18th International Conference on Information Fusion
Washington, DC - July 6-9, 2015

MCMC and MHT Approaches to Multi-INT Surveillance

Stefano Coraluppi, Craig Carthel, William Kreamer, and Alan Willsky
Systems & Technology Research
600 West Cummings Park, Suite 6500, Woburn, MA 01801 USA
Stefano.Coraluppi@STResearch.com, Craig.Carthel@STResearch.com, William.Kreamer@STResearch.com, willsky@mit.edu

***Abstract*—This paper proposes two track fusion methodologies for challenging multi-target tracking (MTT) settings where sensors have highly disparate characteristics and target density is high, leading to many competing tracking solutions. Though distributed multiple hypothesis tracking (MHT) is known to provide a viable solution paradigm, its applicability is limited to medium-size scenarios due to the need for deep hypothesis trees. For large-scale scenarios, a computationally efficient min-cost flow solution paradigm has been proposed that works well for kinematic sensor data, but is not applicable to multi-INT data that includes identity information that does not degrade over time. This paper introduces two approaches to the problem. The first is a natural extension to the MHT paradigm, and seeks to improve performance by considering out-of-sequence processing: the asynchronous MHT (A-MHT). The second adapts a recently proposed Markov Chain Monte Carlo (MCMC) approach to target tracking to multi-INT track fusion: the MCMC Data Fuser (MCMC-DF). A-MHT and MCMC-DF results are promising against an MHT baseline.**

Keywords*—multi-target tracking (MTT); multiple-hypothesis tracking (MHT); track-oriented MHT (TO-MHT); asynchronous MHT (A-MHT); Markov Chain Monte Carlo (MCMC).

I. Introduction

Multi-target tracking (MTT) poses significant technical challenges principally due to the unknown time-varying number of targets as well as to measurement provenance uncertainty, i.e. which measurement originates from which target, and which measurements are false alarms [1]. These challenges are generally not found in classical detection, estimation, and nonlinear filtering problems. Many approaches have emerged over the years; among labelled-tracking approaches, *multiple-hypothesis tracking* (MHT) is generally acknowledged as the most powerful paradigm [2]. First proposed in hypothesis-oriented form, it was later extended to tractable track-oriented form by researchers at ALPHATECH in the 1980s (see [3] and references therein).

Perhaps surprisingly, multi-stage processing has been shown to outperform centralized processing in many settings. The most obvious benefit is in distributed sensor settings, where bandwidth and latency limitations require multi-stage solutions, as well as in settings where single-sensor legacy-system outputs are to be integrated. Further, distributed processing is an effective means to introduce robustness to registration errors and to the target-fading effects that exist in many applications [4-5]. Indeed, detection streaks in space and time are best exploited on an individual sensor basis. Even in the absence of these complications, multi-stage MHT provides an efficient means for hypothesis management that allows significant performance benefits in multi-sensor settings [6-8].

Our specific setting of interest is one for which multi-stage MHT vastly improves upon single-stage (centralized) MHT. We consider a high revisit rate kinematic sensor providing detection-level data with noisy detection and localization statistics. Additionally, we have occasional passive emissions from targets, providing identity information. For simplicity, we assume unambiguous identity information and no spurious (false) emissions. Localization information for the identity sensor tends to be lower than for the kinematic sensor.

First-stage kinematic tracking leads to a fragmented set of high-purity tracks that are reliably associated with the same target; this is achieved by cautious tracking that does not extend track when high ambiguity exists. In second-stage tracking, multiple association hypotheses exist between the identity tracks and kinematic tracks. Resolving these ambiguities via MHT processing requires deep but sparse hypothesis trees. Figure 1 provides a notional illustration of the second-stage track fusion challenge.

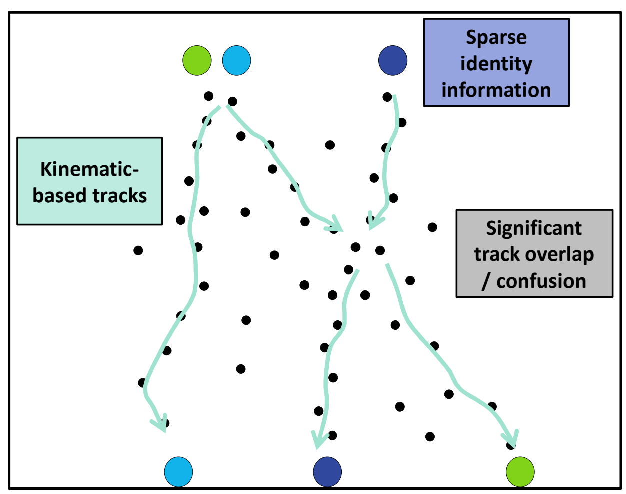

Figure 1. The multi-INT track fusion challenge.

The challenge in multi-INT processing is to consider very deep hypothesis trees for dense target scenarios, for which a large number of track association hypotheses exist. We are interested in application where many nearby targets are present, and when the highly-informative target emissions are potentially infrequent.

An interesting approach to track stitching that avoids hypothesis explosion relies on a Markovian assumption that simplifies likelihood computations as noted in eqn. (1), where y_i represents a track (a sequence of associated measurements) and $y^i = (y_1, \ldots, y_i)$.

$$L(y^n) = L(y_1) \prod_{i=2,\ldots,n} L(y_i | y^{i-1})$$
$$\approx L(y_1) \prod_{i=2,\ldots,n} L(y_i | y_{i-1}). \quad (1)$$

This simplification is generally valid for kinematic data, but is not applicable for identity data as identity information does not degrade over time. Also, the approximation in principle assumes temporally non-overlapping tracks. This approximation has been exploited with promising results for a kinematic sensor large-scale tracking application via a min-cost flow formulation [9]. Unfortunately, the methodology is not applicable to the multi-INT challenge of interest here. While identity information may be used as information overlay to provide additional context as described in [10], the information is not exploited in the data-association solution.

This paper proposes two approaches to multi-INT fusion. Both approaches adopt a distributed tracking paradigm: we consider the output of a first-stage kinematic tracker, and seek an MTT solution by reasoning over this set of tracks as well as identity tracks with infrequent updates. Also, both approaches assume that the problem can be cast as a batch-processing one: we seek a tracking solution for forensic analysis. Tracking with short-delay hypothesis resolution may be performed for computational reasons, but is not required.

Our first approach adopts an MHT paradigm but recasts the problem via an asynchronous data association formulation; this is discussed in Section II. The approach is quite novel, in that out-of-sequence processing is generally addressed when forced upon us due to communication latencies in distributed surveillance settings, and the focus is generally limited to out-of-sequence filtering that does not contend with modifications to association decisions [1]. Asynchronous *global nearest neighbor* (GNN) processing is introduced in [11]. Here, we address the more general problem with target births and deaths, false alarms, missed detections, and the solution is based on a more powerful multiple-hypothesis decision-making framework, leading to the *asynchronous MHT* (A-MHT) solution.

Our second approach leverages the growing body of research on *Markov Chain Monte Carlo* (MCMC) methods and their application to target tracking [12-13]. Our MCMC solution uses both the Metropolis Hastings algorithm and Gibbs sampling, thus allowing for a stochastic optimization approach that operates directly in global hypothesis space. Like the A-MHT, our *MCMC Data Fuser* (MCMC-DF) operates on track-level data.

While the global hypothesis space is necessarily much larger than the track hypothesis space considered in MHT-based approaches, MCMC methods have demonstrated impressive convergence times in many applications. However, it is worth noting that the detection-level MCMC tracking results in [13] are somewhat overstated in that comparison is not made to a track-oriented MHT but rather to a hypothesis-oriented MHT. Here, we compare both the A-MHT and the MCMC-DF algorithms to a baseline track-oriented MHT algorithm applied to the same kinematic and identity track-level data.

Statistical benchmarking of MTT algorithms is a challenge of its own. We adopt a set of performance metrics that rely on an optimal scan-by-scan track-truth assignment, followed by evaluation of completeness and purity metrics for targets and tracks, as well as average localization error. *Track* (or *target*) *quality* measures the fraction of aggregate track (or target) lifetime mapped to a target (or track). *Track* (or *target*) *purity* measures the fraction of track (or target) that is consistent with the mode (i.e. most frequent) assignment. Thus, the quality of the truth trajectories can be thought of as track-level detection probability. The quality of the target tracks is the fraction of track that is target originated. Truth purity is inversely proportional to track fragmentation, and track purity is inversely proportional to track swap occurrences.

Our simulation studies are based on synthetic ground truth with a Poisson birth/death process for target existence, and a 2nd order *Ornstein-Uhlenbeck* (OU) process that governs target evolution. The multi-target OU process is a stable model that admits stationary statistics and long-duration simulations with consistent target densities and motion; a more detailed discussion of the 2nd order OU model and multi-target statistical stationarity may be found in [14].

Figure 2 illustrates the simulation framework for ground truth, data, first-level tracking, and the competing track fusion solutions of interest.

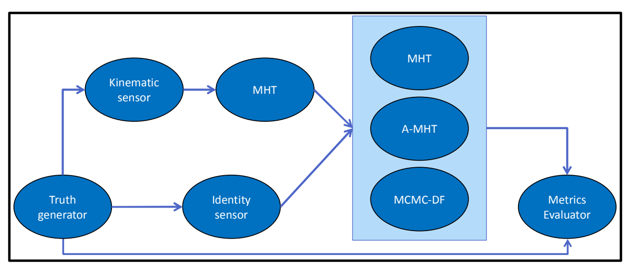

Figure 2. Simluation framework for multi-INT fusion evaluation.

The paper is organized as follows. The multi-target OU process and sensor modeling are described in Section II. The A-MHT and MCMC-DF algorithms are discussed at in Sections III-IV. Section V presents initial simulation results, and conclusions and directions for future work are in Section VI.

II. STATISTICALLY STATIONARY MULTI-TARGET MODEL

We assume a continuous time birth-death process with exponentially distributed target inter-arrival (birth) times with parameter λ_b, and exponentially distributed target lifetime with parameter λ_χ. Discrete-time statistics may be readily obtained, leading to a Poisson distributed number of births with mean $\mu_b(t)$ and death probability $p_\chi(t)$ over an interval of duration t. The discrete-time expressions are given in eqns. (2-3).

$$\mu_b(t) = \frac{\lambda_b}{\lambda_\chi}\left(1 - e^{-\lambda_\chi t}\right), \tag{2}$$

$$p_\chi(t) = 1 - e^{-\lambda_\chi t}. \tag{3}$$

We denote by t_0 the initial time at which no targets are present, and assuming sensors scans at time $(t_i, i \geq 1)$. Statistical stationarity may be achieved by letting $t_0 \to -\infty$, so that the first-interval birth statistics are consistent with the steady-state expected number of targets $\frac{\lambda_b}{\lambda_\chi}$.

We adopt a 2nd order or *modified Ornstein-Uhlenbeck* (MOU) target motion model that generalizes a number of commonly-adopted motion models in the tracking community, each with its strengths and limitations: (1) The *nearly-constant position* (NCP) model is simple, but position grows unbounded over time and velocity is not defined; (2) The *nearly-constant velocity* (NCV) model has a well-defined velocity, but both velocity and position grow unbounded over time; (3) The *Ornstein-Uhlenbeck* (OU) process has bounded position, but no velocity is defined; (4) The *Integrated Ornstein-Uhlenbeck* (IOU) process 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. Further, 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. This, together with the target existence stationarity discussed above, provides a statistically stationary multi-target process. The continuous-time MOU process is illustrated in Figure 3.

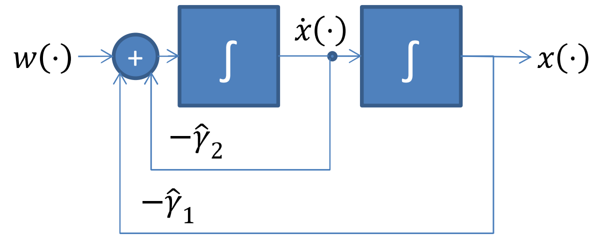

Figure 3. Target motion modeling

MOU dynamics depend on the eigenvalues $-\gamma_1$ and $-\gamma_2$ of the process, which in turn depend on the feedback gains illustrated in Figure 3.

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

The principal case of interest is the real-valued eigenvalue case given by $\sigma_p > \frac{4\sigma_v^3}{q}$. The resulting discrete-time dynamics in each dimension are given by the following, with uncorrelated noises $x_0 \sim N(0, \bar{Q})$, $w_k \sim N(0, Q_k)$, and with $\Delta t_k = t_{k+1} - t_k$.

$$x_{k+1} = A_k x_k + w_k, \tag{5}$$

$$A_k = \frac{1}{\gamma_1 - \gamma_2}\begin{bmatrix} -\gamma_2 \exp(-\gamma_1 \Delta t_k) + \gamma_1 \exp(-\gamma_2 \Delta t_k) & -\exp(-\gamma_1 \Delta t_k) + \exp(-\gamma_2 \Delta t_k) \\ \gamma_1\gamma_2 \exp(-\gamma_1 \Delta t_k) - \gamma_1\gamma_2 \exp(-\gamma_2 \Delta t_k) & \gamma_1 \exp(-\gamma_1 \Delta t_k) - \gamma_2 \exp(-\gamma_2 \Delta t_k) \end{bmatrix}, \tag{6}$$

$$Q_k = \begin{bmatrix} Q_k^{11} & Q_k^{12} \\ Q_k^{12} & Q_k^{22} \end{bmatrix}, \tag{7}$$

$$Q_k^{11} = \frac{q}{(\gamma_1 - \gamma_2)^2}\left(\frac{1 - \exp(-2\gamma_1 \Delta t_k)}{2\gamma_1} + \frac{1 - \exp(-2\gamma_2 \Delta t_k)}{2\gamma_2} - 2\frac{1 - \exp(-(\gamma_1 + \gamma_2)\Delta t_k)}{\gamma_1 + \gamma_2}\right), \tag{8}$$

$$Q_k^{12} = \frac{q}{2(\gamma_1 - \gamma_2)^2}\left(\exp(-2\gamma_1 \Delta t_k) + \exp(-2\gamma_2 \Delta t_k) - 2\exp(-(\gamma_1 + \gamma_2)\Delta t_k)\right), \tag{9}$$

$$Q_k^{22} = \frac{q}{(\gamma_1 - \gamma_2)^2}\left(\frac{\gamma_1}{2}\left(1 - \exp(-2\gamma_1 \Delta t_k)\right) + \frac{\gamma_2}{2}\left(1 - \exp(-2\gamma_2 \Delta t_k)\right) - 2\frac{2\gamma_1\gamma_2}{\gamma_1 + \gamma_2}\left(1 - \exp(-(\gamma_1 + \gamma_2)\Delta t_k)\right)\right), \tag{10}$$

$$\bar{Q} = \begin{bmatrix} \sigma_p^2 & 0 \\ 0 & \sigma_p^2 \end{bmatrix}, \sigma_p = \frac{q}{2\hat{\gamma}_1\hat{\gamma}_2}, \sigma_v = \frac{q}{2\hat{\gamma}_1}. \tag{11}$$

The complex-eigenvalues case under the MOU model and limiting cases of interest that identify the relationship between the MOU, IOU, and NCV models are found in [14].

Our data simulation will assume that kinematic-sensor scans at the sequence of times t^k are characterized by a state-independent detection probability p_d, linear measurements with additive Gaussian noise $v_k \sim N(0, R_k)$, Poisson distributed false alarms with parameter Λ, and false alarms distributed in measurement space according to $N(0, \bar{Q})$. Identity-sensor detections are generated in a similar manner, but with very small p_d and no false alarms. As discussed in Section I, we assume no measurement provenance uncertainty from identity-sensor detections.

Our approach to identity-sensor modeling is essentially to model a passive sensor as an active one. This amounts to the assumption that all targets are equipped with an active emitter and that target inter-emission times are exponentially distributed.

III. ASYNCHRONOUS MHT

The fundamental MHT equation that allows for recursion evaluation of global hypotheses based on local hypothesis computations is given below, where q^k denotes a global hypothesis at time t_k, q_k is an incremental global hypothesis at time t_k, τ is the number of targets at time t_{k-1}, and r, d, b, and χ are the number of returns, number of detections of existing targets, number of new targets, and number of target deaths, respectively, at time t_k. It is worth emphasizing that the factorization of local hypothesis computations relies crucially on the assumption of Poisson-distributed target births. The recent extension to the MHT recursion that relaxes the usual MHT assumption of target detection at birth without increasing algorithmic complexity is in [15]. When target statistics are stationary, the extension leads to an adjustment to the birth statistics to account for previously-undetected births [16].

$$p(q_k|Z^{k-1}, q^{k-1}) = \left\{\frac{\Lambda^r e^{-\Lambda} e^{-\mu_b}}{r!}\right\} p_\chi^\chi \left((1-p_d)(1-p_\chi)\right)^{\tau-d} \left(\frac{p_d}{\Lambda}(1-p_\chi)\right)^d \left(\frac{p_d}{\Lambda^i}\mu_b\right)^b. \quad (12)$$

While MHT processing is effective for kinematic tracking, its application for our multi-INT processing is extremely challenging due to the need for deep hypothesis trees to benefit from highly-informative target emissions. Consider the following illustrative example.

Assume there are N sensor scans, where the first and last scans are due to the low-rate sensor and intervening scans are due to the high-rate sensor. We assume one-dimensional MOU target motion and positional sensor measurements. We consider a number of solution schemes. The first is the *clairvoyant* solution, where measurement provenance is assumed to be known for the high-rate sensor as well. This reduces to a set of linear filtering problems for which the KF provides an optimal solution.

The second solution is to use the *global nearest neighbor* (GNN) assignment with sequential processing of all sensor scans. Note that, in general, data association errors do occur. We recover from such errors at the last scan (from the low-rate sensor), when measurement provenance is known. Naturally, for a fixed number of targets and target density, as the number of scans of data increases, the problem becomes more difficult in the sense that data association errors will accrue prior to the last scan of data.

Can we do better than the GNN solution if we constrain ourselves to maintaining a single global hypothesis? It turns out that improved performance is possible. This is achieved by performing Kalman smoothing based on the current state estimates at time t_k and the final scan of measurements at time t_N, to estimate target positions at time t_k. These estimated positions can be used in defining the GNN assignment matrix, resulting in a more reliable solution than is possible with sequential processing. We call this approach *asynchronous GNN.*

Figure 4 illustrates one realization of target trajectories, along with four candidate solutions. These are the clairvoyant solution, the sequential GNN solution, and two variations on the asynchronous GNN solution – one with scoring based on approximate Kalman smoothing, and one exact scoring based on Kalman filtering. Note the "recovery" at the last scan exhibited by the sequential GNN solution.

Figure 5 illustrates Monte Carlo performance results as a function of the number of sensor scans. When there are only two scans of data, both from the low-rate sensor, all four solutions coincide. The clairvoyant solution improves slightly with an increasing number of scans, due to filter convergence. The three solutions for which measurement provenance on high-rate sensor returns is unavailable all degrade with increasing number of scans, measured in terms of average position estimation error.

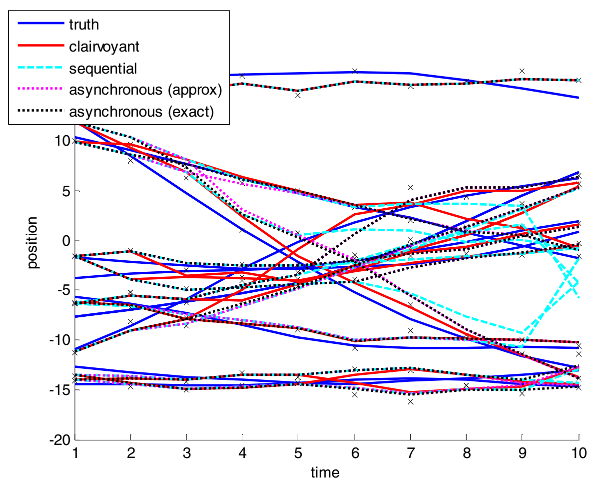

Figure 4. Realization of competing solutions for multi-target filtering.

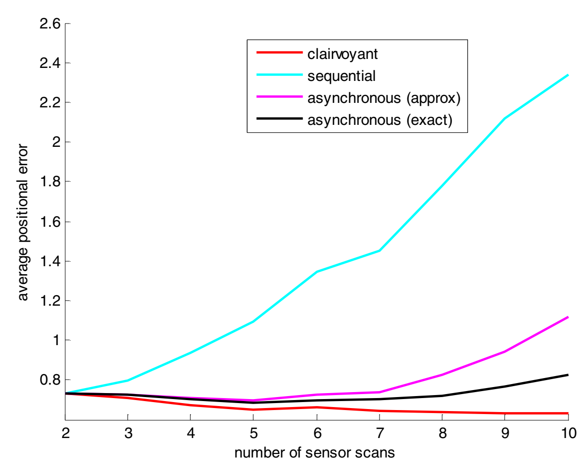

Figure 5. Performance as a function of scenario duration.

We see that the asynchronous GNN provides a dramatic multi-target filtering improvement over sequential GNN, while maintaining the same processing complexity, albeit with the need for Kalman smoothing or an additional Kalman filtering update in defining the GNN assignment matrices. Note that in the asynchronous GNN solutions the information in the final scan is used solely to improve association decisions, and does not impact filter updates. There is no issue of repeated use of final-scan information.

The above result may be applied to the general MTT problem, for which the number of targets is unknown, as well as in MHT processing, where data association decisions are based on sliding window of scans and multiple association hypotheses are maintained. This is best described via a notional example.

Figure 6 illustrates A-MHT processing of track level identity data S1 and kinematic data W1 and W2. The processing proceeds in batch or forensic mode. We initialize the set of track hypothesis trees with all unassociated identity tracks. Next, we proceed to process kinematic tracks sequentially, whereby the entire track is processed in forming and scoring track hypothesis trees; this in analogous to the preceding asynchronous GNN discussion. However, here we consider as well new target hypotheses. Further, we consider several processing steps before pruning the set of track hypothesis trees, as prescribed under the n-scan pruning logic that we adopt [3].

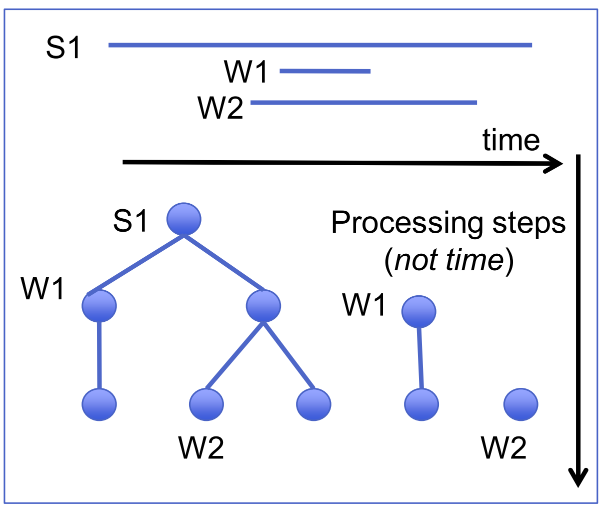

Figure 6. A-MHT hypothesis formation via batch processing of tracks.

The track order for processing is somewhat arbitrary, but for convenience we order kinematic tracks by time, starting with the first track to terminate. This explains why, in the example, we process kinematic track W1 first. We do not consider the hypothesis that both W1 and W2 are due to the same target as they overlap in time. This would require a redundant-measurement sensor model, for which recent developments in MHT are discussed in [17].

A-MHT hypothesis generation logic is different from classical MHT in certain details as well. As an example, there is no need to consider limited track coasting prior to a track termination hypothesis. The A-MHT will allow for arbitrarily long track coasts, since the single-sensor tracks in any fused track hypothesis may exhibit significant temporal separation.

IV. MCMC Data Fuser

MHT-based approaches to multi-INT surveillance – even the A-MHT – will ultimately fail in sufficiently large-scale, high-density, and prolonged inter-emission-time target scenarios. Hence, we consider a stochastic sampling approach that is motivated by recent research on MCMC approaches to multi-target tracking [12-13].

Our approach operates on track-level data. As such, all global hypotheses will necessarily account for all single-sensor tracks. Our initial condition is to hypothesize that all single-sensor tracks originate from a distinct target. Next, we proceed to perturb the current global hypothesis, evaluating new global hypothesis with a batch form of the global hypothesis score that is consistent with MHT scoring identified in eqn. (12). The *MCMC Data Fuser* (MCMC-DF) keeps track of the best global hypothesis encountered in past iterations; this allows the sampling scheme to explore global hypothesis space by occasionally migrating to a worse-performing global-hypothesis state without degrading the current solution. The MCMC-DF proposal density used in sampling must satisfy certain technical ergodicity, reversibility, and detailed balance assumptions to guarantee convergence to the optimal global hypothesis.

The MCMC-DF includes Metropolis-Hastings merge and split moves, as well as Gibbs sampling that cycles through all single-sensor tracks and potentially changes the fused track to which each is associated. Figure 7 illustrates a Metropolis-Hastings split move that extracts a single-sensor track from a fused track, generating a new (singleton) fused track.

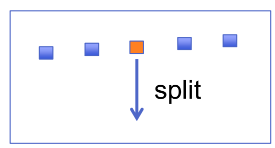

Figure 7. A Metropolis-Hastings split move in MCMC-DF processing.

The potential disadvantage of the MCMC-DF over the A-MHT is that it operates in an enormous global hypothesis space, whereas track-oriented MHT approaches operate in local hypothesis space and rely on efficient relaxation-based optimization approaches to identify a global hypothesis solution. The potential advantage of the MCMC-DF over the A-MHT is that it need not characterize all of hypothesis space and relies on impressive convergence properties that MCMC technology has demonstrated in many application domains.

V. Simulations

We consider a 2D scenario with the following statistical characteristics.

- Target existence: $\lambda_b = 0.01\text{sec}^{-1}$ (birth rate), $\lambda_\chi = 0.001\text{sec}^{-1}$ (death rate); $T = 10\text{min}$ (scenario length).
- Target dynamics (in each dimension): $\sigma_p = 250\text{m}$ (positional std. dev.), $\sigma_v = 6.5\text{m} \cdot \text{sec}^{-1}$ (velocity std. dev.), $q = 5m^2 \cdot \text{sec}^{-3}$ (process noise);

- Sensor characteristics: $\Delta t^W = 2\text{sec}$ (kinematic sensor revisit time), $p_d^W = 0.8$ (kinematic sensor detection probability), $\Lambda = 0$ (kinematic sensor false alarm rate per scan), $\Delta t^S = 0.5\text{sec}$ (identity sensor revisit time), $p_d^S = 0.05$ (identity sensor detection probability); $R = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \text{m}^2$ (measurement error std. dev. for both kinematic and identity sensors).

An illustration of end-to-end processing is in Figure 8. The left pane shows ground truth trajectories (black) and kinematic and identity detections (red & green, respectively). Single-sensor kinematic tracks based on first-stage MHT processing as well as fused multi-sensor tracks are shown in the right pan (red & blue, respectively). These are nearly indistinguishable as kinematic localization of fused tracks is similar to the kinematic tracks, though fragmentation is greatly reduced and associations are aided by identity-sensor information. As discussed in Section I and as illustrated in the bottom pane, our processing architecture includes performance evaluation based on a truth-tracks comparison.

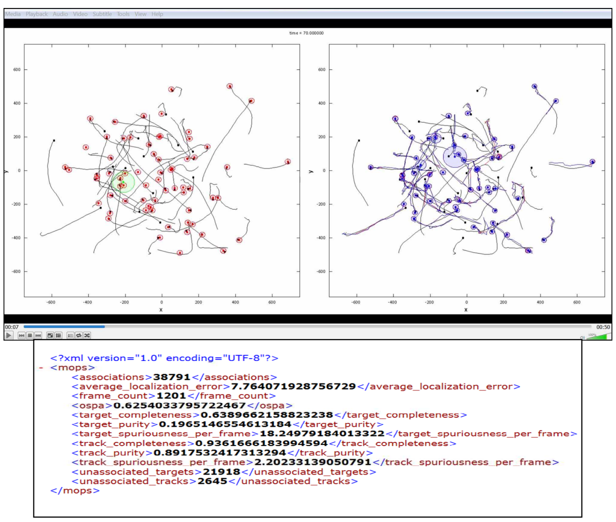

Figure 8. One realization of the benchmark scenario.

In some settings, target and track purity are second order metrics while target and track completeness are of primary importance. On the other hand, for our application, kinematic tracks are well localized and do not incur track swaps, so that track completeness is high. Likewise, target completeness is high due to the high sensor detection probability. Accordingly, our principal focus is on the performance of the MCMC-DF and A-MHT as assessed by target and track purity, relative to an MHT track-fusion baseline.

We first examine MCMC-DF performance in Figure 9, using the Metropolis-Hastings merge move and Gibbs sampling. Recall that the initial condition is for all single-sensor tracks to be unassociated. Hence, track purity is one, and remains nearly one with increasing iteration number. This implies that track associations are performed correctly. Target purity starts low and gradually increases to over 90%, indicating that many track associations are performed and that track fragmentation is greatly reduced.

It is of interest to measure the significance of identity-sensor data on overall performance. We do this by applying the MCMC-DF only to kinematic track-level data. The results are in Figure 10, where we see much lower target purity. Thus, we see that identity information enables much more data fusion, as higher target purity implies lower track fragmentation, while still performing correct association decisions as indicated by high track purity.

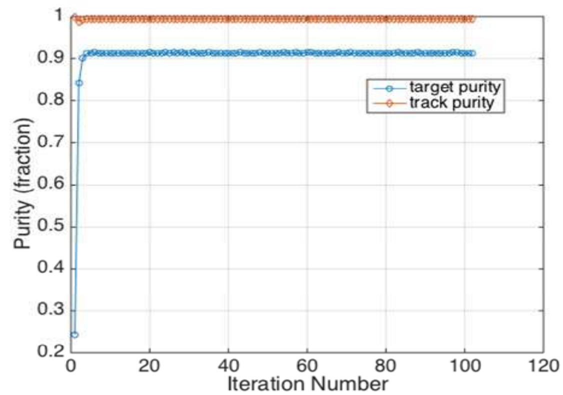

Figure 9. MCMC-DF performance with both Metropolis-Hastings and Gibbs sampling.

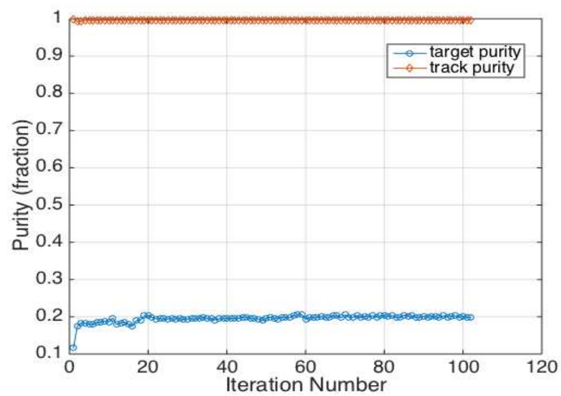

Figure 10. MCMC-DF performance with kinematic data only.

It is surprising that identity information has a significant impact on fragmentation reduction; one might expect that the primary benefit would be lower-error association decision-making. Identity information often establishes that a long-duration target is present: a notional illustration is below.

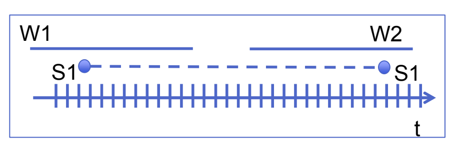

Figure 11. Identity information induces association of kinematic tracks that would not otherwise fuse. The two-target kinematic-only optimal hypothesis is replaced by a one-target multi-sensor optimal hypothesis.

Target visibility under high detection-probability kinematic sensing induces the association of the identity track with multiple kinematic tracks that would not be associated otherwise. Indeed, were the identity track not associated with the kinematic tracks, we would incur a penalty in global hypotheses score with a target that persists for a long time and is not observed by the kinematic sensor.

The MCMC-DF iterations for the results given above are somewhat slow, on the order of 20sec per iteration. This is due to the rather slow Gibbs sampling move that sequentially considers resampling all single-sensor tracks. We have improved both performance and complexity by considering instead Metropolis-Hastings merge and split moves, with no Gibbs sampling. Execution time is reduced to 2-3sec per iteration, and an improved target purity of about 94% is achieved, as illustrated in Figure 12.

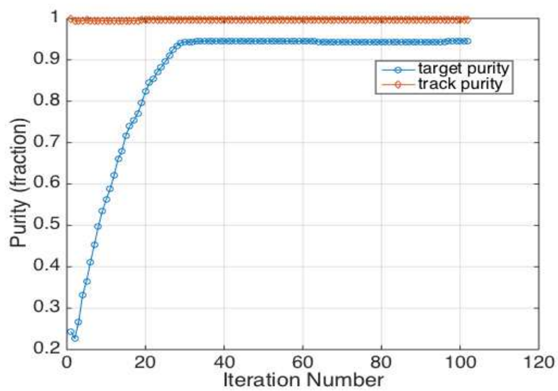

Figure 12. MCMC-DF performance with Metropolis-Hastings moves.

Unfortunately, the MCMC-DF is not able to reach target purity of one, even after prolonged processing – see Figure 13. The apparent lack of convergence to the optimal global hypothesis is a matter of current investigation. Nonetheless, the MCMC-DF achieves comparable performance to baseline MHT processing with a hypothesis tree depth (n-scan) of 10; this is significant and promising performance result for the MCMC-DF approach.

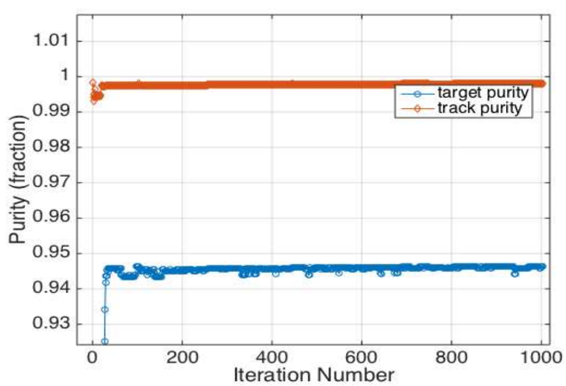

Figure 13. No further improvement with continued MCMC-DF iterations.

We consider MHT and A-MHT on the same scenario as above, though we consider as well a second, more challenging setting, with modifications as noted below.

- Modified settings: $\Lambda = 10$ (kinematic sensor false alarm rate per scan), $R = \begin{bmatrix} 5 & 0 \\ 0 & 5 \end{bmatrix} \mathrm{m}^2$ (measurement error std. dev. for both kinematic and identity sensors).

Target and track purity results are illustrated in Figures 14-17, for a range of n-scan values. Note that the meaning of n-scan is the usual one for MHT processing, i.e. the number of additional temporal scans before global hypothesis resolution. An n-scan of 0 corresponds to GNN processing. On the other hand, n-scan processing in the A-MHT indicates the number of single-sensor tracks that are ingested (in batch form) prior to global hypothesis resolution. Thus, though both MHT and A-MHT processing complexity increase as a function of n-scan, the processing time of the algorithms is not the same for a given n-scan. Nonetheless, examining performance as a function of n-scan is instructive.

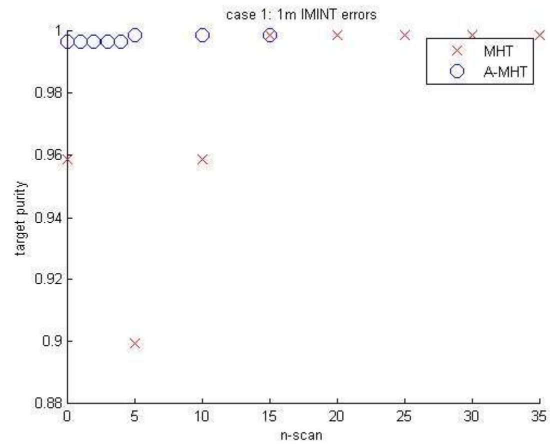

Figure 14. MHT and A-MHT target purity (easier scenario).

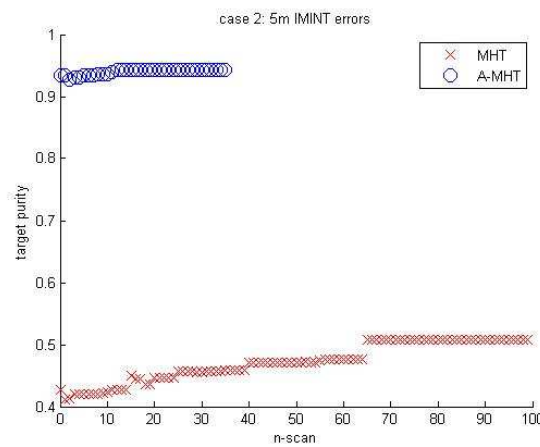

Figure 15. MHT and A-MHT target purity (harder scenario).

For the easier scenario, both the MHT and A-MHT achieve the same level of performance, though a larger n-scan is needed for the MHT to do so. The performance difference, and the potential of the A-MHT approach, becomes clearer with the harder scenario. First, as expected we see that performance is lower for both algorithms than it was with the easier scenario. On the other hand, the A-MHT reaches target purity and track purity levels that the MHT is unable to reach regardless of n-scan. Indeed, the theoretical optimality of MHT for sufficiently-large n-scan is not actually achieved, due to effective but suboptimal track-extraction and track-management logic.

For both the easier and harder scenarios, MHT and A-MHT processing times are fast (on the order of 1sec) for moderate values of n-scan. Nonetheless, processing time ultimately is a bottleneck for sufficiently complex scenarios.

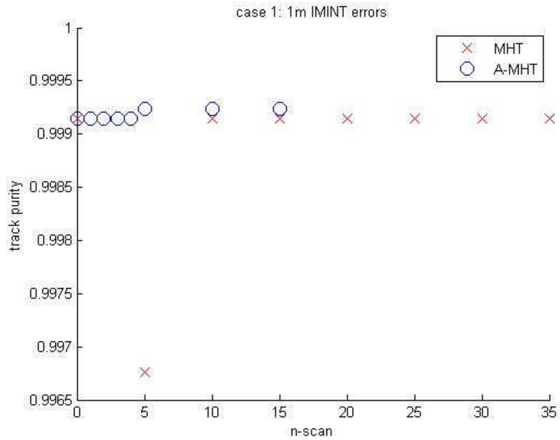

Figure 16. MHT and A-MHT track purity (easier scenario).

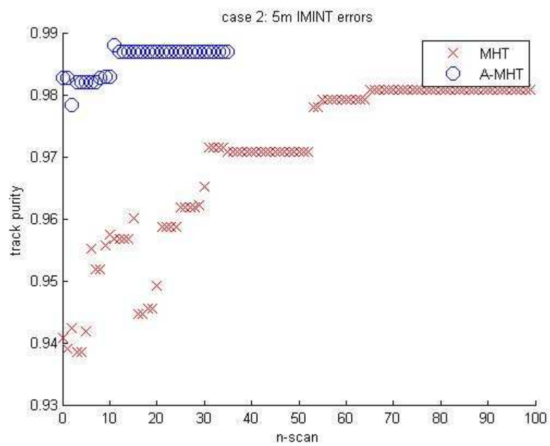

Figure 17. MHT and A-MHT track purity (harder scenario).

VI. CONCLUSIONS

This paper has introduced MCMC-based and MHT-based advances for multi-INT track fusion with complementary kinematic and identity sensors. The challenge is to exploit highly informative but sporadic identity information in complex settings with dense targets and many association possibilities for kinematic single-sensor tracks.

Both the MCMC-DF and the A-MHT show promising results against a highly efficient classical MHT baseline algorithm. We plan to compare performance of the three algorithms in larger-scale and more complex scenarios.

There are a number of directions for future work. More effective single-sensor tracking would maintain higher track continuity and only induce breaks at target group splits. Enhanced passive sensor modeling would allow for inactive, absent, or multiple emitters per target. Optimized management of the association decision-making sequence would improve A-MHT performance. Finally, both optimized MCMC-DF and A-MHT solutions require trajectory smoothing for enhanced forensic performance.

ACKNOWLEDGEMENT

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA). The views, opinions, and/or findings contained in this article are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

REFERENCES

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

[2] D. Reid, An Algorithm for Tracking Multiple Targets, *IEEE Transactions on Automatic Control*, vol. 24(6), December 1979.

[3] 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.

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

[5] W. Blanding, P. Willett, Y. Bar-Shalom, and S. Coraluppi, Multisensor Track Management for Targets with Fluctuating SNR, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 45(4), October 2009.

[6] S. Coraluppi, C. Carthel, M. Luettgen, and S. Lynch, All-Source Track and Identity Fusion, in *Proceedings of the National Symposium on Sensor and Data Fusion*, San Antonio TX, USA, June 2000.

[7] 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.

[8] 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.

[9] G. Castanon and L. Finn, Multi-Target Tracklet Stitching through Network Flows, in *Proceedings of the IEEE Aerospace Conference*, Big Sky MT, USA, March 2011.

[10] M. Martell, C. Smith, L. Finn, M. Pravia, and A. Bigelow, Track Labeling Adds Situational Awareness to a Graph-Based Fused Track Picture, in *Proceedings of the National Symposium on Sensor and Data Fusion*, Springfield VA, USA, October 2014.

[11] S. Coraluppi and C. Carthel, Advances in Multi-Target Filtering of Evasive Targets, in *Proceedings of the IEEE Aerospace Conference*, Big Sky MT, USA, March 2015.

[12] E. Fox, D. Choi, and A. Willsky, Nonparametric Bayesian Methods for Large Scale Multi-Target Tracking, in *Proceedings of the 40th Asilomar Conference on Signals, Systems, and Computers*, November 2007.

[13] S. Oh, Markov Chain Monte Carlo Data Association for Multi-Target Tracking, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54(3), March 2009.

[14] S. Coraluppi and C. Carthel, Stability and Stationarity in Target Kinematic Modeling, in *Proceedings of the IEEE Aerospace Conference*, Big Sky MT, USA, March 2012.

[15] 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), July 2014.

[16] S. Coraluppi, C. Carthel, and M. Guerriero, Hypothesis Structure in Enhanced Multiple-Hypothesis Tracking, in *Proceedings of the IEEE Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove CA, USA, November 2013.

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

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