Multiple-Hypothesis Tracking and Graph-Based Tracking Extensions

STEFANO P. CORALUPPI
CRAIG A. CARTHEL
ALAN S. WILLSKY

This paper reviews key elements in the development of *multiple-hypothesis tracking* (MHT), a leading paradigm for multitarget tracking, as well as *graph-based tracking* (GBT), a scalable version of MHT that has proven effective in kinematic track stitching applications. We introduce a novel MHT/GBT algorithm that we denote as *multi-INT GBT* (MI-GBT). It provides computational benefits over classical MHT, while allowing for static components of the target state that classical GBT does not. Thus, the MI-GBT provides an effective method for multisensor feature-aided track fusion with disparate sensors. We quantify the improved performance over the MHT solution in Monte Carlo studies.

Manuscript received January 28, 2019; revised May 8, 2019, October 30, 2019, and November 5, 2019; released for publication 0, 2019.

Refereeing of this contribution was handled by Jason L. Williams.

S. P. Coraluppi and C. A. Carthel are with Systems and Technology Research, 600 West Cummings Park, Woburn, MA 01801, USA (E-mail: Stefano.Coraluppi@STResearch.com, Craig.Carthel@STResearch.com).
A. S. Willsky is with the Department of Electrical Engineering, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA (E-mail: willsky@mit.edu).

1557-6418/19/$17.00

I. INTRODUCTION

Many approaches have been developed to address the *multitarget tracking* (MTT) problem, whereby an unknown and time-varying set of objects is to be tracked while contending with unknown measurement origin, missed detections, and false alarms [1]. *Multiple-hypothesis tracking* (MHT) was first posed in hypothesis-oriented form [2] and was later shown to admit hypothesis factorization (assuming Poisson-distributed births and clutter) and a more efficient track-oriented formulation [3], consistent with the *integer linear program* (ILP) framework of the MTT problem that had been proposed previously [4].

The ILP can be solved via relaxation approaches [5]–[7], and distributed processing can provide performance and robustness advantages in many settings [8]–[10]. The MHT paradigm has been generalized to consider object births without detection, enabling improved performance in dim-target settings [11]. More recently, extension to allow for multiple measurements per target per scan has been developed to deal with extended objects and multipath phenomena [12], [13].

When a simplifying Markovian (path-independence) assumption may be invoked, significant computational gains can be achieved. In the *hypothesis-oriented MHT* (HO-MHT) context, the simplified formulation may be solved by use of the Viterbi algorithm on an expanded trellis [14]. The more general treatment, with missed detections and clutter, was addressed in a series of papers culminating in [15]. Application of the approach to track-level inputs is discussed in [16]. While these are valuable contributions, unfortunately these approaches do not scale well when the numbers of measurements and targets are large. Also, these papers do not contend with target birth and death phenomena, which appear somewhat cumbersome to include.

Shortly after the publication of [14], the same simplification was introduced in the track-oriented MHT (TO-MHT) context [17]. This established the *graph-based tracking* (GBT) paradigm for MTT. The approach is not generally adopted for remote-sensing applications (e.g., sonar or radar tracking), since the Markovian assumption is too strong in these settings. Nonetheless, GBT for track maintenance with missed detections and clutter is developed in [18]. Perhaps more significantly, application of GBT to track-level inputs is discussed in [19]; this represents an important contribution in that, for track-level kinematic data, the Markovian assumption is quite appropriate.

In some settings, the Markovian assumption inherent in both Viterbi and GBT methods is not appropriate. Indeed, some elements of the target state vector, e.g., object size, color, etc., may be fixed or slowly varying. In such cases, measurement sequences do not exhibit a path-independence property, except when these slowly varying elements are always observed. Accordingly, when feature measurements are temporally

TABLE I
Assumptions and Some Solution Approaches for MTT

Assumptions	General	Partially Markov data	Markov data
General	HO-MHT [2]	Not investigated	Viterbi [14]
Poisson targets and clutter	TO-MHT [3]	MI-GBT [22]	GBT [17]

sparse, until recently it was necessary to resort to MHT solutions.

Recent work has extended the GBT approach to deal with feature states. In [20], a novel GBT multicommodity flow approach is discussed. The approach assumes a known number of objects, a unique object of each type, single-sensor data, and a batch-processing formulation. The methodology is promising in that it leads to a much smaller ILP than with an MHT approach. In [21] and [22], we develop a similar approach—the *multi-INT GBT* (MI-GBT)—for the general multisensor MTT problem, with a temporally sparse identity sensor and one or more kinematic sensors. Refs. [21] and [22] discuss as well a *Markov chain Monte Carlo* (MCMC) approach for the multi-INT (i.e., disparate-sensor) track fusion problem that builds on the work in [23]; for our application with sparse identity data, MCMC provides a promising approach to solution refinement. Related investigations of generalized GBT methods include [24] and [25].

Most recently, in [26] we relax the unique-type assumption in the MI-GBT, allowing for multiple objects of each type, and explore performance for multitarget track maintenance. Here, we discuss a sliding-window approach to ILP formation and resolution, enabling scalable processing of lengthy scenarios that are otherwise computationally prohibitive with earlier, batch-processing solution methods.

Table I provides a summary view of some paradigms for the MTT problem, focusing on hard data association and labeled target tracking. (See [1] and [27] for a discussion of other methods.) Given the computational advantages of the TO-MHT approach that avoids the global-hypothesis enumeration inherent in HO-MHT, we have chosen not to investigate a hybrid HO-MHT/Viterbi approach. This paper makes further progress on the MI-GBT that provides a hybrid TO-MHT/GBT paradigm that exploits path independence when possible, and introduces hypothesis branching when necessary to contend with identity data.

This paper is organized as follows. In Section II, we review salient elements of MHT. Section III discusses both MHT and MI-GBT for the multi-INT track fusion problem with sparse identity data, allowing for multiple objects of each type and with sliding-window processing. Section IV describes performance results for MI-GBT compared to the MHT baseline solution. We establish the superior performance characteristics of MI-GBT over both GBT and MHT. Conclusions are provided in Section V.

II. MULTIPLE-HYPOTHESIS TRACKING

In MTT, we seek a set of trajectories over a sequence of times $t^k = (t_1, \ldots, t_k)$ that we may denote compactly by X^k. Each trajectory in this set has a time of birth, an evolution in target state space, and (possibly) a time of death. Hence, we are interested in identifying the time evolution of an unknown (and time-varying) number of objects. We observe a sequence of sets of measurements $Z^k = (Z_1, \ldots, Z_k)$. The usual simplifying assumption in the MTT problem formulation is that, with each sensor scan, each target gives rise to at most one measurement. It is not known which measurement originates from which object, and there are as well false measurements that are not target originated.

A. MAP Estimation and Hypothesis-Oriented MHT

In statistical estimation theory, it is well established that minimization of the Bayes risk with an underlying cost function that penalizes all estimation errors uniformly is achieved with the conditional mode. Stated another way, the minimum probability of error estimator is given by the *maximum a posteriori* (MAP) estimator [28]. However, use of the MAP criterion for the MTT problem, when applied directly to $p(X^k|Z^k)$, is conceptually problematic [29, pp. 494–500].

One may explicitly consider an explanation for the data, i.e., to specify which measurements are to be rejected as false and how target-originated measurements are to be associated. Let us denote by q^k one such global hypothesis or explanation. This leads to a probabilistic conditioning approach and the following expression for the multitarget posterior probability distribution $p(X^k|Z^k)$[1]:

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

Note that, for any MTT problem of reasonable size, the space of global hypotheses—and hence the summation in Eq. (1)—is enormous.

The MHT paradigm addresses both the conceptual difficulty associated with MAP estimation applied to $p(X^k|Z^k)$ and the computational difficulty associated with the representation given by Eq. (1). In particular, MHT seeks the MAP global hypothesis $\hat{q}^k$ and conditions on this global hypothesis to estimate the set of target trajectories while discarding competing global hypotheses. This is captured in the following equations:

$$\hat{q}^k = \arg\max_{q^k} p\left(q^k|Z^k\right), \qquad (2)$$

[1]Equation (1) is a conceptual expression; a more rigorous treatment with the *random finite set* formalism may be found in [30].

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

Note that X^k is representative of a set of trajectories, each with a time of birth, an evolution in state space, and (possibly) a time of death. We will not discuss here the random finite set treatment of the MTT problem; see [29] and [30] for more details.

Solving Eq. (3) entails the solution to a set of smoothing problems. Most MTT approaches include recursive filtering but do not focus on trajectory smoothing. Indeed, while useful for output reporting, trajectory smoothing does not aid in data association; i.e., it does not contribute to solving Eq. (2).

Though solving Eq. (2) is not conceptually problematic, it remains computationally prohibitive. In practice, most MHT implementations consider a sliding-window formulation and *resolve* (i.e., select) global hypotheses with some delay. Having solved Eq. (2), solving Eq. (3) amounts to solving a set of filtering problems with no measurement-origin uncertainty. It is often beneficial to decouple data association and track extraction (see [31]).

Computational and real-time constraints require that we adopt a recursive formulation of $p(q^k|Z^k)$. The following expression may be derived:

$$p\left(q^k|Z^k\right) = \frac{p\left(Z_k|Z^{k-1}, q^k\right) p\left(q^k|q^{k-1}\right) p\left(q^{k-1}|Z^{k-1}\right)}{p\left(Z_k|Z^{k-1}\right)}. \quad (4)$$

This is the global-hypothesis recursion that expresses $p(q^k|Z^k)$ as a function of $p(q^{k-1}|Z^{k-1})$ and the current scan of data Z_k.

B. Track-Oriented MHT

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

The Poisson assumption is reasonable in many settings. We consider a continuous-time process with exponentially distributed target interarrival (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:

$$\mu_b(t) = \frac{\lambda_b}{\lambda_\chi}\left(1 - e^{-\lambda_\chi t}\right), \quad (5)$$

$$p_\chi(t) = 1 - e^{-\lambda_\chi t}. \quad (6)$$

For simplicity, in the following we will omit the time interval t and use the birth rate and death probability μ_b and p_χ, respectively. Note that the Poisson birth process has an intuitively appealing independence property, whereby the numbers of births in temporally nonoverlapping intervals are independent random variables [32]. Similarly, the Poisson false alarm assumption (with mean Λ) characterizes clutter statistics in many application domains. For target-originated measurements, we assume that, at each scan, each target is detected with probability p_d.

Let τ be the number of tracks in the parent global hypothesis q^{k-1} at time t_{k-1}, let $r = |Z_k|$ be the number of measurements in the current scan at time t_k, and let b, χ, and d be the number of target births, deaths, and measurement updates in global hypothesis q^k at time t_k, respectively.

We now express the global-hypothesis recursion given by Eq. (4) in detail. First, let us consider the factor $p(q^k|q^{k-1})$. For this, we introduce the auxiliary variable ψ_k that specifies the number of births b, the number of target deaths χ, and the number of targets with measurement update d. We use the following conditioning approach that relies on ψ_k:

$$p\left(q^k|q^{k-1}\right) = p\left(\psi_k|q^{k-1}\right) p\left(q^k|q^{k-1}, \psi_k\right). \quad (7)$$

The first factor in Eq. (7) denotes the probability of observing b target births, χ deaths, and d measurement updates from τ targets, and $r - d - b$ false alarms (to account for all remaining measurements). This may be written as follows, noting that we rely on 1) the Poisson distribution to account for b (detected) births and $r - d - b$ false alarms, as well as 2) the binomial distribution for the probability of observing some number of successes in a set of independent trials—this is relevant to the factors that account for χ deaths from τ targets and d detections from the surviving $\tau - \chi$ targets:

$$\begin{aligned} p\left(\psi_k|q^{k-1}\right) = {} & \binom{\tau}{\chi} p_\chi^\chi (1 - p_\chi)^{\tau-\chi} \cdot \binom{\tau - \chi}{d} \\ & \times p_d^d (1 - p_d)^{\tau-\chi-d} \cdot \frac{(p_d \mu_b)^b \exp(-\mu_b)}{b!} \\ & \cdot \frac{\Lambda^{r-d-b} \exp(-\Lambda)}{(r-d-b)!}. \end{aligned} \quad (8)$$

The second factor in Eq. (7) denotes the probability of a particular global hypothesis, conditioned on the parent hypothesis and on the cardinalities associated with ψ_k. As all association probabilities have the same *a priori* probabilities, this factor can be written as follows. Note that the denominator terms quantify the number of ways of selecting the target deaths, the number of ways of selecting which tracks to update, the number of ways of selecting measurements and assigning them to tracks (where ordering matters), and the number of ways of selecting birth measurements among the remaining $r - d$

measurements.

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

Combining Eqs. (8) and (9) according to Eq. (7) yields

$$\begin{aligned}p\left(q^k|q^{k-1}\right)=&\left\{\frac{\exp\left(-\mu_{\mathrm{b}}-\Lambda\right)\Lambda^r}{r!}\right\}p_\chi^\chi\\&\cdot\left((1-p_\chi)(1-p_{\mathrm{d}})\right)^{\tau-\chi-d}\left(\frac{(1-p_\chi)\,p_{\mathrm{d}}}{\Lambda}\right)^d\\&\times\left(\frac{p_{\mathrm{d}}\mu_{\mathrm{b}}}{\Lambda}\right)^b.\end{aligned} \tag{10}$$

The factor $p(Z_k|Z^{k-1},q^k)$ in Eq. (4) accounts for the probability of observing a set of measurements given a global hypothesis. It is a product over filter residual scores; hence, it may be written as follows, where, under q^k, J_{d} is the set of track update measurements, J_{fa} is the set of false alarms, J_{b} is the set of target birth measurements, $f(\cdot|Z^{k-1},q^k)$ is the conditional probability distribution of a target measurement (with no conditioning in the case of object birth), and $f_{\mathrm{fa}}(\cdot)$ is the distribution of false alarms in measurement space:

$$\begin{aligned}p\left(Z_k|Z^{k-1},q^k\right)=&\prod_{z_j\in J_{\mathrm{d}}}f\left(z_j|Z^{k-1},q^k\right)\prod_{z_j\in J_{\mathrm{b}}}f\left(z_j\right)\\&\cdot\prod_{z_j\in J_{\mathrm{fa}}}f_{\mathrm{fa}}\left(z_j\right).\end{aligned} \tag{11}$$

Equations (10) and (11) may be substituted into Eq. (4), resulting in the following TO-MHT recursion, in which we denote by C the factor that is common to all global hypotheses. (This common factor need not be computed for MAP estimation, and indeed its evaluation would rely on $p(Z_k|Z^{k-1})$, requiring summation over all global hypotheses.) The restriction that each measurement be used at most once in track formation, and that all measurements be accounted for, is captured in Eq. (12c), where J_Z is the set of indices for measurement set Z_k.

$$\begin{aligned}p\left(q^k|Z^k\right)=&\,p_\chi^\chi\left(\left(1-p_\chi\right)\left(1-p_{\mathrm{d}}\right)\right)^{\tau-\chi-d}\\&\cdot\prod_{j\in J_{\mathrm{d}}}\frac{\left(1-p_\chi\right)p_{\mathrm{d}}f\left(z_j|Z^{k-1},q^k\right)}{\Lambda f_{\mathrm{fa}}\left(z_j\right)}\\&\cdot\prod_{j\in J_{\mathrm{b}}}\frac{p_{\mathrm{d}}\mu_{\mathrm{b}}f\left(z_j\right)}{\Lambda f_{\mathrm{fa}}\left(z_j\right)}\cdot C\cdot p\left(q^{k-1}|Z^{k-1}\right),\end{aligned} \tag{12a}$$

$$C=\frac{\left\{\frac{\exp(-\mu_{\mathrm{b}}-\Lambda)\Lambda^r}{r!}\right\}\prod_{z_j\in Z_k}f_{\mathrm{fa}}\left(z_j\right)}{p\left(Z_k|Z^{k-1}\right)}, \tag{12b}$$

$$J_{\mathrm{d}}\cap J_{\mathrm{b}}=\oslash,\;\; J_{\mathrm{d}}\cap J_{\mathrm{fa}}=\oslash,\;\; J_{\mathrm{b}}\cap J_{\mathrm{fa}}=\oslash,\;\; J_{\mathrm{d}}\cup J_{\mathrm{b}}\cup J_{\mathrm{fa}}=J_Z. \tag{12c}$$

Equation (12) is of fundamental importance in that it factors the global hypothesis score into (dimensionless) track scores. Accordingly, it is unnecessary to consider each global hypothesis probability explicitly. Indeed, up to the hypothesis-independent factor C, a global hypothesis probability may be evaluated as a product over local hypothesis factors. This in turn allows the determination of $\hat{q}^k$, the solution to Eq. (2), without explicit enumeration of global hypotheses.

Thus, the TO-MHT formalism results in an ILP, with an objective function that may be expressed compactly by Eq. (13), where the cost c_i associated with track hypothesis x_i results from statistics-associated targets and sensors; the variable $x_i\in\{0,1\}$ may be understood as an indicator variable that corresponds to selecting a track hypothesis when setting $x_i=1$. The sum is over all track hypotheses within a hypothesis reasoning window:

$$J=\sum_i c_i x_i, \tag{13}$$

$$Ax\le b. \tag{14}$$

In addition to the objective, the ILP includes constraints captured by Eq. (14) that require that each measurement be used at most once in track formation, and that each resolved track from the start of the reasoning window be accounted for.

III. MULTI-INT TRACK FUSION

Section II described TO-MHT (hereafter, MHT) for detection-level MTT. We now consider the MTT problem, downstream of single-sensor trackers. That is, upstream association decisions have been made, and we assume negligible residual false alarms. The challenge is to perform correct track association over time and across sensors. We assume that tracks are composed of sequences of measurements, so that optimal filtering can be performed without the need to contend with correlated state estimates due to common target process noise [27]. We are interested in both real-time and forensic settings. The principal challenge is how to contend with temporally sparse identity information that is crucial to exploit for high-performance association decisions. After providing some modeling details, we will first discuss the conventional MHT solution and then describe the MI-GBT approach.

A. Some Modeling Details

As noted earlier, we model target existence via a Poisson birth–death process; see Eqs. (5) and (6). For simplicity, we will discuss our work in the context of linear Gaussian dynamics and measurements, though the solution methodologies are applicable more broadly. (The model described here is what we use for the simulation results in Section IV.) Specifically, we will assume independent target dynamics according to a stable,

stationary generalization to *nearly constant velocity* (NCV) motion, as given by a second-order *Ornstein–Uhlenbeck* (OU) process [33].

Each object is of a fixed type, with a probability distribution over the finite set of types given by the vector p_{type}, with element $p_{\text{type}}(i)$ being the probability that an object is of type i. We may handle unique objects of type i by using a small value for $p_{\text{type}}(i)$. This modeling approach is more robust than disallowing multiple objects, as the multiplicity is sometimes necessary to contend with potentially erroneous sliding-window association decisions. Also, all objects must be of some type; hence, we generally include the type "other" to include all those that are not of specific interest.

Kinematic tracks are composed of linear measurements with additive Gaussian noise with $v_k \sim N(0, R_k)$:

$$y_k = C_k x_k + v_k. \tag{15}$$

As noted earlier in the MHT derivation, at each scan targets are detected with probability p_{d}. For simplicity, we do not consider motion-dependent (or, more generally, state-dependent) detection statistics. Further, we do not consider identity-dependent detection statistics, as when objects of certain types are easier to detect than others.

We assume that identity sensors differ from kinematic sensors in two key respects. First, we assume a low revisit time between scans; hence, the identity detections are not associated over time. Second, detections include both kinematic information and precise target-type information. The type information is highly informative but does not provide association information, as there may be multiple objects of the same type. Detection and localization quality (p_{d} and R_k) differ for kinematic and identity sensors. We do not consider false alarms from the identity sensor; this is reasonable in settings where *automatic target recognition* is performed to provide object-type information and to reject spurious detections.

B. MHT Approach

The detection-level TO-MHT recursion given by Eq. (12) yields a dimensionless likelihood ratio associated with each track hypothesis. Normalization is with respect to a null hypothesis whereby all measurements are false alarms. For a track with index i, the negative log of this score yields the coefficient c_i in the objective function to be minimized, as given by Eq. (13).

For track-level association, since all tracks (including identity singleton tracks) are assumed to be target originated, we do not normalize with respect to the same null hypothesis. We may still utilize dimensionless track scores by normalizing with respect to another null hypothesis, whereby all tracks are unassociated.

Let z_j represent a track (i.e., a sequence of previously associated measurements) and let $L(z^n)$ denote the track likelihood associated with a sequence of tracks $z^n = (z_1, \ldots, z_n)$. Note that $L(z^n)$ is the (unnormalized) local hypothesis contribution to the global hypothesis probability. If this sequence corresponds to the ith track hypothesis, we may express the coefficient c_i in Eq. (13) as follows:

$$c_i = -\log L(z^n). \tag{16}$$

The likelihood $L(z^n)$ accounts for target birth, a sequence of detection and missed-detection events, and (possibly) a target death. This score can be computed recursively based on the following probabilistic conditioning:

$$L(z^n) = L(z_1) \prod_{j=2,\ldots,n} L\left(z_j | z^{j-1}\right). \tag{17}$$

We may alternatively adopt a dimensionless track score as given by the following. This is advantageous when solving Eq. (13) with fast greedy track selection methods, in lieu of a relaxation approach.

$$c_i = -\log \frac{L(z^n)}{\prod_{j=1,\ldots,n} L(z_j)}. \tag{18}$$

As with detection-level MHT, most nontrivial track fusion problems entail hypothesis-space reduction via sliding-window processing. That is, with some temporal delay, we resolve ambiguity and identify a single global hypothesis by solving an appropriately defined ILP as in Eq. (13). Then, we ingest further data for processing, and solve a new ILP. The hypothesis tree depth is generally identified as the number of scans of data (n-scan) between the resolved time and current time [12], [31]. *Global nearest-neighbor* (GNN) processing corresponds to n-scan $= 0$ [27].

The MHT track fusion capability is quite general and allows for an arbitrary number of kinematic and identity sensor inputs. Key assumptions include target independence (both existence and dynamics) and correct (but not necessarily complete) upstream association decisions. Hence, we may associate multiple tracks from the same sensor, provided there is no scan with measurements from more than one track.

C. GBT Approach

Computational simplifications may be achieved if local (track) hypotheses satisfy a path-independence assumption, whereby the track likelihood may be factored with pairwise contributions to the likelihood.

The fundamental path-independence assumption that we introduce is appropriate for single-sensor kinematic track-level data. Again, let z_i represent a track (i.e., a sequence of previously associated measurements) and let $L(z^n)$ denote the track likelihood associated with a sequence of tracks. The path-independence assumption

amounts to the following:

$$L(z^n) = L(z_1) \prod_{i=2,\ldots,n} L\left(z_i|z^{i-1}\right)$$
$$\approx L(z_1) \prod_{i=2,\ldots,n} L(z_i|z_{i-1}). \quad (19)$$

The path-independence assumption enables a graph-based representation of the MTT problem with pairwise costs derived from conditional likelihoods.

Under the GBT formalism, we consider a set of kinematic tracks that we represent by a set of nodes V. We consider also source and sink nodes, denoted by v_0 and v_∞, respectively. We define the augmented set of nodes by $\bar{V} = V \cup \{v_0, v_\infty\}$. We consider a directed graph $G = (\bar{V}, A)$, where A is a set of edges. For each feasible edge $(i, j) \in A$, i.e., with no temporal overlap between the corresponding tracks, we define the cost c_{ij} by the negative log conditional likelihood:

$$c_{ij} = -\log L(v_j|v_i). \quad (20)$$

Note that the likelihood $c_{0j} = L(v_j|v_0) = L(v_j)$ accounts for target birth. As all track likelihoods account already for target death, we have $L(v_\infty|v_i) = 1$.

The kinematic GBT formulation leads to the following ILP:

$$J = \sum_{(i,j)\in A} c_{ij}x_{ij}, \quad (21)$$

$$x_{ij} \in \{0, 1\} \quad \forall (i, j) \in A, \quad (22)$$

$$\sum_{i:(i,j)\in A} x_{ij} = 1 \quad \forall v_j \in V, \quad (23)$$

$$\sum_{j:(i,j)\in A} x_{ij} = 1 \quad \forall v_i \in V. \quad (24)$$

We seek the solution that minimizes the objective (21) subject to constraints (22)–(24). Equations (23) and (24) ensure that all nodes be used exactly once, and that flow balance be achieved.

The resulting MAP estimation problem for global hypothesis q^k is over a smaller space than MHT. Indeed, here a second form of factorization is invoked in addition to that of TO-MHT, based on Eq. (19). By exploiting this factorization, we avoid the enumeration of track hypotheses; rather, the ILP is posed over pairwise-association variables.

A nice feature of the single-sensor GBT formulation is that it results in an ILP with special structure: it may be expressed as a *min-cost network flow* (MCNF) problem or, equivalently, as a *bipartite matching* problem. Thus, the problem admits an integer solution and faster solution than a general ILP [34].

As with MHT, we may wish to define dimensionless scores analogous to Eq. (18). Here, we may adopt the following in lieu of Eq. (20):

$$c_{ij} = -\log \frac{L(v_j|v_i)}{L(v_j)} = -\log \frac{L(v_i, v_j)}{L(v_j)L(v_i)}. \quad (25)$$

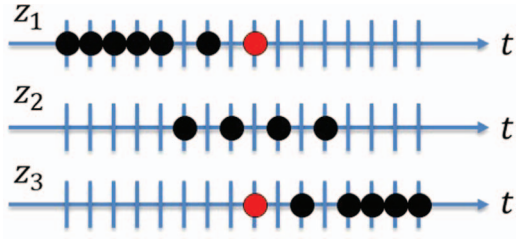

Fig. 1. Tracks z_1 and z_3 cannot originate from the same target, as there is a scan where both have a measurement (shown in red). Avoiding the association in GBT requires a strict condition for pairwise-association feasibility.

Note that the requirement for association feasibility mentioned above—no temporal overlap between the corresponding tracks—is stricter than that in MHT. This is necessary due to the pairwise nature of track scoring. Consider the example in Fig. 1. MHT computes $L(z^3)$ precisely; it must necessarily be zero since tracks z_1 and z_3 share a relevant sensor scan. On the other hand, GBT would allow for the association of the tracks since pairwise feasibility is maintained, were we not to impose the stricter *no temporal overlap* condition.

Indeed, GBT will assume $L(z^3) \approx L(z_3|z_2)L(z_2|z_1)$ $L(z_1)$. In this example, both $L(z_3|z_2)$ and $L(z_2|z_1)$ are nonzero, while $L(z^3)$ must be zero. Hence, the approximation is potentially poor for temporally overlapping tracks, not due to a poor kinematic filtering approximation but rather due to an incorrect accounting for measurement-cardinality information.

D. MI-GBT Approach

We wish to leverage the GBT approach while allowing for multiple kinematic and identity sensors, as well as for identity tracks that violate the Markovian assumption on the data.

Let us first address the need to process multiple sensors, as illustrated in Fig. 2. We do so by exploiting MHT processing with a dedicated kinematic-fusion stage, yielding a single, fused kinematic sensor feed. Only nontemporally overlapping fused tracks (i.e., the strict condition above) will be feasibly associated in downstream MI-GBT processing.

It is worth emphasizing that kinematic processing may introduce undesired measurement association errors. This is not problematic when one or more objects remain in close proximity. On the other hand, when a group of objects splits into two or more, it is important to fragment kinematic-only tracks to enable high-confidence stitching in downstream processing, with the aid of identity-sensor data.

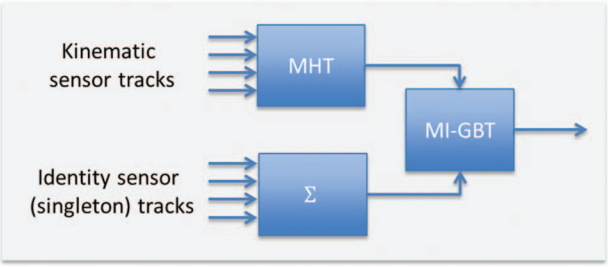

Fig. 2. The MI-GBT is seemingly more restrictive as it requires a single kinematic and identity feed. However, this can be addressed with preprocessing that exploits MHT kinematic tracking, as well the recasting of multiple-measurement identity tracks as unassociated measurements of a vanishingly small type probability, all due to one sensor.

Likewise, there may be multiple identity sensors. However, since we assume single-measurement identity tracks, there is no loss of generality in considering all tracks as originating from a single identity sensor. For identity measurements that potentially observe different aspects of the target state, e.g., object color, object size, etc., we may recast the formulation as a single identity sensor with vector-valued measurements. For simplicity, here we consider scalar object types.

It is important to note that the *unionizing* operation on identity measurements preserves the scan structure of the data. That is, if two identity sensors have sensor scans at the same time, the two scans are kept distinct for downstream processing. The point-target assumption remains crucial as in MHT; i.e., we have at most one measurement per target per scan.

The second need is to handle identity measurements that violate the Markovian approximation in Eq. (29). Indeed, while past kinematic tracks are not relevant to future kinematic association scores, the same is not true for identity tracks that specify the object type. Our approach will be to define an ILP that corresponds to a *multilayer* graph, one for each object type. Path independence holds within a layer of the track, but not across layers. Before defining the ILP, we show an illustrative example.

The advantage of the architecture in Fig. 2 is that we exploit the MHT for what it performs well, namely, multisensor kinematic tracking where small hypothesis tree depths are effective. We defer the disparate-sensor fusion problem, where MHT is severely challenged computationally, to generalized GBT processing.

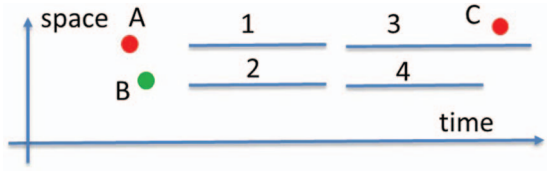

Fig. 3. A multi-INT track fusion example.

E. MHT and MI-GBT Structure

Consider the notional example in Fig. 3. We have an unknown number of objects giving rise to three identity tracks and four kinematic tracks, indicating that there are red and green objects present.

For simplicity of exposition, we assume a forensic surveillance problem, in the sense that all data have been received at the processing center. This enables a compact representation of the association spaces associated with competing solution approaches, with nodes representing input track. Regardless of whether online or forensic analysis is to be performed, the data-association process will necessarily rely on sliding-window processing for computational tractability.

Under the simplifying assumption of no never-observed objects (the usual assumption in MHT), there are at most seven objects present, and there are at least two targets (one red, one green). The MAP solution will depend on target and sensor statistical assumptions, and on the measured data themselves. The corresponding data structure associated with TO-MHT processing may be represented as illustrated in Fig. 4. Note that, for simplicity, we have expressed each path in the MHT track forest as a sequence of tracks. This is not fully reflective of the actual processing sequence, since data are ingested and processed in proper time order. As an example, in the leftmost path, some measurements associated with track 3 follow track C. Track filtering and scoring is performed in proper time sequence.

The MAP solution associated with MHT processing will be that set of paths that accounts for all the data, while minimizing the sum of negative log likelihoods as in Eq. (26). Alternatively, we may use likelihood ratios as in Eq. (28).

Figs. 5 and 6 illustrate the GBT and MI-GBT graphs, respectively. For simplicity, in both graphs we have not drawn the termination node, nor the termination edges from each node to the termination node. (There is no edge directly from the birth node to the termination node.) By default, all edges are *directed* (downward), ex-

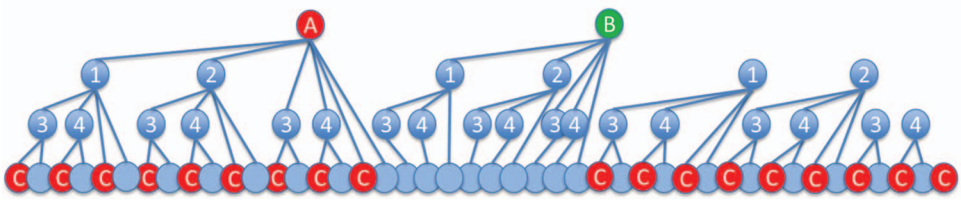

Fig. 4. MHT track forest for the example in Fig. 3.

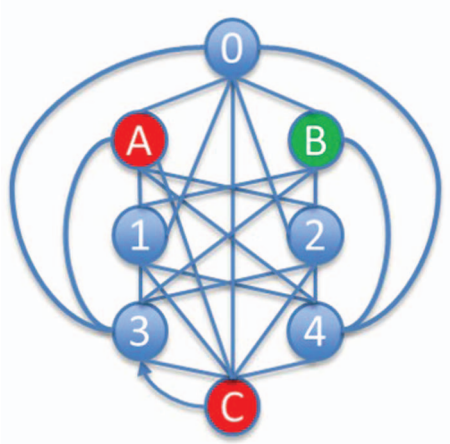

Fig. 5. GBT graph topology for the example in Fig. 3. While the representation supports computationally efficient solutions, these cannot exploitation crucial object-type information.

cept where explicitly denoted, i.e., the second edge between track 3 and node C.

There are several points to note. First, while much more compact than the MHT structure, the GBT graph topology does not allow for exploitation of target-type information (except for the lack of an edge directly connecting node B and node C). There is too much simplification in the problem formulation, so that feature information cannot help the data-association process. Indeed, target-type information does not satisfy the Markovian property that applies to kinematic data.

The MI-GBT structure is more compact than that of MHT, due to the simplifying path-independence assumption. Thus, for instance, each node associated with track 4 is a sufficient representation of kinematic information on the target, without the need for expressing from whence the target originates. At the same time, it is crucial to maintain graph layers (or subgraphs) associated with distinct object types. No flow is permitted between subgraphs, except for flow from the null subgraph to the object-type subgraphs.

The null subgraph captures objects for which no object-type information is known. In the example, these objects may well be red or green—we do not know. Colored edges indicate associations where type information is introduced.

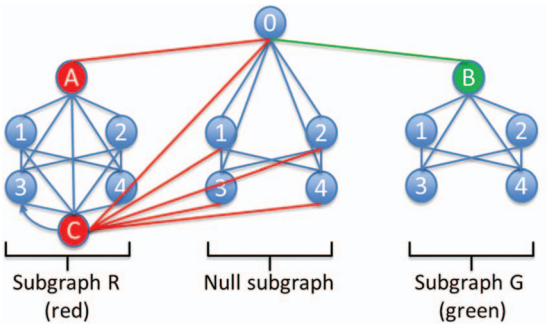

Fig. 6. MI-GBT graph topology for the example in Fig. 3. Colored edges indicate when type information is inserted into a tracking solution. Constraints in the ILP ensure that equivalent vertices are used only once.

The MI-GBT topology distinguishes between nonoverlapping tracks and overlapping (nested) tracks. Indeed, the only temporal overlap that we allow is that between kinematic and identity tracks. This requires the use of double (directional) edges between such tracks when there is temporal overlap. Either both or neither is to be selected.

In the example in Fig. 3, there is temporal overlap between track 3 and track C; hence, if these are associated, we want flow from track 3 to track C, and then from track C to track 3. In this manner, track 3 is relevant to (kinematic) association with any preceding or subsequent tracks, while track C is not. Indeed, it is always the identity track that is temporally nested in the kinematic track, not vice versa.

There is an interesting question of how best to break the symmetry whereby flow might in principle go into track C, then to track 3, and then back to track C. This is not admissible and can be avoided by introducing an inequality constraint that forces flow into node 3 of subgraph R, if the cycle from track 3 to track C is active.

Note that the return flow from track C to track 3 must necessarily be in subgraph R, since the object of interest is necessarily of type R (red). Note also that, for tracks 1, 2, and 4, there is no need for bidirectional flow to any identity track, since none of these identity tracks is temporally nested in these kinematic tracks. Thus, in particular, if there is association between one of the tracks (1, 2, 4) with track C, any associations with subsequent tracks would be from track C.

The data fusion that the MI-GBT permits—that between kinematic tracks and single-measurement identity tracks—does pose a potential hazard. Indeed, we must introduce a mechanism to specifically disallow the fusion of multiple identity measurements at the same time with the same kinematic track. Once more, this can be achieved with a suitable inequality constraint.

The need for this constraint emphasizes a fundamental limitation of graph-based reasoning. It is inherently myopic, in the sense that it reasons only over pairwise-association scores. While this is reasonable for kinematic information, it is not so for cardinality information whereby we wish to disallow fusion of multiple measurements at the same time from the same sensor, hence the need for the constraints noted above. Nor is pairwise reasoning sufficient to exploit object-type information, hence the need for the expanded (multilayer) graph structure in MI-GBT that conventional GBT lacks.

Sliding-window n-scan processing to resolve global hypotheses may be performed on the MI-GBT data structure, in analogous fashion to how it is performed in MHT [1].

As an extension to the example above, consider the scenario in Fig. 7, where we observe an additional kinematic track. The corresponding MI-GBT topology is given in Fig. 8.

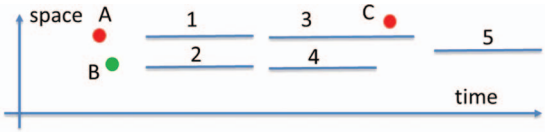

Fig. 7. An extension to the previous multi-INT example.

F. MI-GBT ILP

While our ILP implementation corresponds to the illustration in Fig. 6 (or Fig. 8, for the larger example), it will be easier to describe the ILP associated with the equivalent representation illustrated in Fig. 9. In practice, we prefer the structure in Fig. 6 as we only spawn subgraphs when identity tracks of the corresponding type are to be processed, and not earlier. There is a one-to-one correspondence between the two representations.

Let us denote by V the set of kinematic tracks, and by v_0 and v_∞ the source and sink vertices, respectively. We denote by W_k the set of identity (singleton) tracks of type k, with $k = 1, \ldots, K$.

In the MI-GBT, each kinematic track node may appear on multiple graph subgraphs (or layers). It will appear on all layers if spatial and kinematic gating are not performed; we assume it does so, for ease of presentation. On the other hand, each identity track appears only in one graph layer, e.g., red measurements only in the red layer, etc. We denote by $\bar{V}_k = V \cup W_k \cup \{v_0, v_\infty\}$ the set of vertices in the kth subgraph G_k, and by A_k the set of edges in subgraph G_k, i.e., $G_k = (\bar{V}_k, A_k)$. G_0 is the null subgraph, where no identity tracks are present. The full set of identity tracks is given by $W = \underset{k=1,\ldots,K}{\cup} W_k$. We have $W_0 = \oslash$, as there are no null-type identity measurements.

Recall that in Eq. (25), pairwise scores were indexed by two tracks. Now we have a third index to account for object type. Letting K be the number of object types, and denoting by $k = 0$ the null index (i.e., no object-type information), we have the following, where the likelihood function is understood not to include any contribution from object type. Note there is no need for edges with $i = j$.

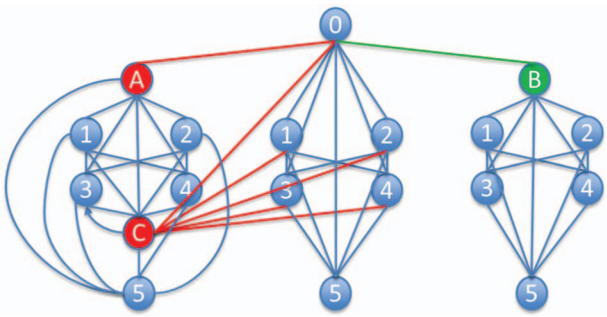

Fig. 8. The MI-GBT topology for the scenario in Fig. 7. Constraints in the ILP ensure that equivalent vertices are used only once.

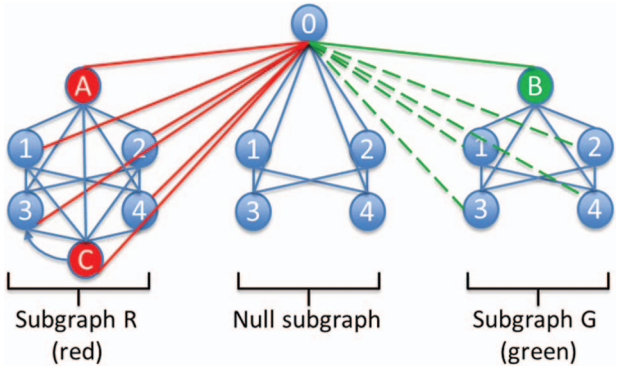

Fig. 9. An equivalent MI-GBT graph topology for the example in Fig. 3. Constraints in the ILP ensure that equivalent vertices are used only once. While our implementation matches the topology of Fig. 6, the ILP is easier to describe for this topology. Dashed edges are not strictly necessary (the solution will not include these edges) but are included for completeness.

$$c_{ijk} = \begin{cases} -\log L\left(v_j\right), & i = 0,\ k = 0, \\ -\log\left(p_{\text{type}}\left(k\right) L\left(v_j\right)\right), & i = 0,\ k \neq 0, \\ -\log L\left(v_j | v_i\right), & i \neq 0,\ j \neq 0, \\ 0, & j = 0. \end{cases} \tag{26}$$

As before, we may alternatively use dimensionless track score based on likelihood ratios. In this case, we have the following:

$$c_{ijk} = \begin{cases} 0, & i = 0,\ k = 0, \\ -\log p_{\text{type}}\left(k\right), & i = 0,\ k \neq 0, \\ -\log \frac{L\left(v_j | v_i\right)}{L\left(v_j\right)}, & i \neq 0,\ j \neq 0, \\ 0, & j = 0. \end{cases} \tag{27}$$

Using either Eq. (26) or Eq. (27), we may then express the ILP as follows:

$$J = \sum_{k=0,\ldots,K} \sum_{(i,j,k) \in A_k} c_{ijk} x_{ijk}, \tag{28}$$

$$x_{ijk} \in \{0, 1\} \quad \forall\, (i, j, k) \in A_k, \quad k = 0, \ldots, K, \tag{29}$$

$$\sum_{k=0,\ldots,K} \sum_{i:(i,j,k) \in A_k} x_{ijk} = 1 \quad \forall v_j \in V \cup W, \tag{30}$$

$$\sum_{j:(j,i,k) \in A_k} x_{jik} - \sum_{j:(i,j,k) \in A_k} x_{ijk} = 0 \quad \forall v_i \in V \cup W_k, \quad k = 0, \ldots, K. \tag{31}$$

We seek the solution that minimizes the objective (28) subject to conditions (29)–(31). Equation (30) ensures that all nodes be used exactly once. Equation (31) ensures that flow balance be achieved in each track node in each subgraph.

Additional constraints are needed to complete the ILP formulation. Indeed, in Section III-E we identified two concerns that must be addressed via appropriate

constraints. The first is that identity track v_j may be temporally nested within track v_i, so that both edge variables x_{ijk} and x_{jik} are defined for $k = 1, \ldots, K$. In this case, we must ensure that both edge variables (or neither) are set to unity in subgraph G_k. Furthermore, if both are set to unity, there must be flow into v_i from another node in addition to v_j. This ensures that, in G_k, flow is into v_i, from v_i to v_j and back, and out of v_i. This is captured in the following constraints:

$$(j, i, k) \in A_k \Rightarrow x_{ijk} = x_{jik} \quad \forall\, (i, j, k) \in A_k, \quad k = 1, \ldots, K, \tag{32}$$

$$(j, i, k) \in A_k,\, v_j \in W_k \Rightarrow \sum_{l:(l,i,k)\in A_k \setminus (j,i,k)} x_{lik} \geq x_{ijk}, \quad k = 1, \ldots, K. \tag{33}$$

A second concern is that we must exclude the association of multiple identity tracks in the same scan with the same kinematic track. This would violate the modeling assumption of at most one detection per target per scan. This can be achieved with an inequality constraint whereby for each identity-sensor scan and each (relevant) kinematic track at most one edge variable is unity.

Each identity scan is at a time $t_m \in t^n$, with t^n the sequence of identity-sensor scan times. Let us denote by $W(t_m)$ the corresponding set of identity (singleton) tracks. For each kinematic track $v_i \in V$, one or both of two cases apply, depending on whether the end of the kinematic track temporally precedes or follows the identity scan at time. Thus, we have

$$\forall v_i \in V,\ \forall t_m \in t^n,\ \begin{cases} \sum_{k=1,\ldots,K} \sum_{j:(i,j,k)\in A_k, v_j \in W(t_m)} x_{ijk} \leq 1, \\ \sum_{k=1,\ldots,K} \sum_{j:(j,i,k)\in A_k, v_j \in W(t_m)} x_{jik} \leq 1. \end{cases} \tag{34}$$

It is worth emphasizing that the MI-GBT solution is fully specified by the ILP defined by Eqs. (28)–(34). The graphical illustrations shown in Figs. 6, 8, and 9 are pictorial aids, but no not capture the required optimization constraints.

G. Solution Complexity

It is useful to have an approximate, analytical assessment of the computational complexity associated with MHT, GBT, and MI-GBT solutions to the multi-INT problem. Here, we estimate the size of the ILP associated with these paradigms. We denote by $\dim(x)$ the length of the solution vector in the objective—Eqs. (13), (21), and (28), respectively.

Given m sets of $|V|$ tracks, and with $|W|$ identity types, the GBT problem size is $|x| = O(m|V|^2)$, while the MI-GBT problem size is $|x| = O(m|V|^2(1 + |W|))$. Both compare favorably to the (track-oriented) MHT approach, for which problem size is $|x| = O(|V|^{m+1}(1 + |W|))$. The solution time associated with the ILP is problem size dependent. Empirically, we observe low-order polynomial times as a function of the solution vector, typically $O(|x|^n)$ with small n for MHT and MI-GBT solutions based on LP relaxation, and $O(|x|^3)$ for the GBT based on MCNF or an equivalent bipartite matching formulation [34].

The MI-GBT provides a good trade-off with its ability to exploit object-type information (like MHT) while maintaining an efficient pairwise-cost formalism (like GBT). For a given hypothesis depth (n-scan), MHT will generally outperform MI-GBT. Likewise, for a given hypothesis depth, GBT will incur lower computational effort. We anticipate that, in disparate-sensor settings where kinematic Markovian assumptions are appropriate, MI-GBT will yield a better complexity versus performance operating curve than both MHT (which *does the right thing*, at great expense) and GBT (which cannot exploit type information).

IV. SIMULATION RESULTS

We now explore the performance of MHT and MI-GBT approaches to multi-INT track fusion. We focus on a simplified version of the general problem while including the key challenge that exposes the differences between the MHT, GBT, and MI-GBT solutions. This will allow us to gain intuition regarding the relative strengths of the methods. It will be of interest to conduct more general MTT performance analysis in subsequent studies.

We consider a fixed number of targets, with no target births or deaths. We assume high track-level detection performance. We simplify the problem further by assuming equivalent-measurement processing that leads to single-measurement kinematic tracks. Thus, our problem may be viewed as one for which we observe a sequence of measurement sets, each containing a detection on all targets, with no false measurements.

We assume that the identity sensor reports twice, at the start and at the end of the scenario. In the interim, we have a number of kinematic scans, each containing positional measurements on all targets. The identity sensor includes precise object-type information with each positional measurement. In general, there are multiple objects of each type, so the association of measurements from the two identity-sensor scans in not known.

We consider GBT, MHT, and MI-GBT solutions to this data-association problem. Ultimately, even for this simplified problem, one would want to compare GBT, MHT, and MI-GBT solutions for a common processing load. Since the computational complexity of MHT grows significantly as a function of scenario duration, we limit processing to an n-scan = 0 solution (i.e., no hypothesis depth) that amounts to GNN processing.

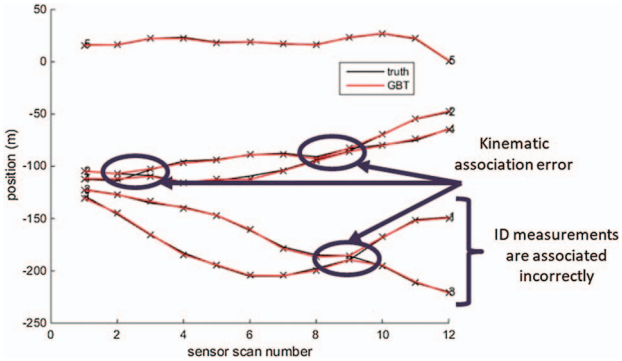

Fig. 10. GBT processing is more error-prone than MHT in kinematic association, due to the fundamental path-independence approximation that degrades kinematic-filtering accuracy. In the 1D case, GBT tracks never cross. Note that ID associations may violate type information: target 1 is of type A and target 3 is of type B.

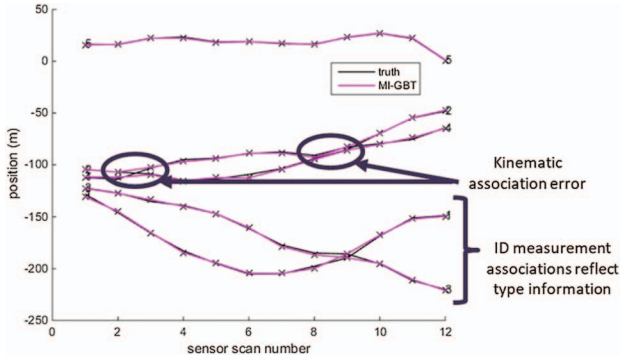

Fig. 12. MI-GBT may perform association errors when consistent with available ID data, but yields solutions that exploit all available ID measurements.

Figs. 10–12 illustrate one realization of 1D target trajectories (black lines), measurement data (black cross symbols), and GBT, MHT, and MI-GBT solutions. Target motion is according to our stable, stationary second-order OU process that generalizes the standard NCV motion. Indeed, note that the positional spread of the trajectories remains roughly the same over time; the same is true in velocity space. Positional measurements include additive Gaussian noise.

The first and last sensor scans are provided by the ID sensor. We consider a five-target scenario. Targets 1 and 2 are of type A, targets 3 and 4 are of type B, and target 5 is of type C. The ID sensor does not exhibit object-type measurement error, but the association of target measurements of the same type is unknown.

It is instructive to consider aspects of the solutions as illustrated in these figures. Note that the GBT solution incurs errors when targets cross, since the solution trajectories do not do so. This can be readily understood, as the GBT reasons over pairwise costs. In the 1D case, it is costly to associate measurements in such a way as to alter the relative ordering of the tracks.

The GBT solution cannot exploit ID information except when there are sequential ID-sensor scans. Hence, when there are kinematic-sensor scans between ID-sensor scans, the association of ID measurements is error-prone. Unlike the GBT solution, the MHT solution relies on recursive filtering so that crossing targets are handled properly in most cases. Corrective action is taken when the second ID-sensor scan is received. Note that full corrective active with MHT may not be possible when kinematic gating disallows sufficiently unlikely associations.

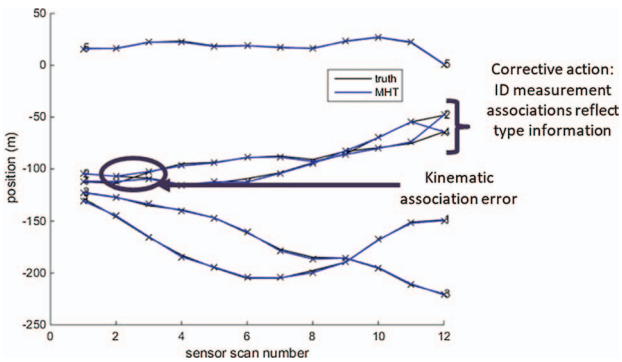

Fig. 11. MHT performs effective kinematic association but is ultimately myopic as it cannot exploit ID measurements that are in the distant future. ID information is part of the track state and, thus, corrective action is taken when prior association errors are detected.

The MI-GBT solution struggles with multiple target crossing in the absence of ID data, but it maintains ID information and is able to deal effectively with single-crossing events between ID reports. More importantly, the MI-GBT exploits ID data in performing associations with the second scan of ID-sensor data. Hence, ID measurement associations do not violate type information, and this is achieved without the corrective action that MHT exhibits. For the same computational load, deeper hypothesis reasoning is possible.

Of course, for sufficiently temporally distant ID measurements, MI-GBT will also require corrective action. Note also that, when focusing only on the MI-GBT tracking solution for targets of the same type, no track crossing occurs. This behavior is consistent with what we observe in the overall solution with all tracks in the GBT solution.

Due to the nonunique nature of target-type measurements, some incorrect ID measurement association decisions are performed by MHT and MI-GBT. Figs. 13 and 14 illustrate a scenario for which both MHT and MI-GBT incorrectly associate some ID measurements, when these are of the same type. We highlight a track that starts with a measurement on target 1, and ultimately associates with one on target 2. Note that both targets are of type A; hence, this error cannot be excluded.

In all cases, the final tracking solution for all paradigms (GBT, MHT, and MI-GBT) includes trajectory smoothing based on the forward–backward implementation of the Kalman smoother. This provides improved localization accuracy that is appropriate for

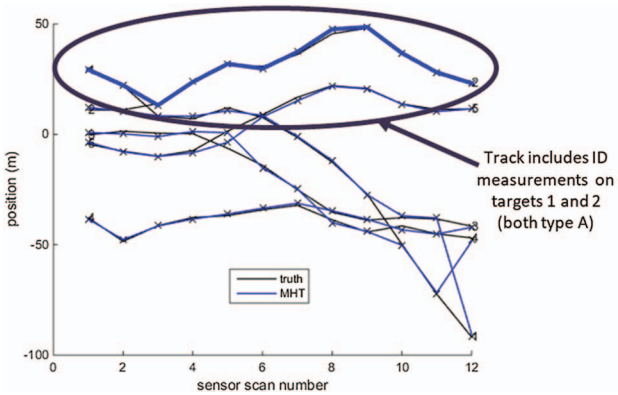

Fig. 13. ID measurement association error in MHT processing, due to common target type (targets 1 and 2).

forensic problems where maintaining small processing latency is not crucial. When it is, this postprocessing step may be omitted.

Figs. 15 and 16 provide performance as a function of the number of kinematic-sensor scans. Note that, with an increasing number of scans, correct measurement association decisions become harder, for all solution approaches. Indeed, the value of temporally distinct identity measurements is more limited than when temporally close identity measurements are available. There are at least two ways that this can be understood. First, knowledge of future location of a target has little bearing on data-association decisions, when the future time is in the distant future. Second, there are typically a larger number of ambiguous object-crossing events over a longer time horizon, providing many similarly scoring tracking solutions. The increasing difficulty of the MTT problem can be seen empirically in the fact that performance under *all* solution paradigms degrades as a function of the number of kinematic-sensor scans, as these lack target-type information.

We compare against a clairvoyant (ideal) algorithm for which measurement association is known *a priori*. We consider both the average track localization error and the fraction of correct data-association decisions.

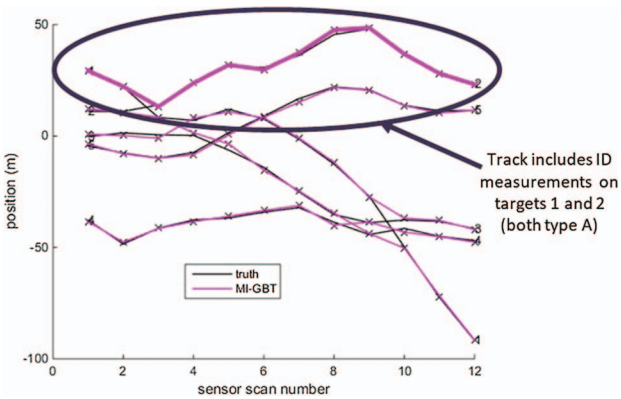

Fig. 14. ID measurement association error in MI-GBT processing, due to common target type (targets 1 and 2).

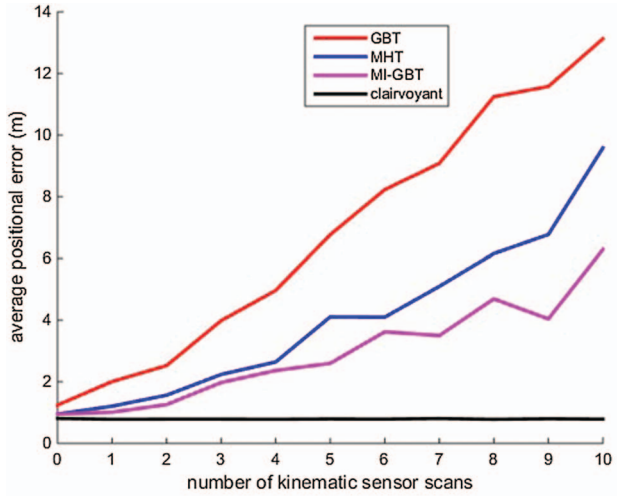

Fig. 15. Positional error of solutions: MI-GBT is best.

Results are based on 100 Monte Carlo realizations for each scenario duration.

These results on idealized scenarios provide confidence in the significant potential of MI-GBT processing for multi-INT surveillance, even when ID measurement associations are only partially constrained. The full MI-GBT solution accounts as well for birth/death phenomena and missed detections. A key enhancement relative to our early work is to relax the unity-flow constraint on the number of objects of each type. In so doing, it is crucial to express ID measurements as nodes in the multilayer graph topology.

V. CONCLUSION

This paper introduces an efficient, generalized GBT scheme for multi-INT track fusion that yields promising performance against an MHT baseline. Crucially, our scheme allows for object identity measurements via a multilayer graph approach, while exploiting kinematic path independence. As such, the approach may be

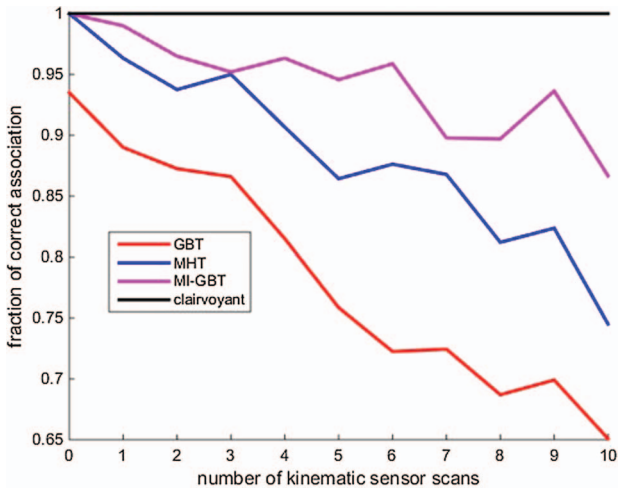

Fig. 16. Association accuracy of solutions: MI-GBT is best.

thought of as a hybrid GBT/MHT approach to data association. We allow for an arbitrary number of objects of each type and achieve scalability via sliding-window hypothesis resolution as is commonly performed in MHT.

It is important to emphasize that our proposed graph-based multisensor fusion algorithm is not fully general, in the sense that we do not directly handle an arbitrary number of kinematic and identity sources. Rather, we rely on upstream processing (see again Fig. 2) and assume a single (consolidated) kinematic source and a single (consolidated) identity source. Further, we do not consider fusion of temporally overlapping kinematic tracks (see again Fig. 1), and we assume singleton (single-measurement) identity tracks. Nonetheless, our work offers promising performance benefits over classical MHT technology in this restricted setting.

In ongoing work, we are investigating use of the MI-GBT on scenarios that exhibit *move–stop–move* target motion cycles and motion-sensitive kinematic sensors. Additionally, further analysis is needed to address slowly varying (nonstatic) feature states and noisy feature measurements.

REFERENCES

1. 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*. Hoboken, NJ: Wiley-IEEE, 2015.
2. D. Reid
"An algorithm for tracking multiple targets,"
IEEE Trans. Autom. Control, vol. 24, no. 6, pp. 843–854, Dec. 1979.
3. T. Kurien
"Issues in the design of practical multitarget tracking algorithms,"
in *Multitarget–Multisensor Tracking: Advanced Applications*, Y Bar-Shalom
Ed. Norwood, MA: Artech House, 1990.
4. C. Morefield
"Application of 0–1 integer programming to multitarget tracking problems,"
IEEE Trans. Autom. Control, vol. 22, no. 3, pp. 302–312, Jun. 1977.
5. S. Coraluppi, C. Carthel, M. Luettgen, and S. Lynch
"All-source track and identity fusion,"
in *Proc. MSS Nat. Symp. Sensor Data Fusion*, San Antonio, TX, USA, Jun. 2000.
6. A. Poore and N. Rijavec
"A Lagrangian relaxation algorithm for multidimensional assignment problems arising from multitarget tracking,"
SIAM J. Optim., vol. 3, no. 3, pp. 544–563, 1993.
7. P. Storms and F. Spieksma
"An LP-based algorithm for the data association problem in multitarget tracking,"
in *Proc. 3rd Int. Conf. Inf. Fusion*, Paris, France, Jul. 2000.
8. C.-Y. Chong, G. Castañón, N. Cooprider, S. Mori, R. Ravichandran, and R. Macior
"Efficient multiple hypothesis tracking by track segment graph,"
in *Proc. 12th Int. Conf. Inf. Fusion*, Seattle, WA, USA, Jul. 2009.
9. S. Coraluppi, C. Carthel, W. Kreamer, G. Titi, and J. Vasquez
"Advanced multi-stage multiple-hypothesis GMTI tracking,"
in *Proc. MSS Nat. Symp. Sensor Data Fusion*, Gaithersburg, MD, USA, Jun. 2016.
10. S. Coraluppi, C. Carthel, B. Zimmerman, T. Allen, J. Douglas, and J. Muka
"Multi-stage MHT with airborne and ground sensors,"
in *Proc. IEEE Aerosp. Conf.*, Big Sky, MT, USA, Mar. 2018.
11. 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 Trans. Aerosp. Electron. Syst., vol. 50, no. 3, pp. 2379–2388, Jul. 2014.
12. S. Coraluppi and C. Carthel
"Multiple-hypothesis tracking for targets producing multiple measurements,"
IEEE Trans. Aerosp. Electron. Syst., vol. 54, no. 3, pp. 1485–1498, Jun. 2018.
13. T. Sathyan, T.-J. Chin, S. Arulampalam, and D. Suter
"A multiple hypothesis tracker for multitarget tracking with multiple simultaneous measurements,"
IEEE J. Sel. Topics Signal Process., vol. 7, no. 3, pp. 448–460, Jun. 2013.
14. J. Wolf, A. Viterbi, and G. Dixon
"Finding the best set of K paths through a trellis with application to multitarget tracking,"
IEEE Trans. Aerosp. Electron. Syst., vol. 25, no. 2, pp. 287–295, Mar. 1989.
15. G. Pulford and B. La Scala
"Multihypothesis Viterbi data association: Algorithm development and assessment,"
IEEE Trans. Aerosp. Electron. Syst., vol. 46, no. 2, pp. 583–609, Apr. 2010.
16. L. van der Merwe and J. de Villiers
"A comparative investigation into Viterbi based and multiple hypothesis based track stitching,"
IET Radar, Sonar Navigation, vol. 10, no. 9, pp. 1575–1582, Dec. 2016.
17. D. Castañón
"Efficient algorithms for finding the K best paths through a trellis,"
IEEE Trans. Aerosp. Electron. Syst., vol. 26, no. 2, pp. 405–410, Mar. 1990.
18. G. Battistelli, L. Chisci, F. Papi, A. Benavoli, A. Farina, and A. Graziano
"Optimal flow models for multiscan data association,"
IEEE Trans. Aerosp. Electron. Syst., vol. 47, no. 4, pp. 2405–2422, Oct. 2011.
19. G. Castañón and L. Finn
"Multi-target tracklet stitching through network flows,"
in *Proc. IEEE Aerosp. Conf.*, Big Sky, MT, USA, Mar. 2011.
20. H. Shirit, J. Berclaz, F. Fleuret, and P. Fua
"Multi-commodity network flow for tracking multiple people,"
IEEE Trans. Pattern Anal. Mach. Intell., vol. 36, no. 8, pp. 1614–1627, Aug. 2014.
21. S. Coraluppi, C. Carthel, G. Castañón, and A. Willsky
"Multiple-hypothesis and graph-based tracking for kinematic and identity fusion,"
in *Proc. IEEE Aerosp. Conf.*, Big Sky, MT, USA, Mar. 2018.
22. S. Coraluppi, C. Carthel, W. Kreamer, and A. Willsky
"New graph-based and MCMC approaches to multi-INT surveillance,"
in *Proc. 19th Int. Conf. Inf. Fusion*, Heidelberg, Germany, Jul. 2016.
23. S. Oh, S. Russell, and S. Sastry

"Markov chain Monte Carlo data association for multi-target tracking,"
IEEE Trans. Aerosp. Electron. Syst., vol. 54, no. 3, pp. 481–497, Mar. 2009.
24. L. Finn and P. Kingston
"Unifying multi-hypothesis and graph-based tracking with approximate track automata,"
in *Proc. IEEE Aerosp. Conf.*, Big Sky, MT, USA, Mar. 2019.
25. T. Tran, J. Foley, G. Godfrey, and J. Ferry
"Graph-based multi-INT tracklet stitching,"
in *Proc. MSS Nat. Symp. Sensor Data Fusion*, Springfield, VA, USA, Oct. 2018.
26. S. Coraluppi, C. Carthel, and A. Willsky
"Graph-based tracking with uncertain ID measurement associations,"
Proc. 22nd Int. Conf. Inf. Fusion, Ottawa, Canada, Jul. 2019.
27. Y. Bar-Shalom, P. Willett, and X. Tian
Tracking and Data Fusion: A Handbook of Algorithms.
Storrs, CT: YBS Publishing, 2011.
28. V. Poor
An Introduction to Signal Detection and Estimation. Berlin, Germany: Springer, 1988.
29. R. Mahler
Statistical Multisource--Multitarget Information Fusion.
Norwood, MA: Artech House, 2007.
30. J. Williams
"Marginal multi-Bernoulli filters: RFS derivation of MHT, JIPDA, and association-based member,"
IEEE Trans. Aerosp. Electron. Syst., vol. 51, no. 3, pp. 1664–1687, Jul. 2015.
31. S. Coraluppi and C. Carthel
"Track management in multiple-hypothesis tracking,"
in *Proc. 10th IEEE Sensor Array Multichannel Signal Process. Workshop*, Sheffield, U.K., Jul. 2018.
32. A. Papoulis
Probability, Random Variables, and Stochastic Processes.
New York, NY: McGraw-Hill, 1991.
33. S. Coraluppi and C. Carthel
"Stability and stationarity in target kinematic modeling,"
in *Proc. IEEE Aerosp. Conf.*, Big Sky, MT, USA, Mar. 2012.
34. C.-Y. Chong
"Graph approaches for data association,"
in *Proc. 15th Int. Conf. Inf. Fusion*, Singapore, Jul. 2012.

Stefano P. Coraluppi received the B.S. degree in electrical engineering and mathematics from Carnegie Mellon University, Pittsburgh, PA, USA, in 1990, and the M.S. and Ph.D. degrees in electrical engineering from the University of Maryland, College Park, MD, USA, in 1992 and 1997, respectively. He is currently a Chief Scientist with Systems and Technology Research, Woburn, MA, USA. He has held research positions with ALPHATECH, Inc. (1997–2002), the NATO Undersea Research Centre (2002–2010), Compunetix, Inc. (2010–2014), and STR (since 2014). His primary research interests are multitarget tracking and multisensor data fusion. He serves on the Board of Governors of the IEEE Aerospace and Electronic Systems Society (AESS) and the Board of Directors of the International Society of Information Fusion (ISIF). He is Associate Editor-in-Chief for both the IEEE *Transactions on Aerospace and Electronic Systems* and the ISIF *Journal of Advances in Information Fusion*. He is a Senior Member of IEEE.

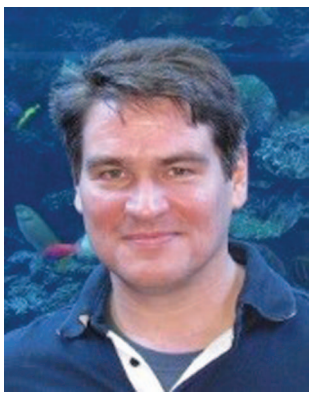

Craig A. Carthel received the B.S. degree in physics and mathematics, and the M.S. and Ph.D. degrees in mathematics from the University of Houston, Houston, TX, USA, in 1988, 1992, and 1995, respectively, where he did research in numerical analysis and optimization theory. He is currently a Lead Scientist with Systems and Technology Research, Woburn, MA, USA. From 1995 to 1997, he was with the Institute for Industrial Mathematics, Johannes Kepler University, Linz, Austria, where he worked on parameter identification and inverse problems. From 1998 to 2002, he was a Senior Mathematician with AlphaTech, Inc., Burlington, MA, USA, where he worked on image processing, multisensor data fusion, and ground target tracking. From 2002 to 2010, he was a Senior Scientist with the NATO Undersea Research Centre, La Spezia, Italy, where he worked on military operations research, simulation, optimization, and data fusion problems associated with maritime environments. From 2010 to 2014, he was a Principal Scientist with Compunetix, Inc., Monroeville, PA, USA. In 2006, he was the Technical Program Chair for FUSION 2006.

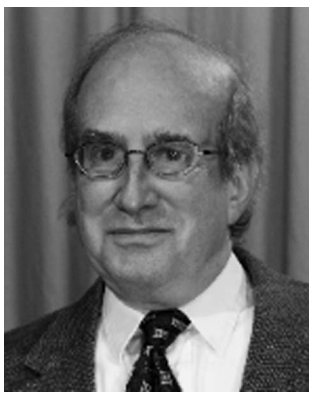

Alan S. Willsky received the S.B. and Ph.D. degrees from the Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA, USA, in 1969 and 1973, respectively. In 1973, he joined the faculty of Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, continued full time through June 2014, and continues now as Edwin Sibley Webster Professor (retired). He served in leadership roles and from 2009 to 2014 as Director of the Laboratory for Information and Decision Systems. He has held visiting positions with Imperial College London (1977), Universite de Paris-Sud (1980–1981), and INRIA (1988). He was a founder, member of the Board of Directors, and Chief Scientific Consultant with AlphaTech, Inc. (where his work on statistical signal processing found applications in a variety of domains including multiobject tracking) and served for 4 years as a member of the U.S. Air Force Scientific Advisory Board. He has authored more than 220 journal papers and 375 conference papers, as well as 2 books, including the widely used undergraduate text *Signals and Systems*. He has received numerous awards, including the 1975 American Automatic Control Council Donald P. Eckman Award, the 1980 IEEE Browder J. Thompson Memorial Award and 1979 ASCE Alfred Noble Prize (both for a *Proceedings of the IEEE* paper on control-signal processing synergies), the 2004 IEEE Donald G. Fink Prize Paper Award (for a widely cited paper in the *Proceedings of the IEEE* that provided a survey and tutorial on multiscale statistical models for signal and image processing), a number of other best paper awards, and an honorary doctorate from Université de Rennes. He has given numerous plenary and keynote lectures at major meetings and received the 2009 Technical Achievement Award from the IEEE Signal Processing Society as well as the 2014 IEEE Signal Processing Society Award. In 2010, he was elected to the National Academy of Engineering and he is the recipient of the 2019 IEEE Jack S. Kilby Signal Processing Medal. His early work on methods for failure detection in dynamic systems is still widely cited and used in practice, and his research on statistical and multiresolution and curve-evolution methods has found application in diverse fields, including target tracking, computer vision, biomedical image analysis, oceanographic remote sensing, and groundwater hydrology. He is a Life Fellow of IEEE.