sequences of state estimates) is helpful to avoid distributed estimation suboptimality, the examples in Sec. III serve to illustrate some of the difficulties with hard upstream data association decisions. In a sense, the same difficulty exists if we were to require that the tracker not discard measurements deemed to be detections in upstream processing. No false alarms could be discarded by the tracker!

A more helpful paradigm is to view all tracking processes as (generalized) filters that discard spurious measurements and retain those deemed to be target originated. The association decisions established as part of this filtering process is an ancillary tool that need not impact subsequent processing. Stated another way, we may simply ignore all track labels and reason over all available detection data.

A reasonable objection might be offered, in that by ignoring previously established association decisions, inefficiency is introduced, and more complex tracking is needed. Referring for instance to Fig. 6, the downstream fuser will need to consider a larger set of track hypotheses if measurements sequences associated with first-stage kinematic tracks need to be reassessed. On the other hand, most upstream association decisions incur little ambiguity with the confounding clutter having been removed. On balance, the modest increase in downstream complexity is easily offset by no longer having to contend with the many difficulties and instances of solution non-robustness identified above.

Accordingly, all MHT modules in distributed architectures as in Figs. 1-3, 6, 10 etc. may be viewed as *detection-level trackers*. As such, the classical track oriented MHT recursion (as found, e.g., in [13]) is directly applicable, with suitably defined target and sensor statistics. In particular, the sensor detection probability p_d following an initial tracking stage is likely well approximated by the raw sensor detection probability, as relatively few target-originated measurements will typically have been discarded. On the other hand, the false alarm rate Λ (or false alarm density λ) will be much lower, reflecting the decluttering associated with the tracking process.

V. Exploiting Upstream Track Labels

In some applications, the target state may differ substantially among tracking processes in the distributed architecture. This, in turn, may warrant accounting for the track labels established in another tracking process. An example of this would be target emission frequency that might have been exploited (in combination with geometric information) to establish measurement association decisions in one process. If the emission frequency measurement component is not shared with other processes, there is a nontrivial information loss that may hamper association decisions in a downstream tracking process.

Difference in target state assumptions depending on the stage of processing is particularly relevant in large-scale surveillance applications. Like a high-level decisionmaker, the fusion center may need to arrive to an overall surveillance solution, without being hampered with a high-dimensional target state associated with the union of all single-sensor target states.

In such a setting, how to avoid the two extremes of receiving overly prescriptive upstream track labels from an upstream tracking process, and receiving limited-information unlabeled measurements, i.e., those that survive an upstream tracking process? We offer here a potential way forward. Track labels may be cast as a *feature state* that informs downstream processing without imposing problematic hard constraints.

More precisely, we model upstream labels as a fully observed target state that evolves according to a Markovian process as illustrated in Fig. 13.

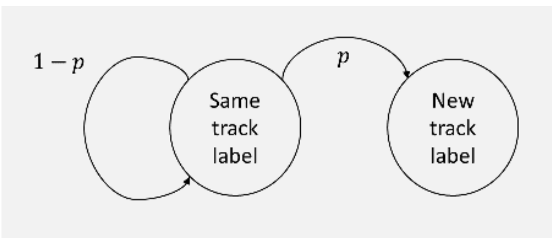

Figure 13. Temporal evolution of track feature.

For each target, one such state exists for each upstream tracker providing measurements. For example, for the distributed MHT architecture in Fig. 2, with two MHTs providing input to the downstream MHT module, the downstream module would have an augmented target state $\tilde{X}_k$ that includes the target state X_k as well as the upstream track label states $L_{I,k}$ and $L_{II,k}$:

$$\tilde{X}_k = \begin{bmatrix} X_k \\ L_{I,k} \\ L_{II,k} \end{bmatrix} \tag{1}$$

We specify below the extensions to classical, track oriented MHT filtering and scoring equations to incorporate such feature states. For simplicity, our description assumes one track feature state as in eqn. (2), but the generalization to multiple feature states is immediate.

$$\tilde{X}_k = \begin{bmatrix} X_k \\ L_k \end{bmatrix} \tag{2}$$

A. *Track filtering*

Let $\{L_k, Z_k, R_k\}$ denote the track label, target measurement, and measurement covariance triple with which downstream track initialization is to be performed.

Assuming Gaussian statistics, track initialization and update result in a sequence of state estimates and associated covariances, denoted by $X(k|k)$ and $P(k|k)$. As we assume perfect label state observation, the label state is given immediately by the current label measurement:

$$L(k|k) = L_k \tag{3}$$

B. *Track scoring*

In track oriented MHT, track scoring relies on the discrete-time Poisson birth rate $\mu_b(\Delta t)$, death probability $p_\chi(\Delta t)$, surveillance volume V with detection probability p_d, and Poisson false alarm rate Λ. Birth and death quantities rely on underlying continuous time birth and death rates λ_b and λ_χ, respectively. False alarms are assumed to be uniformly distributed in measurement space, though this assumption can be relaxed easily. The false alarm rate and false alarm density Λ and λ are related via the surveillance volume V.

$$\mu_b(\Delta t) = \frac{\lambda_b}{\lambda_\chi}\left(1 - e^{-\lambda_\chi \Delta t}\right) \tag{4}$$

$$p_\chi(\Delta t) = 1 - e^{-\lambda_\chi \Delta t} \tag{5}$$

$$\lambda = \frac{\Lambda}{V} \tag{6}$$

Track score initialization is given by eqn. (7); this is the same classical expression as with unlabeled detection data.

$$s_1 = \frac{p_d \mu_b(\Delta t_0)}{\Lambda} \tag{7}$$

In the case of a track update, denoting by C_k the observation matrix, when $L_{k+1} = L_k$, we have the following.

$$s_{k+1} = s_k \frac{\left(1-p_\chi(\Delta t_k)\right) p_d g(Z_{k+1})(1-p)}{\lambda} \tag{8}$$

$$g(Z_{k+1}) = N(Z_{k+1}; CX(k+1|k), \\ C_{k+1}P(k+1|k)C_{k+1}^T + R_{Z_{k+1}}) \tag{9}$$

For $L_{k+1} \neq L_k$, we have:

$$s_{k+1} = s_k \frac{\left(1-p_\chi(\Delta t_k)\right) p_d g(Z_{k+1}) p}{\lambda} \tag{10}$$

For a track coast in or outside the sensor field or regard, we have the following expressions, respectively. These, not surprisingly, are not affected by the track label state.

$$s_{k+1} = s_k \left(1 - p_\chi(\Delta t_k)\right)(1 - p_d) \tag{11}$$

$$s_{k+1} = s_k \left(1 - p_\chi(\Delta t_k)\right) \tag{12}$$

Finally, and as in the unlabeled data setting, we have the following track termination score:

$$s_{k+1} = s_k p_\chi(\Delta t_k) \tag{13}$$

C. *Notes*

With the filtering and scoring extensions noted above, existing hypothesis resolution and track management schemes are directly applicable [10].

It is worth emphasizing that, as specified in eqns. (8-10), the feature enhanced downstream MHT appropriately penalizes both label switches (with factor $1-p$) and non-switches (with factor p). As such, all track-continuation hypotheses are penalized relative to label-free MHT processing. Untracked detection-level data may be viewed as the special case of all single-measurement tracks, with $p = 1$. Track birth hypotheses are not penalized in any manner, regardless of the value of p, as neither label switch nor label consistency are relevant to initiation.

The selection of $p = 0$ in the downstream MHT module is somewhat pathological, though feasible. It corresponds to no label switches being allowed, as a result of which no same-sensor track fusion is performed downstream. (Cross-sensor track associations would still be allowed.)

There is no explicit penalty associated with breaking an upstream track association. That said, the impact of track breaks is reflected in ensuing label switches. As an example, consider Fig. 14 with reference to the architecture in Fig. 2. Assume that MHT I forms tracks IA and IB, while MHT II forms tracks IIA and IIB. IA and IB cross paths, while IIA and IIB do not. The downstream MHT III takes these tracks as input, and we assume that processing results in tracks IIIA (that includes portions of IA, IIA, and IIB) and IIIB (that includes portions of IB, IIA, and IIB). The scores for both IIIA and IIIB will include penalties for the label switch associated with the inclusion of both IIA and IIB. In this sense, the fragmentation of these input tracks is captured indirectly while not being penalized explicitly.

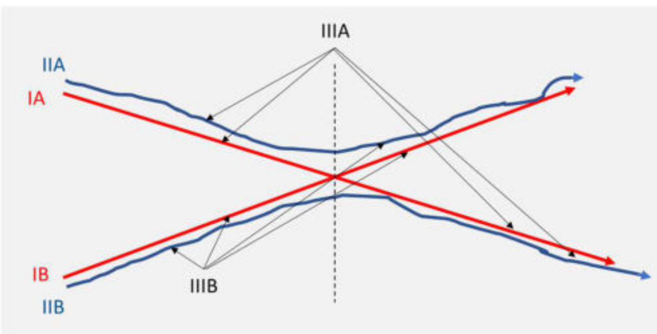

Figure 14. Full track continuity despite conflicting input track data.

As a second example, consider Fig. 15 and the single-sensor distributed architecture in Fig. 1. Here we see that fusion of IA and IB incurs three label penalties. Note that increasingly overlapping and interleaved upstream tracks are more difficult to fuse. This is consistent with relying on upstream labels as governed by the switching parameter p, without enforcing hard constraints.

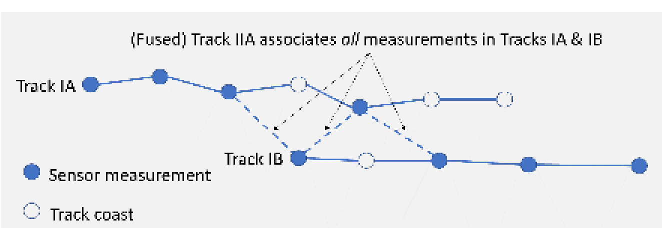

Figure 15. Track fusion score encodes label switches (three in this case).

VI. Calibration

Assuming that calibration for target and sensor statistics has been performed, the MHT scoring equations enable determination of the MAP global data association solution and the corresponding MTT solution. Under our simplified distributed tracking paradigm, a relevant question is how to select MHT parameters for downstream detection-level tracking. We must specify how the *downstream* detection statistics $\tilde{p}_d$ and $\tilde{\Lambda}$ for use in MHT scoring are to be set as a function of *upstream* detection statistics p_d and Λ, assuming upstream track confirmation with an *M-of-N* criterion and track termination with K missed detections. Further, in the case of label-enhanced downstream tracking, a methodology to set the track label switch parameter p is needed. Other parameters relating to target existence and evolution, as well as sensor measurement accuracy, remain unchained across processing stages.

Tracker performance modeling is quite complex, as the impact of incorrect association to false alarms and the presence of closely spaced targets on performance is most easily assessed via empirical investigation. See [1] for a discussion of performance modeling, and [20-21] for *probability of correct association* (PCA) analysis with some extension to multi-scan processing. For our purposes here, we will apply significant simplifying assumptions to derive compact expressions for $\tilde{p}_d$, $\tilde{\Lambda}$, and p.

For simplicity, we assume that the tracker properly disambiguates target-originated measurements from false alarms and does not introduce cross-target association errors. We assume track confirmation and termination based on sequential increments of N and K frames of detection data. Accordingly, probability of confirmation and deletion can be computed as follows [22].

$$p_c = \sum_{k=M}^{N} \binom{N}{k} p_d^k (1-p_d)^{N-k} \tag{14}$$

$$p_t = (1-p_d)^K \tag{15}$$

Further, average time to confirmation and time to termination, respectively denoted by T_c and T_t, can be computed based on geometrically distributed statistics.

$$T_c = \frac{N}{p_c} \tag{16}$$

$$T_t = \frac{K}{p_t} \tag{17}$$

Assuming target existence statistics such that targets are present for a large number of scans relative to confirmation and termination timelines (i.e., small p_χ), we can estimate directly the fraction of time during which an active track exists for any target. This fraction corresponds precisely to the reduction in target detection probability due to a stage of tracking. Specifically, we have:

$$\tilde{p}_d \approx p_d \frac{T_t}{T_c + T_t} \tag{18}$$

For each target, a sequence of tracks will be formed based on the confirmation and termination statistics noted above. Accordingly, we can approximate as follows the probability of label switch for two successive measurements associated with the same track.

$$p \approx \frac{1}{p_d T_t} \tag{19}$$

Finally, we model the probability of false-track update as the probability of clutter in the association gate; the *probability of false update* (p_{fu}) relies on the analysis in [20]. In eqn. (20), γ depends on the association gate size and the measurement-space dimension [23]. For instance, for a 99% association gate, γ can be shown to take values 6.636 ($1D$), 9.21 ($2D$), or 11.34 ($3D$), respectively. See p. 90 of [23] for gate size expressions as a function of measurement-space dimension and gate probability.

For each false alarm, a sequence of false track updates will generate a false track. This *probability of false track* (p_{ft}) may be used directly to model the clutter reduction that results from a stage a tracking, as given by eqn. (22).

$$p_{fu} = 1 - \exp\left(-\Lambda \cdot \gamma^{D/2}\right) \tag{20}$$

$$p_{ft} = \sum_{k=M-1}^{N-1} \binom{N-1}{k} p_{fu}^{k} \left(1 - p_{fu}\right)^{N-1-k} \tag{21}$$

$$\widetilde{\Lambda} \approx p_{ft}\Lambda \tag{22}$$

VII. Why Distributed Tracking?

As we have noted, there is significant evidence in the published literature for the benefit of distributed tracking to achieve performance and robustness. Generally, the focus is on settings with passive or disparate-sensor settings, possibly with further confounding effects (nontrivial bias errors, fading targets, etc.).

The evidence for the benefit of distributed tracking in a single active-sensor setting is much more limited. Here, we address the comparison of one-stage and two-stage tracking performance via model-based analysis. This provides intuition for the benefits that can be expected in multi-stage processing.

For dim targets in clutter, absent any problem structure, it is difficult to motivate why two-stage processing is beneficial. We will impose the reasonable assumption that target motion statistics are non-stationary and exhibit low-maneuver and high-maneuver modes. This structure suggests that performance gains may be achieved in two-stage processing, as illustrated notionally in Fig. 16.

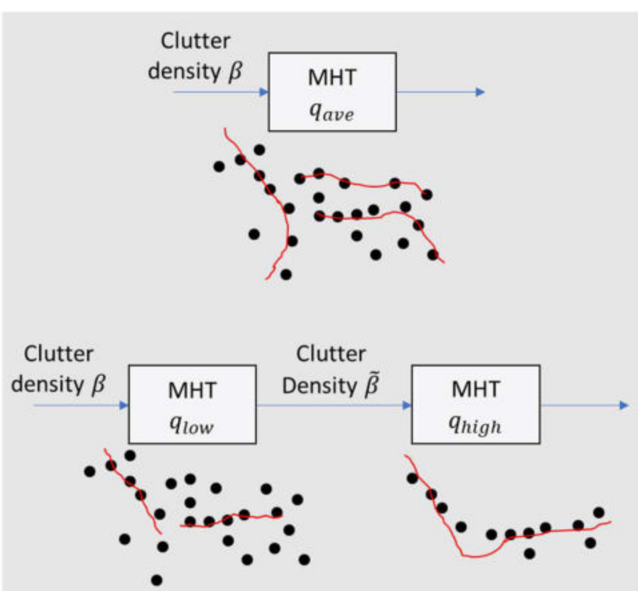

Figure 16. Motivation for two-stage processing for targets in clutter exhibiting alternating low-maneuver and high-maneuver modes.

Let us assume target exhibiting nearly constant velocity (NCV) motion in 2D, with a time varying process noise parameter $q_k \in \{q_{low}, q_{high}\}$, with discrete-time transition probabilities p_{12} (low to high) and p_{21} (high to low), and linear positional measurements. We have $\Delta t_k = t_{k+1} - t_k$.

$$x_{k+1} = A_k x_k + w_k,\ E[w_k w_k^T] = Q_k,\ k = 0, \ldots \tag{23}$$

$$A_k = \begin{bmatrix} 1 & 0 & \Delta t_k & 0 \\ 0 & 1 & 0 & \Delta t_k \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{24}$$

$$Q_k = q_k \begin{bmatrix} \frac{(\Delta t_k)^3}{3} & 0 & \frac{(\Delta t_k)^2}{2} & 0 \\ 0 & \frac{(\Delta t_k)^3}{3} & 0 & \frac{(\Delta t_k)^2}{2} \\ \frac{(\Delta t_k)^2}{2} & 0 & \Delta t_k & 0 \\ 0 & \frac{(\Delta t_k)^2}{2} & 0 & \Delta t_k \end{bmatrix} \tag{25}$$

$$z_k = C x_k + v_k,\ E[v_k v_k^T] = R,\ k = 0, \ldots \tag{26}$$

$$C = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \tag{27}$$

We address the multi-target tracking (MTT) problem via the track-oriented multiple-hypothesis tracking (MHT) solution paradigm. As shown in Fig. 14, in one-stage MHT we use a weighted-average process noise parameter, while in two-stage MHT we use low process noise in the first stage and high process noise in the second stage.

$$q_{ave} = \frac{p_{12} q_{high} + p_{21} q_{low}}{p_{12} + p_{21}} \tag{28}$$

Let us focus on the ability to track successfully a single target in clutter. We rely again on the *Mori Chang Chong* (MCC) exponential model for *probability of correct association* (PCA) [20], which leads to the following simple model for PCA and *track life* (TL), where C_m is a constant that depends on measurement space dimension ($C_2 = \pi$), β is the clutter density, $S(q, \Delta t)$ is the filter residual covariance matrix, and $P(q, \Delta t)$ is the (steady-state) filter prediction covariance that is the solution to the *discrete algebraic Riccati equation* (DARE).

$$PCA(\beta, q, \Delta t) = \exp\left(-C_m \beta |S(q, \Delta t)|^{1/2}\right) \tag{29}$$

$$S(q, \Delta t) = CP(q, \Delta t)C^T + R \tag{30}$$

$$PCA_1 = PCA(\beta, q_{ave}, \Delta t) \tag{31}$$

$$TL_1 = \frac{1}{1 - PCA_1} \tag{32}$$

For two-stage tracking, considering a sequence of N frames of data, we have approximately n_1 frames with the target in a low-maneuver mode, and n_2 frames in the high-maneuver mode. This leads to the following, where $\tilde{\beta}$ is the residual clutter density as input to second stage MHT, P_G is the gate probability as a function of gate size G (eqn. (41) is for the 2D case only), V_G is the corresponding gate volume, p_{fu} is the probability of false update, and p_{ft} is the probability of false track (relying on m consecutive associated measurements). As noted earlier, we typically set $G = 9.21$ which corresponds to $P_G = 0.99$.

$$n_1 = \left\lfloor \frac{p_{21} N}{p_{12} + p_{21}} \right\rfloor \tag{33}$$

$$n_2 = N - n_1 \tag{34}$$

$$PCA_2 = \left(\begin{array}{c} PCA(\beta, q_{low}, \Delta t)^{n_1} \\ \cdot\, PCA(\tilde{\beta}, q_{high}, (n_2 - 1)\Delta t) \end{array} \right)^{1/N} \tag{35}$$

$$TL_2 = \frac{1}{1 - PCA_2} \tag{36}$$

$$\tilde{\beta} = p_{ft}\beta \tag{37}$$

$$p_{ft} = p_{fu}^m \tag{38}$$

$$p_{fu} = 1 - \exp(-\beta V_G) \tag{39}$$

$$V_G = \pi G |S(q_{low}, \Delta t)|^{1/2} \tag{40}$$

$$P_G = 1 - \exp\left(-\frac{G}{2}\right) \tag{41}$$

To summarize, one-stage model-based performance is captured by the pair of measures (PCA_1, TL_1) and two-stage performance by (PCA_2, TL_2). For notional values of target and sensor parameter, we can assess these measures over a range of values for the clutter density β. Performance curves are given in Figs. 17-18.

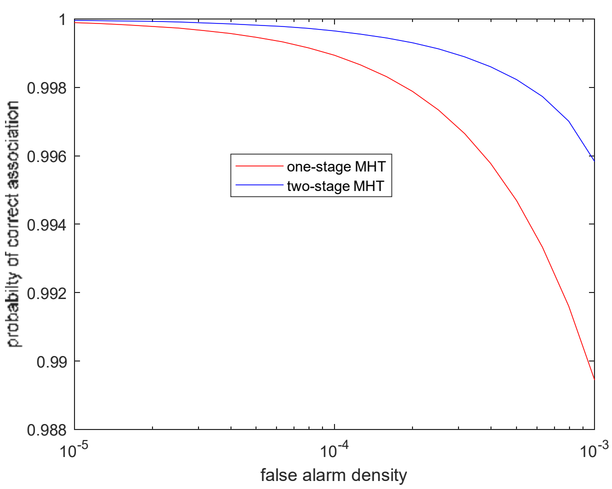

Figure 17. Model-based single-scan PCA as a function of clutter density.

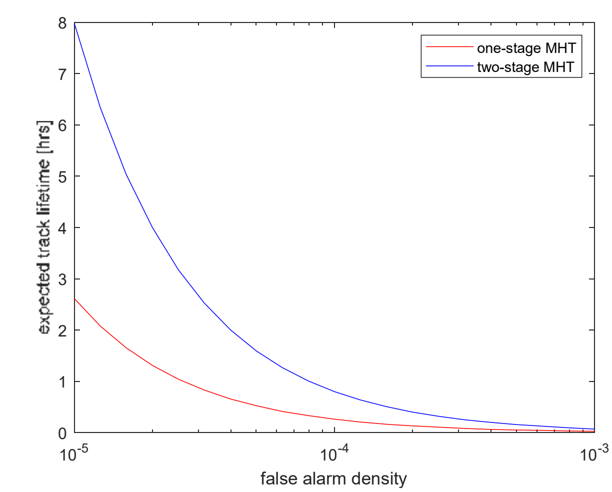

Figure 18. Model-based track life performance as a function of clutter density.

These results suggest that first-stage tracking focused on extracting target-induced measurements during low-maneuver times can be effective, with downstream measurement association with higher maneuvering index in a depleted clutter field of density $\tilde{\beta}$. The model-based assessment matches the intuition and notional illustration of Fig. 16.

VIII. Conclusions

Distributed MHT is a well-established paradigm for advanced surveillance tracking solutions. It avoids the suboptimality of distributed estimation by viewing tracks as sequences of associated measurements. Measurement transmission enables lossless communications for downstream processing and can be performed despite bandwidth restricts, as most measurements do not survive a first stage of tracking. In some settings (including complex nonlinear estimation problems, limited-observability settings, and large-scale surveillance), the use of equivalent measurements is appropriate.

In many cases, distributed MHT outperforms what can be achieved with centralized processing. On the other hand, the robustness properties of distributed MHT may be hampered by overly prescriptive upstream association decisions.

This paper illustrates many settings where difficulties arise, and offers a straightforward, simplified distributed tracking paradigm that overcomes these challenges, albeit with a modest increase in downstream processing as upstream association decisions must be re-established. (The processing increase is limited, as most clutter measurements will have been suppressed upstream.)

With our simplified approach, all MHT modules may be viewed as detection-level trackers with suitably calibrated target and sensor statistics. This provides a drastic reduction in conceptual and implementation complexity that otherwise requires customized logic to recover from a myriad of upstream (or in-stage) association errors.

We extend the simplified distributed MHT paradigm by modeling upstream track labels as feature measurements. This allows for enhanced exploitation of upstream processes without incurring the difficulties of hard exogenous data association decisions. We provide simple expressions for detection statistics following a first stage of tracking. Simulation studies validating the promise of our proposed approach are ongoing and results are forthcoming.

It is of particular interest to validate empirically the model-based analysis of Sec. VII. That is, it is important to determine experimentally whether single-stage active-sensor tracking can be improved upon by judicious multi-stage processing.

We stress again that most evidence of successful distributed tracking in the literature is for multi-sensor, passive sensor, or other complex settings with measurement biases or fading detection statistics. When faced with a simpler single-sensor active-sensor setting (albeit with targets in heavy clutter), can centralized MHT truly be improved upon? Our analysis suggests that this is so, provided that targets exhibit non-stationary dynamical modes (maneuver and non-maneuver) that can be exploited.

References

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

[2] B.-N. Vo, M. Mallick, Y. Bar-Shalom, S. Coraluppi, R. Osborne, R. Mahler, and B.-T. Vo, Multitarget Tracking, in *Wiley Encyclopedia of Electrical and Electronics Engineering,* Wiley-IEEE, 2015.

[3] D. Reid, An Algorithm for Tracking Multiple Targets, *IEEE Transactions on Automatic Control*, vol. 24(6), pp. 843-854, 1979.

[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] 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.

[6] 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.

[7] S. Coraluppi, C. Carthel, and A. Willsky, Multiple-Hypothesis Tracking and Graph-Based Tracking Extensions, *ISIF Journal of Advances in Information Fusion*, vol. 14(2), pp. 152-166, December 2019.

[8] C.-Y. Chong, S. Mori, and D. Reid, Forty Years of Multiple Hypothesis Tracking, *ISIF Journal of Advances in Information Fusion*, vol. 14(2), pp. 131-151, December 2019.

[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, Track Management in Multiple-Hypothesis Tracking, in *Proceedings of the 10th IEEE Sensor Array and Multichannel Signal Processing Workshop*, Sheffield, United Kingdom, July 2018.

[11] S. Mori, K.-C. Chang, and C.-Y. Chong, Tracking Aircraft by Acoustic Sensors - Multiple Hypothesis Approach Applied to Possibly Unresolved Measurements, in *Proceedings of the American Control Conference*, Minneapolis MN, USA, June 1987.

[12] W. Koch and G. van Keuk, Multiple Hypothesis Track Maintenance with Possibly Unresolved Measurements, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 33(3), pp. 883-892, July 1997.

[13] S. Coraluppi and C. Carthel, Multiple-Hypothesis Tracking for Targets Producing Multiple Measurements, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54(3), pp. 1485-1498, June 2018.

[14] S. Coraluppi and C. Carthel, Track-Oriented MHT with Unresolved Measurements, in *Proceedings of the 13th IEEE/ISIF Symposium on Sensor Data Fusion: Trends, Solutions, and Applications*, Bonn, Germany, October 2019.

[15] S. Coraluppi, C. Rago, and C. Carthel, Multiple-Hypothesis Group Tracking, in *Proceedings of the IEEE Radar Conference*, Florence, Italy, September 2020.

[16] R. Osborne and Y. Bar-Shalom, Statistical Efficiency of Composite Position Measurements from Passive Sensors, *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 49(4), pp. 2799-2806, October 2013.

[17] S. Coraluppi, C. Rago, C. Carthel, and B. Bale, Distributed MHT with Passive Sensors, in *Proceedings of the 24th International Conference on Information Fusion*, Sun City, South Africa, November 2021.

[18] G. Mellen, M. Pachter, and J. Raquet, Closed-Form Solution for Determining Emitter Location Using Time Difference of Arrival Measurements, *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 39(3), pp. 1056-1058, July 2003.

[19] P. Varshney, *Distributed Detection and Data Fusion*, Springer, 1997.

[20] S. Mori, K.-C. Chang, and C.-Y. Chong, Performance Analysis of Optimal Data Association with Application to Multiple Target Tracking, in Y. Bar-Shalom (Ed.), *Multitarget-Multisensor Tracking: Applications and Advances Vol. II*, Artech House, 1992.

[21] S. Coraluppi, L. Vertatschitsch, and C. Carthel, Decision Sequencing and Distributed Data Association, in *Proceedings of the 22nd International Conference on Information Fusion*, Ottawa, Canada, July 2019.

[22] A. Reibman and L. Nolte, Optimal Detection and Performance of Distributed Sensor Systems, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 23(1), pp. 24-30, January 1987.

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