

Track Coalescence and Repulsion in Multitarget Tracking: An Analysis of MHT, JPDA, and Belief Propagation Methods

THOMAS KROPFREITER ¹ (Member, IEEE), FLORIAN MEYER ¹ (Member, IEEE),
DAVID F. CROUSE ² (Senior Member, IEEE), STEFANO CORALUPPI ³ (Fellow, IEEE),
FRANZ HLAWATSCH ⁴ (Life Fellow, IEEE), AND PETER WILLETT ⁵ (Fellow, IEEE)

¹ Scripps Institution of Oceanography and Department of Electrical and Computer Engineering, University of California San Diego, La Jolla, CA 92037 USA

² Naval Research Laboratory, Washington, DC 20375 USA

³ Systems & Technology Research, Woburn, MA 01801 USA

⁴ Institute of Telecommunications, TU Wien, 1040 Vienna, Austria

⁵ Department of Electrical and Computer Engineering, University of Connecticut, Storrs, CT 06269 USA

CORRESPONDING AUTHOR: THOMAS KROPFREITER (e-mail: tkropfreiter@ucsd.edu).

An earlier version of this paper was presented in part at FUSION-2021, Sun City, South Africa, Nov. 2021 [DOI: 10.23919/FUSION49465.2021.9626958].

This work was supported in part by the University of California San Diego, in part by the Office of Naval Research under Grant N00014-23-1-2284, and in part by the Austrian Science Fund under Grant P32055-N31 and Grant J4726-N.

ABSTRACT Joint probabilistic data association (JPDA) filter methods and multiple hypothesis tracking (MHT) methods are widely used for multitarget tracking (MTT). However, they are known to exhibit undesirable behavior in tracking scenarios with targets in close proximity: JPDA filter methods suffer from the track coalescence effect, i.e., the estimated tracks of targets in close proximity tend to merge and can become indistinguishable, while MHT methods suffer from an opposite effect known as track repulsion, i.e., the estimated tracks of targets in close proximity tend to repel each other in the sense that their separation is larger than the actual distance between the targets. In this paper, we review the JPDA filter and MHT methods and discuss the track coalescence and track repulsion effects. We also consider a more recent methodology for MTT that is based on the belief propagation (BP) algorithm. We argue that BP-based MTT does not exhibit track repulsion because it is not based on maximum a posteriori estimation, and that it exhibits significantly reduced track coalescence because certain properties of the BP messages related to data association encourage separation of target state estimates. Our theoretical arguments are confirmed by numerical results for four representative simulation scenarios.

INDEX TERMS Multitarget tracking, joint probabilistic data association, JPDA, multiple hypothesis tracking, MHT, belief propagation, track coalescence, track repulsion.

I. INTRODUCTION

A. BACKGROUND AND MOTIVATION

The goal of multitarget tracking (MTT) is to recursively estimate the time-dependent number and states of multiple objects (“targets”) based on noisy sensor measurements [1], [2], [3], [4], [5], [6]. This enables situational awareness in a wide variety of applications, including aerospace and maritime surveillance [2], [7], [8], [9], [10], [11], [12] and robotics [13], [14], [15]. However, MTT is a challenging task due to clutter, missed detections, and measurement-origin uncertainty [1], [3], [4], [5], [6]. MTT methods are

typically placed in the framework of Bayesian inference and can be broadly classified as vector-type and set-type methods. Vector-type methods [1], [2], [7], [8], [9], [16], [17], [18], [19], [20], [21], [22], [23], [24], [25], [26], [27], [28] model the multitarget state by a random vector, whereas set-type methods [4], [29], [30], [31], [32], [33], [34], [35], [36], [37] model it by a random finite set.

Two classical and popular instances of vector-type MTT methods are the joint probabilistic data association (JPDA) filter [1], [7], [9] and multiple hypothesis tracking (MHT) methods [2], [16], [21], [22], [23], [24], [25]. The JPDA filter

is a single-scan method that seeks to calculate minimum mean square error (MMSE) state estimates for a known number of targets at each time step. An unknown number of targets can be addressed using heuristics [1] or a probabilistic model for target existence as used for the development of the joint integrated probabilistic data association (JIPDA) filter [17]. Data association (DA) hypotheses are incorporated in a soft (probabilistic) manner by calculating weighted sums over all possible associations at each time step. This is equivalent to “marginalizing out” the unknown DA vector in the posterior probability density function (pdf) of the target states, which can also be seen as an exhaustive averaging over DA vectors. In contrast, MHT methods are multiscan methods that achieve hard DA by performing approximate maximum a posteriori (MAP) detection of an entire sequence of DA vectors using optimization techniques such as the Auction algorithm or multiple frame assignment [38]. Subsequently, the target states are estimated individually, e.g., using Kalman smoothing. Both MHT and JPDA methods are capable of maintaining track continuity, i.e., target identification over consecutive time steps. Besides the aspect of performing exhaustive averaging as in JPDA filter methods or approximate maximization as in MHT methods, an additional aspect is the representation of a single potential target or “track”. Typical representations are by a single Gaussian [2], [7], a Gaussian mixture [39], [40], or a set of particles [41], [42]. We note that certain set-type filters, e.g., [32] and [37], also use exhaustive averaging or approximate maximization as their DA strategy.

Experimental evidence has demonstrated undesirable behavior of JPDA and MHT type methods in tracking scenarios with targets in close proximity [43], [44]. More specifically, JPDA filter methods suffer from the track coalescence effect [43], which means that the estimated tracks of closely spaced targets that move in parallel tend to come together and merge, thereby becoming indistinguishable. In MHT methods, on the other hand, an opposite effect called track repulsion can be observed [44]: the estimated tracks of closely spaced targets that move in parallel repel each other in the sense that their separation is larger than the actual distance between the targets. Methodological modifications intended to mitigate these detrimental effects were proposed in [45], [46], [47], [48], [49], [50]. In particular, track coalescence can be mitigated by using a specific hypothesis pruning strategy [45], [46] or filter designs based on the optimal sub-pattern assignment (OSPA) metric [48], [49] or on variational methods [50], and track repulsion by choosing a DA hypothesis from a MAP equivalence class rather than the single MAP solution [47]. However, these modifications are unsuitable for large-scale tracking scenarios with an unknown and time-varying number of targets.

Among the set-type MTT methods, some avoid DA at the cost of using a suboptimal postprocessing step for track formation and, possibly, exhibiting a reduced target detection and tracking performance [29], [30], [31], [33]. Others rely on DA strategies similar to those used by the JPDA filter and

hence also suffer from track coalescence [32], [51]. Inspired by the concept of set-type MTT, track coalescence has also been investigated in the context of permutation-invariant estimation [48], [52], [53].

A recently developed methodology for vector-type MTT describes filtering and DA by a factor graph that provides a blueprint for applying belief propagation (BP) [6], [12], [15], [42], [54], [55]. BP, also known as the *sum-product algorithm* [56], [57], is based on a factor graph [56], [57] and aims at computing the marginal posterior pdfs or probability mass functions (pmfs) in an efficient way. For each node of the factor graph, certain messages are calculated, each of which is passed to some of the adjacent nodes. If the factor graph is a tree, i.e., without loops, then the solutions provided by BP—referred to as beliefs—are equal to the marginal posterior pdfs or pmfs. For factor graphs with loops, BP has to be performed iteratively, and the beliefs are only approximations of the respective marginal posterior pdfs or pmfs; these approximations have been observed to be very accurate in many applications [56], [57]. BP can be applied for the inference of discrete or continuous random variables. If the belief represents a marginal pmf, its values can be interpreted as probabilities. In the context of probabilistic DA in MTT, BP is guaranteed to converge [54].

By performing efficient approximate marginalization operations within the overall MMSE estimation framework of the JPDA filter, the BP algorithm provides scalable solutions to both probabilistic DA—accounting for measurement-origin uncertainty, clutter, and missed detections—and nonlinear state estimation for randomly appearing and disappearing targets. BP-based MTT methods exhibit excellent scalability in the number of targets and in the number of measurements because the BP algorithm systematically exploits conditional independence properties of the underlying statistical model for a reduction of computational complexity. BP-based methods are also capable of maintaining track continuity. Since the nodes of the graph corresponding to the target states are ordered, BP-based methods are quite different from the set-type approach that aims at performing permutation-invariant or unordered estimation. BP methods rely on efficient approximations to exhaustive averaging and are thus related to J(I)PDA filters. However, contrary to J(I)PDA filters, BP methods employ an improved birth model that, together with the efficient DA solution, allows the detection and tracking of a potentially very large number of targets even if the targets are in close vicinity in space and time.

B. CONTRIBUTIONS AND PAPER ORGANIZATION

In this paper, we review the track coalescence and track repulsion effects and investigate them in the context of the JPDA filter and MHT methods as well as the BP-based MTT methods. Our analysis suggests that BP-based MTT methods do not suffer from the track repulsion effect, while the track coalescence effect is significantly smaller than in JPDA filter methods. These findings are confirmed by

simulation results for four different MTT scenarios, which demonstrate a significant reduction of track coalescence and the absence of track repulsion in BP-based methods relative to JPDA filter methods and MHT methods, respectively. In contrast to permutation-invariant estimation [48], [52], [53], where track coalescence is avoided by an alternative formulation of the estimation problem, in BP-based MTT, track coalescence is reduced because certain properties of the DA-related BP messages encourage separation of target state estimates.

The main contributions of the paper are as follows:

- We discuss and analyze the track coalescence effect and the track repulsion effect in the context of the JPDA filter and MHT methods, respectively.
- We demonstrate numerically that BP-based MTT does not suffer from the track repulsion effect, while the track coalescence effect is much smaller than in the JPDA filter.
- We provide evidence that the reduced coalescence in BP-based MTT is due to the nature of the DA solution and not due to the choice of track representation.

The remainder of this paper is organized as follows. The basic notation is described in the next subsection. In Section II, we present the MTT system model used in this paper as well as a statistical formulation. Section III describes the resulting joint posterior distribution and the corresponding factor graph. The track repulsion effect in MHT methods and the track coalescence effect in JPDA filter methods are considered in Sections IV and V, respectively. In Section VI, we describe the BP approach to MTT. The reduction of the track coalescence effect achieved by particle-based BP methods is analyzed in Section VII. Finally, simulation results are presented in Section VIII. We note that this paper advances beyond our earlier conference publication [58] in that it contains an improved introduction to MHT and JPDA filter methods; it presents a simple illustration of the track repulsion effect due to MHT-like hard DA; it provides a detailed analysis of the track coalescence effect and an explanation why BP methods are less susceptible to it; and it presents a more extensive simulation analysis.

C. NOTATION

Random variables are displayed by sans serif, upright print (e.g., \mathbf{x}) and their realizations by serif, italic print (e.g., x). Vectors are denoted by bold lowercase letters (e.g., \mathbf{x} or \mathbf{x}) and matrices by bold uppercase letters (e.g., \mathbf{A}). Furthermore, \mathbf{x}^T denotes the transpose of vector \mathbf{x} , α indicates equality up to a normalization factor, and $f(\mathbf{x})$ denotes the pdf and $p(\mathbf{x})$ the pmf of the, respectively, continuous and discrete random vector \mathbf{x} . We also use $f(\mathbf{x})$ to denote the hybrid pdf/pmf of a mixed continuous-discrete random vector \mathbf{x} . Finally, $1(a)$ is defined to be 1 if $a = 0$ and 0 otherwise, and $\delta(x)$ denotes the Dirac delta function.

II. SYSTEM MODEL

We consider the MTT system model from [1], [3], [6], [24]. This model and the corresponding statistical formulation will be briefly described for completeness.

A. POTENTIAL TARGET STATES AND STATE EVOLUTION

The number of targets is time-varying and unknown, which is accounted for by using the concept of potential targets (PTs) [6], [42]. At discrete time $k = 0, 1, \dots$, the number of PTs J_k is the maximum possible number of targets, which is determined as described in Section II-B. The state $[\mathbf{x}_k^{(j)T} \ r_k^{(j)}]^T$ of PT $j \in \{1, \dots, J_k\}$ at time k consists of a kinematic state $\mathbf{x}_k^{(j)} \in \mathbb{R}^n$, which comprises the PT's position and possibly further motion-related parameters, and a binary existence indicator variable $r_k^{(j)} \in \{0, 1\}$ that is 1 if PT j exists at time k and 0 otherwise. This existence indicator variable was first introduced in the context of the JIPDA filter [17]. The kinematic state $\mathbf{x}_k^{(j)}$ of a nonexisting PT—i.e., of a PT with $r_k^{(j)} = 0$ —is undefined; its distribution will therefore be described by an arbitrary “dummy pdf” $f_D(\mathbf{x}_k^{(j)})$. Accordingly, all pdfs of PT states, $f(\mathbf{x}_k^{(j)}, r_k^{(j)})$, have the property that $f(\mathbf{x}_k^{(j)}, r_k^{(j)} = 0) = f_k^{(j)} f_D(\mathbf{x}_k^{(j)})$, where $f_k^{(j)} \in [0, 1]$ is the probability of nonexistence. Only targets that exist can generate measurements.

For each PT state $[\mathbf{x}_{k-1}^{(j)T} \ r_{k-1}^{(j)}]^T$, $j \in \{1, \dots, J_{k-1}\}$ at time $k - 1$, there is one “legacy” PT state $[\mathbf{x}_k^{(j)T} \ r_k^{(j)}]^T$ at time k . The motion and potential disappearance of PT j are modeled statistically by the single-target state-transition pdf/pmf $f(\mathbf{x}_k^{(j)}, r_k^{(j)} | \mathbf{x}_{k-1}^{(j)}, r_{k-1}^{(j)})$, which involves the kinematic state-transition pdf $f(\mathbf{x}_k^{(j)} | \mathbf{x}_{k-1}^{(j)})$ and the probability of target survival p_s as described in, e.g., [6, Section VIII-C]. All PT states evolve independently, and they are independent at time $k = 0$. In the absence of any information about the number of PTs at time $k = 0$, we set $J_0 = 0$.

B. MEASUREMENTS, NEW PTS, AND DATA ASSOCIATION

At each time (or “scan”) $k \geq 1$, a sensor produces M_k measurements $\mathbf{z}_k^{(m)}$, $m \in \{1, \dots, M_k\}$. Here, M_k is modeled as a random variable whose value (realization) M_k is known once the measurements have been observed. We define the joint vector of all measurements at time k as $\mathbf{z}_k \triangleq [\mathbf{z}_k^{(1)T} \dots \mathbf{z}_k^{(M_k)T}]^T$. Each measurement $\mathbf{z}_k^{(m)}$ has one of three possible origins: (i) a new PT representing a target that generates a measurement for the first time, or (ii) a legacy PT representing a target that generated at least one measurement before, or (iii) clutter.

The birth of new targets is modeled by a Poisson point process with mean parameter μ_b (which is the mean number of newborn targets) and spatial pdf $f_b(\bar{\mathbf{x}}_k)$. To account for the birth of new targets, at time k , M_k new PT states $[\bar{\mathbf{x}}_k^{(m)T} \ \bar{r}_k^{(m)}]^T$, $m \in \{1, \dots, M_k\}$ are introduced, where $\bar{r}_k^{(m)} = 1$ means that measurement $\mathbf{z}_k^{(m)}$ has origin (i) and $\bar{r}_k^{(m)} = 0$ means that it has origin (ii) or (iii).

We define the joint vector of all kinematic states at time k as $\mathbf{x}_k \triangleq [\mathbf{x}_k^{(1)} \dots \mathbf{x}_k^{(J_k)}]^T = [\mathbf{x}_k^{(1)T} \dots \mathbf{x}_k^{(J_{k-1})T} \ \bar{\mathbf{x}}_k^{(1)T} \dots \bar{\mathbf{x}}_k^{(M_k)T}]^T$,

and the joint vector of all existence indicators at time k as $\mathbf{r}_k \triangleq [\mathbf{r}_k^{(1)} \dots \mathbf{r}_k^{(J_k)}]^T = [\mathbf{r}_k^{(1)} \dots \mathbf{r}_k^{(J_{k-1})} \mathbf{r}_k^{(1)} \dots \mathbf{r}_k^{(M_k)}]^T$. It follows that once the measurements have been observed and, thus, M_k is known, the total number of PTs—both legacy PTs and new PTs—is $J_k = J_{k-1} + M_k$.

The target represented by PT j is detected by the sensor, in the sense that it generates a measurement $\mathbf{z}_k^{(m)}$, with probability p_d . The dependence of a target-generated measurement $\mathbf{z}_k^{(m)}$ on the kinematic state $\mathbf{x}_k^{(j)}$ of the corresponding detected PT is modeled statistically by the conditional pdf $f(\mathbf{z}_k^{(m)}|\mathbf{x}_k^{(j)})$, which can be derived from the measurement model of the sensor. Clutter measurements are modeled by a Poisson point process with mean parameter μ_c (which is the mean number of clutter measurements) and spatial pdf $f_c(\mathbf{z}_k^{(m)})$. Note that if $f_c(\mathbf{z}_k^{(m)})$ is uniform, then the Poisson point process has a constant spatial density.

The origin of each measurement is unknown, i.e., it is not known if the measurement is generated by a PT, and by which PT, or if it is clutter. We use the conventional DA assumption, which postulates that a target can generate at most one measurement, and a measurement can be generated by at most one target [1], [4], [6]. The association between the M_k measurements and the J_{k-1} legacy PTs can be modeled by the “target-oriented” DA vector $\mathbf{a}_k = [\mathbf{a}_k^{(1)} \dots \mathbf{a}_k^{(J_{k-1})}]^T$ whose j th entry $\mathbf{a}_k^{(j)}$ is $m \in \{1, \dots, M_k\}$ if legacy PT j generated measurement m and zero if it did not generate any measurement [1], [6]. Since only targets that exist can generate measurements, $\mathbf{a}_k^{(j)} \in \{1, \dots, M_k\}$ implies that $\mathbf{r}_k^{(j)} = 1$. However, since there can be a missed detection, we can have $\mathbf{a}_k^{(j)} = 0$ for $\mathbf{r}_k^{(j)} = 1$. Note that $\mathbf{a}_k^{(j)} = m \in \{1, \dots, M_k\}$ implies that measurement $\mathbf{z}_k^{(m)}$ has origin (ii). A DA vector \mathbf{a}_k is considered “valid” if it satisfies the DA assumption mentioned above. Due to the DA assumption, a DA vector \mathbf{a}_k is valid if and only if at most one entry is equal to any specific $m \in \{1, \dots, M_k\}$. Also note that due to the DA assumption, $\mathbf{a}_k^{(j)} = m$ implies $\bar{\mathbf{r}}_k^{(m)} = 0$, because if measurement m was generated by a legacy PT, the corresponding new PT cannot correspond to an actual target.

We also introduce the “measurement-oriented” DA vector $\mathbf{b}_k = [\mathbf{b}_k^{(1)} \dots \mathbf{b}_k^{(M_k)}]^T$ whose m th entry $\mathbf{b}_k^{(m)}$ is $j \in \{1, \dots, J_{k-1}\}$ if measurement m was generated by legacy PT j and zero if it was generated by a new PT or by clutter. Together, \mathbf{a}_k and \mathbf{b}_k constitute a redundant DA representation because for any given valid \mathbf{a}_k , there is exactly one valid \mathbf{b}_k and for any given valid \mathbf{b}_k , there is exactly one valid \mathbf{a}_k . However, this redundant representation makes it possible to develop scalable MTT methods that exploit the structure of the DA problem for a substantial reduction of computational complexity [6], [54] (see Section VI).

III. JOINT POSTERIOR DISTRIBUTION AND FACTOR GRAPH

Let $\mathbf{x}_{0:k}$ denote the vector obtained by stacking the vectors \mathbf{x}_0 through \mathbf{x}_k , and similarly for $\mathbf{r}_{0:k}$, $\mathbf{a}_{1:k}$, $\mathbf{b}_{1:k}$, and $\mathbf{z}_{1:k}$ as well

as their random counterparts. Using common assumptions, it has been shown in [6], [32], [42], [54] that the joint posterior pdf/pmf of $\mathbf{x}_{0:k}$, $\mathbf{r}_{0:k}$, $\mathbf{a}_{1:k}$, and $\mathbf{b}_{1:k}$ conditioned on the observed (and thus fixed) measurements $\mathbf{z}_{1:k}$ is given by

$$f(\mathbf{x}_{0:k}, \mathbf{r}_{0:k}, \mathbf{a}_{1:k}, \mathbf{b}_{1:k} | \mathbf{z}_{1:k}) \propto \left(\prod_{j=1}^{J_0} f(\mathbf{x}_0^{(j)}, \mathbf{r}_0^{(j)}) \right) \prod_{k'=1}^k g(\mathbf{x}_{k'}, \mathbf{r}_{k'}, \mathbf{a}_{k'}, \mathbf{b}_{k'}, \mathbf{x}_{k'-1}, \mathbf{r}_{k'-1}; \mathbf{z}_{k'}), \quad (1)$$

with

$$\begin{aligned} g(\mathbf{x}_k, \mathbf{r}_k, \mathbf{a}_k, \mathbf{b}_k, \mathbf{x}_{k-1}, \mathbf{r}_{k-1}; \mathbf{z}_k) &\triangleq \left(\prod_{j'=1}^{J_{k-1}} f(\mathbf{x}_k^{(j')}, \mathbf{r}_k^{(j')} | \mathbf{x}_{k-1}^{(j')}, \mathbf{r}_{k-1}^{(j')}) \right) \\ &\times \left(\prod_{j=1}^{J_{k-1}} q(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)}, \mathbf{a}_k^{(j)}; \mathbf{z}_k) \prod_{m'=1}^{M_k} \Psi_{j,m'}(\mathbf{a}_k^{(j)}, \mathbf{b}_k^{(m')}) \right) \\ &\times \prod_{m=1}^{M_k} v(\bar{\mathbf{x}}_k^{(m)}, \bar{\mathbf{r}}_k^{(m)}, \mathbf{b}_k^{(m)}; \mathbf{z}_k^{(m)}). \end{aligned} \quad (2)$$

Here, the factor $q(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)}, \mathbf{a}_k^{(j)}; \mathbf{z}_k)$ is given by

$$\begin{aligned} q(\mathbf{x}_k^{(j)}, 1, \mathbf{a}_k^{(j)}; \mathbf{z}_k) &= \begin{cases} \frac{p_d}{\mu_c f_c(\mathbf{z}_k^{(m)})} f(\mathbf{z}_k^{(m)} | \mathbf{x}_k^{(j)}), & \mathbf{a}_k^{(j)} = m \in \{1, \dots, M_k\} \\ 1 - p_d, & \mathbf{a}_k^{(j)} = 0 \end{cases} \end{aligned} \quad (3)$$

and $q(\mathbf{x}_k^{(j)}, 0, \mathbf{a}_k^{(j)}; \mathbf{z}_k) = 1(a_k^{(j)} = 0)$, and the factor $v(\bar{\mathbf{x}}_k^{(m)}, \bar{\mathbf{r}}_k^{(m)}, \mathbf{b}_k^{(m)}; \mathbf{z}_k^{(m)})$ is given by

$$\begin{aligned} v(\bar{\mathbf{x}}_k^{(m)}, 1, \mathbf{b}_k^{(m)}; \mathbf{z}_k^{(m)}) &= \begin{cases} 0, & \mathbf{b}_k^{(m)} \in \{1, \dots, J_{k-1}\} \\ \frac{p_d \mu_b f_b(\bar{\mathbf{x}}_k^{(m)})}{\mu_c f_c(\mathbf{z}_k^{(m)})} f(\mathbf{z}_k^{(m)} | \bar{\mathbf{x}}_k^{(m)}), & \mathbf{b}_k^{(m)} = 0 \end{cases} \end{aligned} \quad (4)$$

and $v(\bar{\mathbf{x}}_k^{(m)}, 0, \mathbf{b}_k^{(m)}; \mathbf{z}_k^{(m)}) = f_D(\bar{\mathbf{x}}_k^{(m)})$. Furthermore, $\Psi_{j,m}(\mathbf{a}_k^{(j)}, \mathbf{b}_k^{(m)})$ is a binary indicator function that checks the consistency of the target-oriented DA variable $\mathbf{a}_k^{(j)}$ and the measurement-oriented DA variable $\mathbf{b}_k^{(m)}$: it is zero if $\mathbf{a}_k^{(j)} = m$, $\mathbf{b}_k^{(m)} \neq j$ or $\mathbf{b}_k^{(m)} = j$, $\mathbf{a}_k^{(j)} \neq m$ and one otherwise (see [6], [54] for details). This consistency of the DA variables $\mathbf{a}_k^{(j)}$ and $\mathbf{b}_k^{(m)}$ guarantees that the DA assumption is not violated. The \propto sign in (1) has to be interpreted in the sense that the right-hand side, when multiplied by a normalization constant that does not depend on $\mathbf{x}_{0:k}$, $\mathbf{r}_{0:k}$, $\mathbf{a}_{1:k}$, or $\mathbf{b}_{1:k}$, is equal to the joint posterior pdf/pmf $f(\mathbf{x}_{0:k}, \mathbf{r}_{0:k}, \mathbf{a}_{1:k}, \mathbf{b}_{1:k} | \mathbf{z}_{1:k})$. We note that the factor $g(\mathbf{x}_k, \mathbf{r}_k, \mathbf{a}_k, \mathbf{b}_k, \mathbf{x}_{k-1}, \mathbf{r}_{k-1}; \mathbf{z}_k)$ in (1) groups all factors corresponding to a single time step; this representation will support future discussions in Section VI.

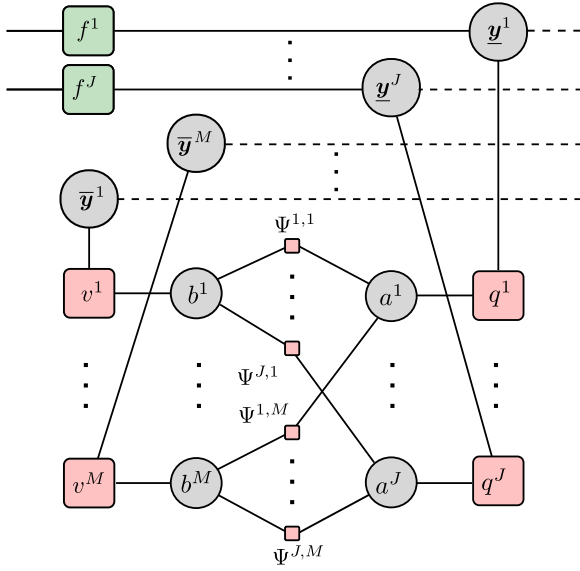


FIGURE 1. Factor graph representing the joint posterior pdf/pmf $f(\mathbf{x}_{0:k}, \mathbf{r}_{0:k}, \mathbf{a}_{1:k}, \mathbf{b}_{1:k} | \mathbf{z}_{1:k})$ in (1), (2) for one time step k . The time index k is omitted, and the following short notations are used: $M \triangleq M_k, J \triangleq J_{k-1}, \mathbf{y}^j \triangleq [\mathbf{x}_k^{(j)T} \mathbf{r}_k^{(j)T}]^T, \bar{\mathbf{y}}^m \triangleq [\bar{\mathbf{x}}_k^{(m)T} \bar{\mathbf{r}}_k^{(m)T}]^T, \mathbf{a}^j \triangleq \mathbf{a}_k^{(j)}, \mathbf{b}^m \triangleq \mathbf{b}_k^{(m)}, f^j \triangleq f(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)} | \mathbf{x}_{k-1}^{(j)}, \mathbf{r}_{k-1}^{(j)}), q^j \triangleq q(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)}, \mathbf{a}_k^{(j)}, \mathbf{z}_k), v^m \triangleq v(\bar{\mathbf{x}}_k^{(m)}, \bar{\mathbf{r}}_k^{(m)}, \mathbf{b}_k^{(m)}, \mathbf{z}_k^{(m)}), \Psi_{j,m} \triangleq \Psi_{j,m}(\mathbf{a}_k^{(j)}, \mathbf{b}_k^{(m)})$.

A detailed derivation of expressions (1)–(4) is provided in [6, Sec. VIII-G].

The factorization in (1) presupposes that the measurements $\mathbf{z}_{1:k}$ are observed; accordingly, we use the semicolon “;” in (2)–(4) to indicate that the argument to the right of “;” is fixed. The observed measurements $\mathbf{z}_{1:k}$ reveal additional factorization structures unavailable in the statistical model with $\mathbf{z}_{1:k}$ random. Since for Bayesian estimation we are either interested in marginal posterior pdfs/pmfs or in values of $\mathbf{x}_{0:k}, \mathbf{r}_{0:k}, \mathbf{a}_{1:k}$, and $\mathbf{b}_{1:k}$ that maximize the joint posterior pdf/pmf, we do not need to compute the normalization constant. In case we want to obtain marginal posterior pdfs/pmfs, we can first marginalize the unnormalized version of the joint posterior pdf/pmf on the right-hand side of (1) and then normalize the result after marginalization. Since the marginalization step reduces the dimensionality, computing a normalization constant after marginalization always has a lower computational complexity.

The factorization of the joint posterior pdf/pmf $f(\mathbf{x}_{0:k}, \mathbf{r}_{0:k}, \mathbf{a}_{1:k}, \mathbf{b}_{1:k} | \mathbf{z}_{1:k})$ as given by (1), (2) is represented graphically by the factor graph [56], [57] shown in Fig. 1. In this factor graph, the random variables and pdfs/pmfs involved in $f(\mathbf{x}_{0:k}, \mathbf{r}_{0:k}, \mathbf{a}_{1:k}, \mathbf{b}_{1:k} | \mathbf{z}_{1:k})$ are represented by circles and squares, respectively [56]. A circle is connected with a square if the random variable represented by the circle is involved in the pdf/pmf represented by the square [56]. Note that in the factor graph, we combine the kinematic state of a single target and the corresponding existence indicator variable into a single state variable $\mathbf{y}^j \triangleq [\mathbf{x}_k^{(j)T} \mathbf{r}_k^{(j)T}]^T$ or

$\bar{\mathbf{y}}^m \triangleq [\bar{\mathbf{x}}_k^{(m)T} \bar{\mathbf{r}}_k^{(m)T}]^T$, which is represented by a single variable node.

Marginalizing out the redundant measurement-oriented DA vector $\mathbf{b}_{1:k}$ from $f(\mathbf{x}_{0:k}, \mathbf{r}_{0:k}, \mathbf{a}_{1:k}, \mathbf{b}_{1:k} | \mathbf{z}_{1:k})$ yields the posterior pdf/pmf

$$f(\mathbf{x}_{0:k}, \mathbf{r}_{0:k}, \mathbf{a}_{1:k} | \mathbf{z}_{1:k}) = \sum_{\mathbf{b}_{1:k}} f(\mathbf{x}_{0:k}, \mathbf{r}_{0:k}, \mathbf{a}_{1:k}, \mathbf{b}_{1:k} | \mathbf{z}_{1:k}). \quad (5)$$

This marginalization is in fact trivial since for a given $\mathbf{a}_{1:k}$ there is exactly one $\mathbf{b}_{1:k}$ for which (1) is nonzero. The posterior pdf/pmf $f(\mathbf{x}_{0:k}, \mathbf{r}_{0:k}, \mathbf{a}_{1:k} | \mathbf{z}_{1:k})$ in (5) together with expression (1), (2) represents a system model that is identical to the one used by MHT methods¹ and also is a generalization of the one used by JPDA filter methods.

IV. MHT METHODS AND TRACK REPULSION EFFECT

MHT methods aim at calculating the MAP sequence estimate of $\mathbf{x}_{0:k}, \mathbf{r}_{0:k}$, and $\mathbf{a}_{1:k}$ from $\mathbf{z}_{1:k}$. This is done in two steps. First, the joint MAP sequence estimate of $\mathbf{r}_{0:k}$ and $\mathbf{a}_{1:k}$ is calculated, i.e.,

$$(\hat{\mathbf{r}}_{0:k}^{\text{MAP}}, \hat{\mathbf{a}}_{1:k}^{\text{MAP}}) = \underset{\mathbf{r}_{0:k}, \mathbf{a}_{1:k}}{\operatorname{argmax}} \int f(\mathbf{x}_{0:k}, \mathbf{r}_{0:k}, \mathbf{a}_{1:k} | \mathbf{z}_{1:k}) d\mathbf{x}_{0:k}. \quad (6)$$

This is equivalent to searching for the PT-measurement association $(\mathbf{r}_{0:k}, \mathbf{a}_{1:k})$ that is most probable given the measurements $\mathbf{z}_{1:k}$ [2]. The marginalization operation $\int d\mathbf{x}_{0:k}$ can be performed sequentially and, assuming linear-Gaussian state-transition and measurement models, also in closed form. On the other hand, the number of PT-measurement associations $(\mathbf{r}_{0:k}, \mathbf{a}_{1:k})$ that must be searched to perform the maximization (6) grows exponentially with k .

In the second step, the MAP estimate of $\mathbf{x}_{0:k}$ given $\mathbf{r}_{0:k} = \hat{\mathbf{r}}_{0:k}^{\text{MAP}}$ and $\mathbf{a}_{1:k} = \hat{\mathbf{a}}_{1:k}^{\text{MAP}}$ is calculated, i.e.,

$$\begin{aligned} \hat{\mathbf{x}}_{0:k}^{\text{MAP}} &= \underset{\mathbf{x}_{0:k}}{\operatorname{argmax}} f(\mathbf{x}_{0:k} | \hat{\mathbf{r}}_{0:k}^{\text{MAP}}, \hat{\mathbf{a}}_{1:k}^{\text{MAP}}, \mathbf{z}_{1:k}) \\ &= \underset{\mathbf{x}_{0:k}}{\operatorname{argmax}} f(\mathbf{x}_{0:k}, \hat{\mathbf{r}}_{0:k}^{\text{MAP}}, \hat{\mathbf{a}}_{1:k}^{\text{MAP}} | \mathbf{z}_{1:k}). \end{aligned} \quad (7)$$

Here, for $(\mathbf{r}_{0:k}, \mathbf{a}_{1:k}) = (\hat{\mathbf{r}}_{0:k}^{\text{MAP}}, \hat{\mathbf{a}}_{1:k}^{\text{MAP}})$ fixed, the joint posterior $f(\mathbf{x}_{0:k}, \hat{\mathbf{r}}_{0:k}^{\text{MAP}}, \hat{\mathbf{a}}_{1:k}^{\text{MAP}} | \mathbf{z}_{1:k})$ in (5) factorizes into the posterior pdfs/pmfs $f(\mathbf{x}_{0:k}^{(j)}, \hat{\mathbf{r}}_{0:k}^{(j) \text{MAP}}, \hat{\mathbf{a}}_{1:k}^{(j) \text{MAP}} | \mathbf{z}_{1:k})$, one for each PT j . The maximization (7) is then equivalent to a Bayesian smoothing operation—carried out by using, e.g., the Kalman smoother—for each PT j in parallel with, and independently of, all the other PTs.

In a practical implementation, to reduce the computational complexity, the MAP estimates of the PT-measurement association sequence in (6) and of the kinematic PT state sequence in (7) are computed only over a sliding window of N consecutive times steps. A hard decision on the PT-measurement association is made only at the oldest step within the window.

¹ In the MHT literature, the two vectors $\mathbf{r}_{0:k}$ and $\mathbf{a}_{1:k}$ are typically represented by a single equivalent vector $\mathbf{q}_{0:k}$ referred to as a global hypothesis.

The maximization in (6) is performed using the Auction algorithm or multiple frame assignment [38].

MHT can be formulated in a hypothesis-oriented [2] and a track-oriented [16], [24], [25] manner. The two formulations are equivalent under the assumption that target births and clutter follow Poisson point processes [59]; however, the track-oriented MHT formulation offers a more compact representation of the DA problem. The computational complexity of the original hypothesis-oriented MHT formulation is still problematic due to the high number of hypotheses. However, it can be reduced by discarding unlikely hypotheses using, e.g., an efficient m -best assignment algorithm [22], [23]. The more efficient track-oriented MHT methods [16], [24], [25] represent the DA hypotheses by a set of tree structures, where each tree represents the possible DA histories of a single PT. The most likely hypothesis is then found by choosing a leaf node from each single-PT tree so that no measurement is used by more than one PT. A fast hypothesis search is enabled by combinatorial optimization methods [26], [60], [61], [62].

A limitation of MHT methods known as track repulsion can arise when targets come in close proximity. The estimated tracks then tend to have a larger distance than the true tracks. This effect is a consequence of performing hard DA by considering for PT state estimation (see (7)) only the single PT-measurement association ($\hat{p}_{0:k}^{\text{MAP}}, \hat{a}_{1:k}^{\text{MAP}}$) in (6). For an illustration, consider two targets that move on parallel tracks. We assume that the states of the targets include the targets' positions; furthermore, each target generates one measurement, which is the target's position plus Gaussian measurement noise. The detection probability is assumed to be one and there are no clutter measurements. The target tracks are supposed to be in close proximity in the sense that their distance is significantly smaller than the standard deviation of the measurement noise. The MAP decision rule (6) now always assigns to each measurement the nearest target. However, due to the small distance between the two targets, this association is incorrect if the measurement closest to a given target originated from the respective other target. A number of consecutive incorrect associations will then cause the position estimates in (7) to be more distant than the actual targets; this effect is known as track repulsion. The track repulsion effect is related to the use of MAP estimation and not to the choice of a specific optimization method or the number of hypotheses. It gets worse as the distance between the targets decreases.

The reason why “hard” DA strategies such as those performed by MHT cause repulsion between close targets can be illustrated in a simple static one-dimensional (1-D) scenario. Consider two static 1-D target states $x^{(1)} = d/2$ and $x^{(2)} = -d/2$, separated by distance d . The targets generate the measurements $z^{(1)} \sim \mathcal{N}(x^{(1)}, 1) = \mathcal{N}(d/2, 1)$ and $z^{(2)} \sim \mathcal{N}(x^{(2)}, 1) = \mathcal{N}(-d/2, 1)$, which are assumed statistically independent. Fig. 2 visualizes the underlying geometry in an equivalent 2-D state space representing the 2-D joint target state $\mathbf{x} \triangleq [x^{(1)} \ x^{(2)}]^T = [d/2 \ -d/2]^T$ and the distribution of the 2-D joint measurement $[z^{(1)} \ z^{(2)}]^T$. We estimate the joint target state $\hat{\mathbf{x}}$ from the measurements $z^{(1)}$ and $z^{(2)}$ using




FIGURE 2. Equivalent 2-D geometry of the static 1-D example employing the “hard association” estimator $\hat{\mathbf{x}}$, which is $[z^{(1)} \ z^{(2)}]^T$ for $z^{(1)} \geq z^{(2)}$ and $[z^{(2)} \ z^{(1)}]^T$ for $z^{(1)} < z^{(2)}$. The level-curves visualize the distribution of $\hat{\mathbf{x}}$. The true state \mathbf{x} and its mirror image (mirrored on the $z_1 = z_2$ line) are shown as a blue and a white bullet, respectively. When the 2-D joint measurement $[z^{(1)} \ z^{(2)}]^T$ crosses over the $z_1 = z_2$ line, the level-curves are reflected, indicating the possibility of an incorrect hard association.




FIGURE 3. Same as Fig. 2, but \mathbf{x} rotated by 45° .

a “hard association” estimator $\hat{\mathbf{x}} = [\hat{x}^{(1)} \ \hat{x}^{(2)}]^T$ defined as $\hat{\mathbf{x}} = [z^{(1)} \ z^{(2)}]^T$ for $z^{(1)} \geq z^{(2)}$ and $\hat{\mathbf{x}} = [z^{(2)} \ z^{(1)}]^T$ for $z^{(1)} < z^{(2)}$. This estimator, which is also known as the global nearest neighbor (GNN) decision rule [1], employs a hard association strategy similar to MHT.

To demonstrate track repulsion effects in the estimator $\hat{\mathbf{x}}$, we consider the expected difference between $\hat{x}^{(1)}$ and $\hat{x}^{(2)}$, i.e., $E\{\hat{x}^{(1)} - \hat{x}^{(2)}\}$. To this end, we first calculate $E\{\hat{\mathbf{x}}\}$. This calculation can be simplified by exploiting the rotational symmetry of the pdf of $[z^{(1)} \ z^{(2)}]^T$ and rotating $\mathbf{x} = [d/2 \ -d/2]^T$ and $\hat{\mathbf{x}}$ by 45° to obtain $\mathbf{x}^r \triangleq [d/\sqrt{2} \ 0]^T$ and $\hat{\mathbf{x}}^r$, respectively. The result is shown in Fig. 3. The expected value of the rotated estimate $\hat{\mathbf{x}}^r$ is given by $E\{\hat{\mathbf{x}}^r\} = \mathbf{x}^r + [\Delta^r \ 0]^T$ with

$$\Delta^r = -2 \int_{-\infty}^0 y \frac{1}{\sqrt{2\pi}} e^{-(y - \frac{d}{\sqrt{2}})^2/2} dy.$$

This can be rewritten as $\Delta^r = \sqrt{\frac{2}{\pi}} e^{-d^2/4} - \sqrt{2}d Q\left(\frac{d}{\sqrt{2}}\right)$, where $Q(\alpha) \triangleq \frac{1}{\sqrt{2\pi}} \int_{\alpha}^{\infty} e^{-\zeta^2/2} d\zeta$ is the probability that a zero-mean Gaussian random variable with unit variance is larger than α . After rotating back, we obtain the expectation of the original estimator as $E\{\hat{\mathbf{x}}\} = [E\{\hat{\mathbf{x}}^{(1)}\} \ E\{\hat{\mathbf{x}}^{(2)}\}]^T$ with $E\{\hat{\mathbf{x}}^{(1)}\} = \frac{d}{2} + \Delta = x^{(1)} + \Delta$ and $E\{\hat{\mathbf{x}}^{(2)}\} = -\frac{d}{2} - \Delta = x^{(2)} - \Delta$, where

$$\Delta = \frac{\Delta^r}{\sqrt{2}} = \frac{1}{\sqrt{\pi}} e^{-d^2/4} - d Q\left(\frac{d}{\sqrt{2}}\right).$$

Thus, the expected distance between $\hat{\mathbf{x}}^{(1)}$ and $\hat{\mathbf{x}}^{(2)}$ is

$$E\{\hat{\mathbf{x}}^{(1)} - \hat{\mathbf{x}}^{(2)}\} = E\{\hat{\mathbf{x}}^{(1)}\} - E\{\hat{\mathbf{x}}^{(2)}\} = d + 2\Delta,$$

which is seen to be larger by 2Δ than the actual distance between the target states, $x^{(1)} - x^{(2)} = d$. We note that Δ is positive and increases with decreasing d . This analysis indicates that the considered hard association estimator leads to a repulsion of the target state estimates $\hat{\mathbf{x}}^{(1)}$ and $\hat{\mathbf{x}}^{(2)}$.

The track repulsion effect can increase the estimation error and lead to track swaps (or identity switches) [63], [64]. In particular, in a scenario with crossing tracks, the estimated tracks tend to “bounce” more often than they cross. A method for mitigating track repulsion was presented in [47].

V. JPDA FILTER METHODS AND TRACK COALESCENCE EFFECT

The JPDA filter aims to calculate the MMSE estimates of the kinematic target states $\mathbf{x}_k^{(j)}$, i.e.,

$$\hat{\mathbf{x}}_k^{(j)\text{MMSE}} = \int \mathbf{x}_k^{(j)} f(\mathbf{x}_k^{(j)} | \mathbf{z}_{1:k}) d\mathbf{x}_k^{(j)}, \quad j = 1, \dots, L_k. \quad (8)$$

Here, L_k denotes the number of actual targets (not PTs) at time k . The posterior pdf of the joint kinematic target state $\mathbf{x}_k = [\mathbf{x}_k^{(1)T} \dots \mathbf{x}_k^{(L_k)T}]^T$ is given by

$$\begin{aligned} f(\mathbf{x}_k | \mathbf{z}_{1:k}) &= \sum_{\mathbf{a}_k} f(\mathbf{x}_k, \mathbf{a}_k | \mathbf{z}_{1:k}) \\ &= \sum_{\mathbf{a}_k} f(\mathbf{x}_k | \mathbf{a}_k, \mathbf{z}_{1:k}) p(\mathbf{a}_k | \mathbf{z}_{1:k}). \end{aligned} \quad (9)$$

This sum over all possible association events \mathbf{a}_k is also referred to as exhaustive averaging.

The JPDA filter is based on the approximation of the posterior DA pmf $p(\mathbf{a}_k | \mathbf{z}_{1:k})$ involved in (9) by the product of its marginals, i.e. [1], [6]

$$p(\mathbf{a}_k | \mathbf{z}_{1:k}) \approx \prod_{j=1}^{L_k} p(a_k^{(j)} | \mathbf{z}_{1:k}), \quad (10)$$

with

$$p(a_k^{(j)} | \mathbf{z}_{1:k}) = \sum_{\mathbf{a}_k^{(\sim j)}} p(\mathbf{a}_k | \mathbf{z}_{1:k}). \quad (11)$$

Here, $\mathbf{a}_k^{(\sim j)}$ denotes the vector \mathbf{a}_k with the j th entry $a_k^{(j)}$ removed. The posterior DA pmf $p(\mathbf{a}_k | \mathbf{z}_{1:k})$ involved in (9) and

(11) is given as [6]

$$p(\mathbf{a}_k | \mathbf{z}_{1:k}) \propto \prod_{j=1}^{L_k} w_k^{(j, a_k^{(j)})}, \quad (12)$$

with

$$w_k^{(j, a_k^{(j)})} \triangleq \int q(\mathbf{x}_k^{(j)}, 1, a_k^{(j)}; \mathbf{z}_k) f(\mathbf{x}_k^{(j)} | \mathbf{z}_{1:k-1}) d\mathbf{x}_k^{(j)}, \quad (13)$$

for $j = 1, \dots, L_k$ and $a_k^{(j)} = 0, \dots, M_k$. Here, the function $q(\mathbf{x}_k^{(j)}, 1, a_k^{(j)}; \mathbf{z}_k)$ is given by (3) with $\mathbf{x}_k^{(j)}$ replaced by $\mathbf{x}_k^{(j)}$, and $f(\mathbf{x}_k^{(j)} | \mathbf{z}_{1:k-1})$ is the “predicted” posterior pdf of $\mathbf{x}_k^{(j)}$. Furthermore, based on the assumption that, for $a_k^{(j)} = m$, the measurement $\mathbf{z}_k^{(m)}$ is conditionally independent given $\mathbf{x}_k^{(j)}$ of all past and future measurements $\mathbf{z}_{k'}^{(m')}$ and states $\mathbf{x}_{k'}^{(j)}$ with $k' \neq k$, it is shown in [6] that the pdf $f(\mathbf{x}_k | \mathbf{a}_k, \mathbf{z}_{1:k})$ involved in (9) factorizes as

$$f(\mathbf{x}_k | \mathbf{a}_k, \mathbf{z}_{1:k}) = \prod_{j=1}^{L_k} f(\mathbf{x}_k^{(j)} | a_k^{(j)}, \mathbf{z}_{1:k}). \quad (14)$$

Then, by inserting (10) into (9) and using the factorization (14), we arrive at the approximate posterior pdf

$$\begin{aligned} f(\mathbf{x}_k | \mathbf{z}_{1:k}) &\approx \sum_{\mathbf{a}_k} \prod_{j=1}^{L_k} f(\mathbf{x}_k^{(j)} | a_k^{(j)}, \mathbf{z}_{1:k}) p(a_k^{(j)} | \mathbf{z}_{1:k}) \\ &= \prod_{j=1}^{L_k} \sum_{a_k^{(j)}=0}^{M_k} f(\mathbf{x}_k^{(j)} | a_k^{(j)}, \mathbf{z}_{1:k}) p(a_k^{(j)} | \mathbf{z}_{1:k}) \\ &= \prod_{j=1}^{L_k} \tilde{f}(\mathbf{x}_k^{(j)} | \mathbf{z}_{1:k}), \end{aligned} \quad (15)$$

with

$$\tilde{f}(\mathbf{x}_k^{(j)} | \mathbf{z}_{1:k}) \triangleq \sum_{a_k^{(j)}=0}^{M_k} f(\mathbf{x}_k^{(j)} | a_k^{(j)}, \mathbf{z}_{1:k}) p(a_k^{(j)} | \mathbf{z}_{1:k}). \quad (16)$$

Here, the conditional pdf $f(\mathbf{x}_k^{(j)} | a_k^{(j)}, \mathbf{z}_{1:k})$ is given as follows [6]: for $a_k^{(j)} = m \in \{1, \dots, M_k\}$,

$$\begin{aligned} f(\mathbf{x}_k^{(j)} | a_k^{(j)} = m, \mathbf{z}_{1:k}) &= \frac{f(\mathbf{z}_k^{(m)} | \mathbf{x}_k^{(j)}) f(\mathbf{x}_k^{(j)} | \mathbf{z}_{1:k-1})}{\int f(\mathbf{z}_k^{(m)} | \mathbf{x}_k^{(j')}) f(\mathbf{x}_k^{(j')} | \mathbf{z}_{1:k-1}) d\mathbf{x}_k^{(j')}, \quad m \in \{1, \dots, M_k\}, \end{aligned} \quad (17)$$

and for $a_k^{(j)} = 0$,

$$f(\mathbf{x}_k^{(j)} | a_k^{(j)} = 0, \mathbf{z}_{1:k}) = f(\mathbf{x}_k^{(j)} | \mathbf{z}_{1:k-1}). \quad (18)$$

We note that, based on the approximation (10), the pdfs $\tilde{f}(\mathbf{x}_k^{(j)} | \mathbf{z}_{1:k})$ in (15) and (16) approximate the marginal posterior pdfs $f(\mathbf{x}_k^{(j)} | \mathbf{z}_{1:k})$. An important, often overlooked property

of the JPDA solution in (16) is that if we have two identical predicted pdfs $f(\mathbf{x}_k^{(j)}|\mathbf{z}_{1:k-1})$ —which enter (16) via (17) and (18)—for two different targets $j \neq j'$, then we also have identical approximate marginal posterior pdfs [1], [6], i.e.,

$$\begin{aligned} f(\mathbf{x}_k^{(j')}|\mathbf{z}_{1:k-1}) &= f(\mathbf{x}_k^{(j)}|\mathbf{z}_{1:k-1}) \\ \Rightarrow \tilde{f}(\mathbf{x}_k^{(j')}|\mathbf{z}_{1:k}) &= \tilde{f}(\mathbf{x}_k^{(j)}|\mathbf{z}_{1:k}). \end{aligned} \quad (19)$$

This is due to symmetry in the computation of (11) and (16) and can easily be verified.

In conventional JPDA filter methods, differently from MHT methods, the sequence of existence indicators $\mathbf{r}_{0:k}$, and thus also the number of targets L_k , are assumed known. In practical implementations, a heuristic for track initialization and termination provides an estimate of $\mathbf{r}_{0:k}$. For linear and Gaussian measurement and motion models and a Gaussian prior for the kinematic states $\mathbf{x}_k^{(j)}$, as assumed in the original formulation of the JPDA filter [1], [7], [9], the approximate marginal posterior pdfs $\tilde{f}(\mathbf{x}_k^{(j)}|\mathbf{z}_{1:k})$ in (16) are Gaussian mixture pdfs, which are further approximated by Gaussian pdfs.

A deficiency of the JPDA filter is the track coalescence effect: when targets come close to each other, the estimated tracks tend to merge and become indistinguishable. This behavior is due to the facts that (i) the predicted marginal posterior pdfs $f(\mathbf{x}_k^{(j)}|\mathbf{z}_{1:k-1})$ involved in (17) and (18) become similar when the targets come close to each other, and (ii) the approximate marginal posterior pdfs $\tilde{f}(\mathbf{x}_k^{(j)}|\mathbf{z}_{1:k})$ are calculated via expression (16). More specifically, performing “soft DA” by means of the summation over all possible associations $a_k^{(j)} \in \{0, 1, \dots, M_k\}$, as done in (16), has the effect that targets j with similar predicted marginal posterior pdfs $f(\mathbf{x}_k^{(j)}|\mathbf{z}_{1:k-1})$ —which enter (16) via (17) and (18)—tend to also have similar approximate marginal posterior pdfs $\tilde{f}(\mathbf{x}_k^{(j)}|\mathbf{z}_{1:k})$ (cf. (19)), which is due to symmetry in the computations (11) and (16). This, in turn, results in similar MMSE state estimates according to (8) (with $f(\mathbf{x}_k^{(j)}|\mathbf{z}_{1:k})$ replaced by $\tilde{f}(\mathbf{x}_k^{(j)}|\mathbf{z}_{1:k})$). A detailed analysis of the track coalescence effect will be presented in Section VII.

Variants and extensions of the JPDA filter include the JIPDA filter [3], [17], the JPDA* filter [45], and the set JPDA (SJPDA) filter [48]. The JIPDA filter extends the JPDA filter by a binary existence indicator to systematically account for target existence, but suffers from track coalescence just as the conventional JPDA filter. The JPDA* filter mitigates track coalescence by pruning target-measurement associations. However, this comes at the cost of a reduced tracking performance in more challenging tracking scenarios with a high number of clutter measurements and missed detections. The SJPDA filter is based on the OSPA estimator [48], which is in turn based on the minimization of the mean OSPA metric [65]. Similar to the JPDA* filter, the SJPDA filter exhibits reduced track coalescence effects. Efficient implementations are based on convex optimization techniques [66].

To maintain track continuity, the SJPDA filter relies on post-processing techniques [67], which are not required for the other JPDA filter variants or MHT methods.

A potential limitation of the JPDA filter and its variants is the fact that their complexity scales exponentially with the number of targets and the number of measurements, which is due to the marginalization operation in (11) or equivalent marginalization operations. Gating and clustering strategies can significantly reduce the complexity in many cases, but fail when many objects come in close proximity. A potential solution in such scenarios is to switch to JPDA-type tracking methods of lower complexity such as the methods proposed in [19] and [68]. Other solutions rely on the pruning of target-measurement associations—which is inherently done in the JPDA* filter—or exploit the potential independence of target-measurement associations [20].

VI. BP METHOD

BP-based MTT methods [6], [12], [15], [36], [42], [55], [69], [70] aim at computing the marginal posterior pdf/pmf $f(\mathbf{x}_k^{(j)}, r_k^{(j)}|\mathbf{z}_{1:k})$ for each PT $j \in \{1, \dots, J_k\}$. This marginal posterior pdf/pmf is then used to perform target detection and MMSE state estimation. In what follows, we will consider the specific BP method introduced in [6]. For target detection, the marginal posterior pmf $p(r_k^{(j)}|\mathbf{z}_{1:k})$ of the existence indicator $r_k^{(j)}$ is obtained from $f(\mathbf{x}_k^{(j)}, r_k^{(j)}|\mathbf{z}_{1:k})$ as

$$p(r_k^{(j)}|\mathbf{z}_{1:k}) = \int f(\mathbf{x}_k^{(j)}, r_k^{(j)}|\mathbf{z}_{1:k}) d\mathbf{x}_k^{(j)}.$$

PT j is then detected—i.e., declared to exist—if $p(r_k^{(j)}|\mathbf{z}_{1:k})$ is larger than a predefined threshold P_{th} . Next, for all PTs that are declared to exist, MMSE state estimation is performed according to

$$\hat{\mathbf{x}}_k^{(j)\text{MMSE}} = \int \mathbf{x}_k^{(j)} f(\mathbf{x}_k^{(j)}, r_k^{(j)}|\mathbf{z}_{1:k}) d\mathbf{x}_k^{(j)}, \quad (20)$$

with

$$f(\mathbf{x}_k^{(j)}|r_k^{(j)} = 1, \mathbf{z}_{1:k}) = \frac{f(\mathbf{x}_k^{(j)}, r_k^{(j)} = 1|\mathbf{z}_{1:k})}{p(r_k^{(j)} = 1|\mathbf{z}_{1:k})}. \quad (21)$$

It remains to calculate $f(\mathbf{x}_k^{(j)}, r_k^{(j)}|\mathbf{z}_{1:k})$. We have

$$f(\mathbf{x}_k^{(j)}, r_k^{(j)}|\mathbf{z}_{1:k}) = \int \sum_{r_k^{(\sim j)}} f(\mathbf{x}_k, \mathbf{r}_k|\mathbf{z}_{1:k}) d\mathbf{x}_k^{(\sim j)},$$

where $\mathbf{r}_k^{(\sim j)}$ denotes \mathbf{r}_k with the entry $r_k^{(j)}$ removed, $\mathbf{x}_k^{(\sim j)}$ denotes \mathbf{x}_k with the subvector $\mathbf{x}_k^{(j)}$ removed, and $f(\mathbf{x}_k, \mathbf{r}_k|\mathbf{z}_{1:k})$ is a marginal pdf/pmf of the joint posterior pdf/pmf $f(\mathbf{x}_{0:k}, \mathbf{r}_{0:k}, \mathbf{a}_{1:k}, \mathbf{b}_{1:k}|\mathbf{z}_{1:k})$, i.e.,

$$\begin{aligned} &f(\mathbf{x}_k, \mathbf{r}_k|\mathbf{z}_{1:k}) \\ &= \int \sum_{\mathbf{r}_{0:k-1}} \sum_{\mathbf{a}_{1:k}} \sum_{\mathbf{b}_{1:k}} f(\mathbf{x}_{0:k}, \mathbf{r}_{0:k}, \mathbf{a}_{1:k}, \mathbf{b}_{1:k}|\mathbf{z}_{1:k}) d\mathbf{x}_{0:k-1}. \end{aligned} \quad (22)$$

By inserting the factorization (1), (2) of $f(\mathbf{x}_{0:k}, \mathbf{r}_{0:k}, \mathbf{a}_{1:k}, \mathbf{b}_{1:k} | \mathbf{z}_{1:k})$ and carrying out the marginalizations with respect to $\mathbf{x}_{0:k-2}, \mathbf{r}_{0:k-2}, \mathbf{a}_{1:k-1}$, and $\mathbf{b}_{1:k-1}$, it can be shown that [6]

$$f(\mathbf{x}_k, \mathbf{r}_k | \mathbf{z}_{1:k}) = \int \sum_{\mathbf{r}_{k-1}} \sum_{\mathbf{a}_k} \sum_{\mathbf{b}_k} f(\mathbf{x}_{k-1}, \mathbf{r}_{k-1} | \mathbf{z}_{1:k-1}) \times g(\mathbf{x}_k, \mathbf{r}_k, \mathbf{a}_k, \mathbf{b}_k, \mathbf{x}_{k-1}, \mathbf{r}_{k-1}; \mathbf{z}_k) d\mathbf{x}_{k-1}, \quad (23)$$

where $g(\mathbf{x}_k, \mathbf{r}_k, \mathbf{a}_k, \mathbf{b}_k, \mathbf{x}_{k-1}, \mathbf{r}_{k-1}; \mathbf{z}_k)$ was defined in (2).

Two observations can be made at this point. First, it can be concluded from (20)–(23) that the MMSE estimator $\hat{\mathbf{x}}_k^{(j)\text{MMSE}}$ in (20) is a “single-scan solution” in the sense that, as shown by (23), the marginal posterior pdf/pmf $f(\mathbf{x}_k, \mathbf{r}_k | \mathbf{z}_{1:k})$ at time k can be directly obtained from the marginal posterior pdf/pmf at time $k-1$, $f(\mathbf{x}_{k-1}, \mathbf{r}_{k-1} | \mathbf{z}_{1:k-1})$, and the current measurement \mathbf{z}_k (which enters via $g(\mathbf{x}_k, \mathbf{r}_k, \mathbf{a}_k, \mathbf{b}_k, \mathbf{x}_{k-1}, \mathbf{r}_{k-1}; \mathbf{z}_k)$). In other words, $f(\mathbf{x}_{k-1}, \mathbf{r}_{k-1} | \mathbf{z}_{1:k-1})$ subsumes and provides all the relevant information from past time steps. We note that a single-scan method based on MMSE estimation does not necessarily perform worse than a multiscan method based on MAP estimation (such as an MHT method).

Second, the computational complexity of evaluating the expressions (23) and (2) is much smaller than that of directly performing the marginalization in (22). Indeed, the summation over $\mathbf{r}_{0:k-1}$ in (22) involves $2^{J_0} \dots 2^{J_{k-1}} = 2^{\sum_{k'=0}^{k-1} J_{k'}}$ terms, whereas the repeated application of (23) from 0 to $k-1$ only involves $\sum_{k'=0}^{k-1} 2^{J_{k'}}$ terms. This complexity reduction is a consequence of the temporal factorization structure of the MTT problem. However, we still have to perform the marginalizations in (23), whose complexity scales exponentially in the number of PTs and the number of measurements. These marginalizations can be computed in an efficient (though approximate) manner using the BP approach described in the following, which has only a linear complexity scaling.

The BP method operates on a factor graph representing the statistical model of the considered estimation problem [56], [57]. The factor graph for our statistical model is shown in Fig. 1. The marginal pdfs/pmfs $f(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)} | \mathbf{z}_{1:k})$, $j = 1, \dots, J_k$ needed for target detection and state estimation as discussed above are calculated efficiently by performing local operations corresponding to the individual graph nodes and exchanging the results of these local operations—called “messages”—along the graph edges [6], [42]. In the following, we outline the calculation of these messages and the resulting beliefs (see also Fig. 1). More details are provided in [6], and a particle-based implementation is described in [42].

- 1) *Prediction*: First, state prediction is performed separately for each legacy PT $j \in \{1, \dots, J_{k-1}\}$ by sending a message $\alpha_k^{(j)}(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)})$ from factor node “ $f(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)} | \mathbf{x}_{k-1}^{(j)}, \mathbf{r}_{k-1}^{(j)})$ ” to variable node

“ $[\mathbf{x}_k^{(j)\text{T}}, \mathbf{r}_k^{(j)\text{T}}]^\text{T}$ ”. This message is calculated according to

$$\alpha_k^{(j)}(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)}) = \sum_{\mathbf{r}_{k-1}^{(j)} \in \{0,1\}} \int f(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)} | \mathbf{x}_{k-1}^{(j)}, \mathbf{r}_{k-1}^{(j)}) \times f(\mathbf{x}_{k-1}^{(j)}, \mathbf{r}_{k-1}^{(j)} | \mathbf{z}_{1:k-1}) d\mathbf{x}_{k-1}^{(j)}.$$

Here, $f(\mathbf{x}_{k-1}^{(j)}, \mathbf{r}_{k-1}^{(j)} | \mathbf{z}_{1:k-1})$ is the marginal posterior pdf/pmf of legacy PT $j \in \{1, \dots, J_{k-1}\}$ at time $k-1$.

- 2) *Measurement Evaluation*: Next, for each legacy PT $j \in \{1, \dots, J_{k-1}\}$, the following message is passed from factor node “ $q(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)}, \mathbf{a}_k^{(j)}; \mathbf{z}_k)$ ” to variable node “ $\mathbf{a}_k^{(j)}$ ”:

$$\beta_k^{(j)}(\mathbf{a}_k^{(j)}) = \sum_{\mathbf{r}_k^{(j)} \in \{0,1\}} \int q(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)}, \mathbf{a}_k^{(j)}; \mathbf{z}_k) \alpha_k^{(j)}(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)}) d\mathbf{x}_k^{(j)}.$$

Similarly, for each new PT $m \in \{1, \dots, M_k\}$, the following message is passed from factor node “ $v(\bar{\mathbf{x}}_k^{(m)}, \bar{\mathbf{r}}_k^{(m)}, \mathbf{b}_k^{(m)}; \mathbf{z}_k)$ ” to variable node “ $\mathbf{b}_k^{(m)}$ ”:

$$\xi_k^{(m)}(\mathbf{b}_k^{(m)}) = \sum_{\bar{\mathbf{r}}_k^{(m)} \in \{0,1\}} \int v(\bar{\mathbf{x}}_k^{(m)}, \bar{\mathbf{r}}_k^{(m)}, \mathbf{b}_k^{(m)}; \mathbf{z}_k) d\bar{\mathbf{x}}_k^{(m)}.$$

- 3) *Probabilistic DA*: Probabilistic DA is now performed iteratively by sending, at message passing iteration $p \in \{1, \dots, P\}$, a message $\varphi_k^{[p](j \rightarrow m)}$ from variable node “ $\mathbf{a}_k^{(j)}$ ” via factor node “ $\Psi_{j,m}(\mathbf{a}_k^{(j)}, \mathbf{b}_k^{(m)})$ ” to variable node “ $\mathbf{b}_k^{(m)}$ ” and a message $\nu_k^{[p](m \rightarrow j)}$ from variable node “ $\mathbf{b}_k^{(m)}$ ” via factor node “ $\Psi_{j,m}(\mathbf{a}_k^{(j)}, \mathbf{b}_k^{(m)})$ ” to variable node “ $\mathbf{a}_k^{(j)}$ ”. These messages are calculated according to

$$\varphi_k^{[p](j \rightarrow m)} = \frac{\beta_k^{(j)}(\mathbf{a}_k^{(j)})}{\beta_k^{(j)}(0) + \sum_{\substack{m'=1 \\ m' \neq m}}^{M_k} \beta_k^{(j)}(\mathbf{a}_k^{(j)}) \nu_k^{[p](m' \rightarrow j)}},$$

$$\nu_k^{[p](m \rightarrow j)} = \frac{\xi_k^{(m)}(\mathbf{b}_k^{(m)})}{\xi_k^{(m)}(0) + \sum_{\substack{j'=1 \\ j' \neq j}}^{J_{k-1}} \xi_k^{(m)}(\mathbf{b}_k^{(m)}) \varphi_k^{[p-1](j' \rightarrow m)}},$$

for $j = 1, \dots, J_{k-1}$ and $m = 1, \dots, M_k$. This iterative procedure is initialized with $\varphi_k^{[0](j \rightarrow m)} = 1$. After the last iteration $p = P$, a message $\kappa_k^{(j)}(\mathbf{a}_k^{(j)})$ is passed from variable node “ $\mathbf{a}_k^{(j)}$ ” to factor node “ $q(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)}, \mathbf{a}_k^{(j)}; \mathbf{z}_k)$ ”. This message is given for $\mathbf{a}_k^{(j)} = m \in \{0, 1, \dots, M_k\}$ by

$$\kappa_k^{(j)}(m) = \frac{\beta_k^{(j)}(\mathbf{a}_k^{(j)}) \nu_k^{[P](m \rightarrow j)}}{\beta_k^{(j)}(0) + \sum_{m'=1}^{M_k} \beta_k^{(j)}(\mathbf{a}_k^{(j)}) \nu_k^{[P](m' \rightarrow j)}}.$$

Similarly, a message $\iota_k^{(m)}(b_k^{(m)})$ is passed from variable node “ $b_k^{(m)}$ ” to factor node “ $v(\bar{\mathbf{x}}_k^{(m)}, \bar{\mathbf{r}}_k^{(m)}, b_k^{(m)}; \mathbf{z}_k^{(m)})$ ”. It is given for $b_k^{(m)} = j \in \{0, 1, \dots, J_{k-1}\}$ by

$$\iota_k^{(m)}(j) = \frac{\xi_k^{(m)}(j) \varphi_k^{[P](j \rightarrow m)}}{\xi_k^{(m)}(0) + \sum_{j'=1}^{J_{k-1}} \xi_k^{(m)}(j') \varphi_k^{[P](j' \rightarrow m)}}.$$

- 4) *Measurement Update*: Measurement update steps are next applied for the legacy PTs and the new PTs. For each legacy PT $j \in \{1, \dots, J_{k-1}\}$, the message

$$\gamma_k^{(j)}(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)}) = \sum_{a_k^{(j)}=0}^{M_k} q(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)}, a_k^{(j)}; \mathbf{z}_k) \kappa_k^{(j)}(a_k^{(j)})$$

is passed from factor node “ $q(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)}, a_k^{(j)}; \mathbf{z}_k)$ ” to variable node “ $[\mathbf{x}_k^{(j)T} \mathbf{r}_k^{(j)T}]^T$ ”. Furthermore, for each new PT $m \in \{1, \dots, M_k\}$, the message

$$\varsigma_k^{(m)}(\bar{\mathbf{x}}_k^{(m)}, \bar{\mathbf{r}}_k^{(m)}) = v(\bar{\mathbf{x}}_k^{(m)}, \bar{\mathbf{r}}_k^{(m)}, b_k^{(m)}; \mathbf{z}_k^{(m)}) \iota_k^{(m)}(0)$$

is passed from factor node “ $v(\bar{\mathbf{x}}_k^{(m)}, \bar{\mathbf{r}}_k^{(m)}, b_k^{(m)}; \mathbf{z}_k^{(m)})$ ” to variable node “ $[\bar{\mathbf{x}}_k^{(m)T} \bar{\mathbf{r}}_k^{(m)T}]^T$ ”.

- 5) *Belief calculation*: Finally, for each legacy PT $j \in \{1, \dots, J_{k-1}\}$, a belief $\tilde{f}(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)})$ approximating the marginal posterior pdf/pmf $f(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)} | \mathbf{z}_{1:k})$ is obtained as

$$\tilde{f}(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)}) = \frac{1}{\underline{C}_k^{(j)}} \alpha_k^{(j)}(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)}) \gamma_k^{(j)}(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)}),$$

with $\underline{C}_k^{(j)} \triangleq \sum_{\mathbf{r}_k^{(j)} \in \{0,1\}} \int \alpha_k^{(j)}(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)}) \gamma_k^{(j)}(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)}) d\mathbf{x}_k^{(j)}$. Analogously, for each new PT $m \in \{1, \dots, M_k\}$, a belief $\tilde{f}(\bar{\mathbf{x}}_k^{(m)}, \bar{\mathbf{r}}_k^{(m)})$ approximating the marginal posterior pdf/pmf $f(\bar{\mathbf{x}}_k^{(m)}, \bar{\mathbf{r}}_k^{(m)} | \mathbf{z}_{1:k})$ is obtained as

$$\tilde{f}(\bar{\mathbf{x}}_k^{(m)}, \bar{\mathbf{r}}_k^{(m)}) = \frac{1}{\bar{C}_k^{(m)}} \varsigma_k^{(m)}(\bar{\mathbf{x}}_k^{(m)}, \bar{\mathbf{r}}_k^{(m)}),$$

$$\text{with } \bar{C}_k^{(m)} \triangleq \sum_{\bar{\mathbf{r}}_k^{(m)} \in \{0,1\}} \int \varsigma_k^{(m)}(\bar{\mathbf{x}}_k^{(m)}, \bar{\mathbf{r}}_k^{(m)}) d\bar{\mathbf{x}}_k^{(m)}.$$

This BP method systematically exploits the conditional independence structure of the involved random variables, as expressed by the factorization structure of (1), (2) and represented by the factor graph of Fig. 1, for a large reduction of computational complexity. Here, in particular, the spatial factorization structure across PTs expressed by (1) and (2) is exploited in addition to the temporal factorization structure. Since the factor graph in Fig. 1 has loops, the beliefs $\tilde{f}(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)})$ and $\tilde{f}(\bar{\mathbf{x}}_k^{(m)}, \bar{\mathbf{r}}_k^{(m)})$ only provide approximations to the marginal posterior pdfs/pmfs $f(\mathbf{x}_k^{(j)}, \mathbf{r}_k^{(j)} | \mathbf{z}_{1:k})$ and $f(\bar{\mathbf{x}}_k^{(m)}, \bar{\mathbf{r}}_k^{(m)} | \mathbf{z}_{1:k})$, respectively [56]. In the case of our factor graph, these approximations are sufficiently accurate to yield excellent tracking performance. However, the loops in the factor graph also lead to overconfident beliefs [71], i.e., the spread of the beliefs underestimates the true posterior uncertainty of the respective random variables.

Since the BP method performs a marginalization similar to JPDA filter methods, one may suspect that it suffers from track coalescence to a similar extent. Surprisingly, this is not the case: as analyzed in the next section and further evidenced by our simulation results in Section VIII, the BP method exhibits track coalescence to a lesser extent than traditional JPDA filtering methods. Moreover, it does not suffer from track repulsion effects.

VII. A CLOSER LOOK AT THE TRACK COALESCENCE EFFECT

In this section, we take a closer look at the track coalescence effect and investigate why BP methods exhibit this effect to a lesser extent than JPDA filter methods.

A. JPDA FILTER

Let us first reconsider the JPDA filter from Section V for the case where the number of targets is fixed and known, i.e., $L_k = L$ for all times k . Here, using (12) and (13), the posterior DA pmf $p(a_k | \mathbf{z}_{1:k})$ in (10) can be expressed as

$$p(a_k | \mathbf{z}_{1:k}) \propto \prod_{j=1}^L \int q(\mathbf{x}_k^{(j)}, 1, a_k^{(j)}; \mathbf{z}_k) f(\mathbf{x}_k^{(j)} | \mathbf{z}_{1:k-1}) d\mathbf{x}_k^{(j)}. \quad (24)$$

If $f(\mathbf{x}_k^{(j)} | \mathbf{z}_{1:k-1})$, i.e., the “prior information” about $\mathbf{x}_k^{(j)}$ at time k , is equal for all $j = 1, \dots, L$, then $f(\mathbf{x}_k^{(j)} | a_k^{(j)}, \mathbf{z}_{1:k})$ in (17), (18) is equal for all $a_k^{(j)} = 0, \dots, M_k$. Furthermore, $p(a_k | \mathbf{z}_{1:k})$ is invariant to a permutation of the entries $a_k^{(j)}$ of \mathbf{a}_k . This, in turn, implies that the marginal posterior DA pmfs $p(a_k^{(j)} | \mathbf{z}_{1:k})$, $j = 1, \dots, L$ calculated from $p(\mathbf{a}_k | \mathbf{z}_{1:k})$ according to (11) are equal and, further, the approximate marginal state pdfs $\tilde{f}(\mathbf{x}_k^{(j)} | \mathbf{z}_{1:k})$, $j = 1, \dots, L$ calculated from $p(a_k^{(j)} | \mathbf{z}_{1:k})$ and $f(\mathbf{x}_k^{(j)} | a_k^{(j)}, \mathbf{z}_{1:k})$ according to (16) are equal as well. Therefore, the target state estimates $\mathbf{x}_k^{(j)\text{MMSE}}$ calculated from $\tilde{f}(\mathbf{x}_k^{(j)} | \mathbf{z}_{1:k})$ according to (8) (with $f(\mathbf{x}_k^{(j)} | \mathbf{z}_{1:k})$ replaced by $\tilde{f}(\mathbf{x}_k^{(j)} | \mathbf{z}_{1:k})$) become equal, which means that the estimated tracks merge and, thus, the track coalescence effect is observed. An inspection of (16) and (24) shows that this indistinguishability of targets with the same prior information $f(\mathbf{x}_k^{(j)} | \mathbf{z}_{1:k-1})$ is a direct consequence of the measurement model with measurement-origin uncertainty. We will illustrate this fact by considering two simple scenarios.

The first scenario demonstrates the track coalescence effect without using the Gaussian assumption underlying the JPDA filter. We consider a 1-D state space and two targets that are close to each other. The target states are the targets’ 1-D positions $x_k^{(j)} \in \{-1, 1\}$ for $j = 1, 2$. The targets are observed by a sensor that generates the measurements $z_k^{(1)} = 1$ and $z_k^{(2)} = -1$. Thus, we have $L_k = L = 2$ and $M_k = 2$. There are no clutter measurements and no measurement noise, and the detection probability is assumed to be one. We further assume that each measurement equals either one of the true target positions $x_k^{(j)}$ with equal probability,

i.e., $z_k^{(1)} = 1$ equals $x_k^{(1)}$ or $x_k^{(2)}$ with equal probability $1/2$, and similarly for $z_k^{(2)} = -1$. From these assumptions, it follows that $p(a_k^{(1)} = m | z_{1:k}) = p(a_k^{(2)} = m | z_{1:k}) = 1/2$ for $m = 1, 2$ and $f(x_k^{(j)} | a_k^{(j)} = 1, z_{1:k}) = \delta(x_k^{(j)} - 1)$ and $f(x_k^{(j)} | a_k^{(j)} = 2, z_{1:k}) = \delta(x_k^{(j)} + 1)$ for $j = 1, 2$. Inserting these expressions into (16) yields $\tilde{f}(x_k^{(j)} | z_{1:k}) = \frac{1}{2}\delta(x_k^{(j)} - 1) + \frac{1}{2}\delta(x_k^{(j)} + 1)$ for $j = 1, 2$. Finally, using (8), the MMSE state estimates are obtained as $\hat{x}_k^{(j)\text{MMSE}} = 0$ for $j = 1, 2$. This means that the tracks are merged and, thus, the track coalescence effect is observed.

For the second scenario, we reconsider the simple tracking scenario previously considered in the context of MHT methods in Section IV. We recall this scenario for convenience. There are two targets, whose states include the targets' positions. Each target generates one measurement, which is the target's position plus Gaussian measurement noise, and there are no clutter measurements. Thus, we again have $L_k = L = 2$ and $M_k = 2$. The targets move on parallel tracks and in close proximity, in the sense that the distance between them is significantly smaller than the standard deviation of the measurement noise. If the two targets move in this way for a sufficiently long time, it can be expected that the predicted posterior pdfs $f(x_k^{(1)} | z_{1:k-1})$ and $f(x_k^{(2)} | z_{1:k-1})$, i.e., the "prior information," are approximately equal and that each of the two measurements is approximately equally likely to have originated from either one of the two targets. This assumption results in $f(x_k^{(1)} | a_k^{(1)} = m, z_{1:k})$ and $f(x_k^{(2)} | a_k^{(2)} = m, z_{1:k})$ in (17) being approximately equal for $m = 1, 2$ and $p(a_k^{(1)} = m | z_{1:k})$ and $p(a_k^{(2)} = m | z_{1:k})$ being also approximately equal for $m = 1, 2$. As a consequence, it can be expected that, according to (16), $\tilde{f}(x_k^{(1)} | z_{1:k})$ and $\tilde{f}(x_k^{(2)} | z_{1:k})$ are almost equal. As explained above, this implies that the target state estimates $\hat{x}_k^{(1)\text{MMSE}}$ and $\hat{x}_k^{(2)\text{MMSE}}$ become approximately equal, which means that the track coalescence effect is observed. The track coalescence effect also occurs when one of the two targets is missed, i.e., $M_k = 1$, or there is clutter, i.e., $M_k > 2$, as long as the distance between the targets is smaller than the standard deviation of the measurement noise. In both cases, again, each of the measurements is approximately equally likely to have originated from either one of the two targets. If the measurements follow a more complicated, possibly nonlinear or non-Gaussian measurement model, the track coalescence effect can still be observed, but a characterization of the case where the targets are "in close proximity" may be more difficult.

B. BP METHOD

Although the BP method considered in Section VI is partly related to the JPDA filtering paradigm, it exhibits significantly reduced track coalescence effects. There are two reasons for this fact: (i) the BP method typically uses a particle representation of the marginal posterior pdfs $f(x_k^{(j)} | r_k^{(j)} = 1, z_{1:k})$ in (21) [42], which preserves the multimodality of these pdfs, and (ii) the overconfident nature of the DA solution provided

by the BP method favors the most likely PT-measurement association and thus a separation of target state estimates.

Reason (i) can be explained as follows. In the original JPDA filter, the pdfs $f(x_k^{(j)} | a_k^{(j)}, z_{1:k})$ in (17) and (18) are Gaussian pdfs. As a consequence, the approximate marginal posterior pdfs $\tilde{f}(x_k^{(j)} | z_{1:k})$ calculated according to (16) are Gaussian mixture pdfs. The JPDA filter approximates these multimodal Gaussian mixture pdfs by Gaussian pdfs. This additional approximation exacerbates the track coalescence effect since pdfs $\tilde{f}(x_k^{(j)} | z_{1:k})$ that are well distinguishable may become indistinguishable after the Gaussian approximation. By contrast, the particle representation of the marginal posterior pdfs $f(x_k^{(j)} | r_k^{(j)} = 1, z_{1:k})$ that is used in the BP method is able to capture the multimodality of $f(x_k^{(j)} | r_k^{(j)} = 1, z_{1:k})$. We note that the role of $f(x_k^{(j)} | r_k^{(j)} = 1, z_{1:k})$ in the BP method is similar to that of $f(x_k^{(j)} | z_{1:k})$ in JPDA filter methods.

Reason (ii) is the overconfident nature of the BP method for a factor graph with loops ("loopy BP") [72]. To illustrate this aspect, we consider modified marginal posterior DA pmfs $p_\rho(a_k^{(j)} | z_{1:k}) \triangleq (p(a_k^{(j)} | z_{1:k}))^\rho / C_k^{(j)}$ with exponent $\rho \in (0, \infty)$, where $C_k^{(j)} \triangleq \sum_{a_k^{(j)}=0}^{M_k} (p(a_k^{(j)} | z_{1:k}))^\rho$ and $p(a_k^{(j)} | z_{1:k})$ is calculated according to (11). For $0 < \rho < 1$, $p_\rho(a_k^{(j)} | z_{1:k})$ is underconfident relative to $p(a_k^{(j)} | z_{1:k})$, in the sense that likely PT-measurement associations (corresponding to large values of $p(a_k^{(j)} | z_{1:k})$) are deemphasized and unlikely ones (corresponding to small values of $p(a_k^{(j)} | z_{1:k})$) are emphasized, resulting in a larger spread of the modified pmf $p_\rho(a_k^{(j)} | z_{1:k})$. In particular, for $\rho \rightarrow 0$, $p_\rho(a_k^{(j)} | z_{1:k})$ is the uniform pmf, i.e., all PT-measurement associations are equally likely. For $\rho > 1$, on the other hand, $p_\rho(a_k^{(j)} | z_{1:k})$ is overconfident relative to $p(a_k^{(j)} | z_{1:k})$, in the sense that likely associations are emphasized and unlikely ones are deemphasized, resulting in a smaller spread of $p_\rho(a_k^{(j)} | z_{1:k})$. In particular, for $\rho \rightarrow \infty$, $p_\rho(a_k^{(j)} | z_{1:k})$ is one for the value of $a_k^{(j)}$ where $p(a_k^{(j)} | z_{1:k})$ is largest—i.e., the MAP estimate of $a_k^{(j)}$ —and zero otherwise. The JPDA filter performs soft DA with $\rho = 1$, which leads to track coalescence as discussed in Section VII-A. By contrast, MHT methods perform hard DA, in that they use only the MAP estimate of $a_k^{(j)}$ instead of the entire posterior distribution of $a_k^{(j)}$, corresponding to $p_\rho(a_k^{(j)} | z_{1:k})$ with $\rho \rightarrow \infty$; this leads to track repulsion as explained in Section IV.

The BP method, just as JPDA filtering methods, performs soft DA in the sense that it calculates approximate marginal posterior state pdfs $\tilde{f}(x_k^{(j)} | z_{1:k})$ via a weighted summation over all possible associations, thereby taking into account the entire DA distribution (see (16)). However, in contrast to JPDA filter methods, it relies on loopy BP and thus computes overconfident approximations $\tilde{p}(a_k^{(j)} | z_{1:k})$ of the marginal posterior DA pmfs $p(a_k^{(j)} | z_{1:k})$. These approximations resemble $p_\rho(a_k^{(j)} | z_{1:k})$ for some $\rho > 1$, which means that the BP method emphasizes likely PT-measurement associations and




FIGURE 4. Target trajectories of scenarios (a) S1, (b) S2, (c) S3, and (d) S4. The initial target positions are indicated by circles.

deemphasizes unlikely ones. Thus, the soft DA performed by the BP method is somewhat closer, in regard to the tracking results, to the hard DA performed by MHT methods (which do not exhibit the track coalescence effect). This explains why the BP method exhibits a reduced track coalescence effect. However, the BP solution is still less confident than the solution provided by MHT, which does not reflect association ambiguity by setting $\rho = \infty$. The price paid for the reduced track coalescence is that the BP solution tends to be overconfident.

VIII. SIMULATION STUDY

Next, we present simulation results assessing and comparing the performance of the considered MTT methods in four different scenarios where targets come in close proximity. In particular, we will demonstrate experimentally that the BP method exhibits no track repulsion effect and a significantly reduced track coalescence effect compared to JPDA filter methods.

A. SIMULATION SETUP

We consider the four simulation scenarios S1 through S4 depicted in Fig. 4. In S1 through S3, there are two targets with deterministically chosen trajectories: in S1, the two targets approach each other, move in parallel close to each other, and separate again; in S2, they move in parallel close to each other from the beginning and then separate; and in S3, they cross each other. In S4, on the other hand, six targets are born with uniform spacing on a circle of radius 150m and evolve randomly according to a nearly constant-velocity motion (NCVM) model [73]; they come in close proximity near the origin and then separate again [42]. Each scenario comprises 300 time steps. The minimum distance between the two targets in S1 and S2 is equal to the standard deviation of the measurement noise, σ_v (defined below). The targets remain close to each other for 100 time steps in S1 and for 150 time steps in S2. In S3 and S4, the two and six targets are in a “ $2\sigma_v$ -neighborhood” for about 50 time steps and 30 time steps, respectively. The region of interest (ROI) is $[-750\text{m}, 750\text{m}] \times [-750\text{m}, 750\text{m}]$ for all scenarios.

The target states consist of 2-D position and velocity, i.e., $\mathbf{x}_k^{(j)} = [x_{1,k}^{(j)} \ x_{2,k}^{(j)} \ \dot{x}_{1,k}^{(j)} \ \dot{x}_{2,k}^{(j)}]^T$. The various MTT methods, for all

scenarios S1 through S4, use the NCVM model [73]

$$\mathbf{x}_k^{(j)} = \mathbf{A}\mathbf{x}_{k-1}^{(j)} + \mathbf{u}_k^{(j)}, \quad (25)$$

where $\mathbf{u}_k^{(j)} \sim \mathcal{N}(\mathbf{0}, \Sigma_u)$ is a sequence of independent and identically distributed (iid) 4-D Gaussian random vectors and \mathbf{A} and Σ_u are given by

$$\mathbf{A} = \begin{pmatrix} 1 & 0 & T & 0 \\ 0 & 1 & 0 & T \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad \Sigma_u = \begin{pmatrix} \frac{T^3}{3} & 0 & \frac{T^2}{2} & 0 \\ 0 & \frac{T^3}{3} & 0 & \frac{T^2}{2} \\ \frac{T^2}{2} & 0 & T & 0 \\ 0 & \frac{T^2}{2} & 0 & T \end{pmatrix} \sigma_u^2.$$

Here, $T = 1\text{s}$ and $\sigma_u^2 = 0.1\text{m}^2/\text{s}^4$ for S1, S2, and S3 and $\sigma_u^2 = 0.0001\text{m}^2/\text{s}^4$ for S4. We note that the changes in direction occurring in the target trajectories in S1 and S2 are not explicitly modeled by the NCVM model in (25); however, this model mismatch is compensated by the relatively high driving noise variance σ_u^2 used by the MTT methods in S1 and S2. The survival probability is set to $p_s = 0.995$. Furthermore, $\mu_b = 0.01$ and $f_b(\bar{\mathbf{x}}_k)$ is uniform over the ROI.

The sensor produces noisy measurements of the target positions. More specifically, the target-generated measurements are given by

$$\mathbf{z}_k^{(m)} = [x_{1,k} \ x_{2,k}]^T + \mathbf{v}_k^{(m)},$$

where $\mathbf{v}_k^{(m)} \sim \mathcal{N}(\mathbf{0}, \sigma_v^2 \mathbf{I}_2)$ with $\sigma_v = 10\text{m}$ is an iid sequence of 2-D Gaussian random vectors. The detection probability is set to $p_d = 0.5$. We assume a sensor with perfect resolution, i.e., independently of their distance, the targets can always be resolved in the sense that two different targets do not lead to the same measurement. In addition to the target-generated measurements, there are also clutter measurements. The mean number of clutter measurements is $\mu_c = 10$, and the clutter pdf $f_c(\mathbf{z}_k^{(m)})$ is uniform on the ROI. These measurement parameters and state evolution parameters are used in the various MTT methods; furthermore, the same measurement parameters and, in scenario S4, also the NCVM state evolution parameters are used to generate the measurements. We note that if larger values of σ_u^2 , σ_v^2 , and μ_c and smaller values of p_s and p_d are used to generate the measurements, then a

poorer tracking accuracy of the various MTT methods can be expected.

For each of our four scenarios, we performed 1000 simulation runs, each comprising 300 time steps. For performance evaluation, we use the generalized optimal subpattern assignment (GOSPA) metric [74] based on the L2-norm, averaged over all simulation runs, with parameters $p = 1$, $c = 50$, and $\beta = 2$. The GOSPA metric accounts for both cardinality and state estimation errors, similar to the OSPA metric [65] and the complete OSPA metric [75]. For S1, we compute two additional performance metrics. The first, termed “D-Tracks,” is the distance between the two tracks, i.e., $|\hat{x}_{2,k}^{(1)} - \hat{x}_{2,k}^{(2)}|$, for each time step k . The second, termed “D-Center,” is the average of the distances of the estimated y-coordinates of the two tracks from the y-center (origin), i.e., $(|\hat{x}_{2,k}^{(1)}| + |\hat{x}_{2,k}^{(2)}|)/2$, again for each time step k . A small D-Tracks metric can indicate track coalescence, while the D-Center metric assesses the accuracy of the tracks’ centroid (which is high if the D-Center metric is small). For example, if D-Tracks and D-Center are both small, the estimated tracks may have coalesced, but at least the estimated “joint track” is quite accurate. On the other hand, if D-Center is large, then also the estimated joint track is inaccurate.

B. REDUCED TRACK COALESCENCE IN THE BP METHOD

We first demonstrate and analyze experimentally the reduction of track coalescence exhibited by the BP method of Section VI. In Section VII-B, we argued that this reduction is due to the particle representation of the marginal posterior state pdfs—which, differently from a Gaussian representation, preserves multimodality—and the overconfident nature of the DA-related messages. We now investigate this numerically by combining two choices of track representation (single Gaussian and set of particles) with the DA strategies employed by the BP method and the JPDA filter. In particular, to demonstrate the influence of the particle representation of the marginal posterior state pdfs, we compare a particle implementation of the BP method with a Gaussian implementation; these implementations will be designated as PART-BP and GAUSS-BP, respectively. Furthermore, to demonstrate the influence of the overconfident nature of the DA-related messages in the BP method, we consider modified versions of the particle and Gaussian implementations in which the BP-based approximate marginalization of the association pmf is replaced by an exact marginalization. These latter methods will be designated as PART-EX and GAUSS-EX. To exclude the influence of target detection errors, we temporarily assume that the birth times and birth positions of the targets are perfectly known by the various tracking algorithms. We note that in this case GAUSS-EX coincides with the JPDA filter of Section V, and thus the corresponding implementation equations of GAUSS-EX can be found, e.g., in [1]. Furthermore, we recall that in GAUSS-BP, the exact marginalization of GAUSS-EX is replaced by the BP-based approximate marginalization. PART-BP and PART-EX use 5000 particles to represent the PT states. The implementation equations of

PART-BP are equal to those in [42] for the case of a single sensor and known birth times and locations. We recall that in PART-EX, the BP-based marginalization of PART-BP is replaced by the exact marginalization. The threshold for target confirmation is $P_{th} = 0.5$.

Fig. 5 presents the results for scenario S1. The D-Tracks curves in Fig. 5(a) show that the tracks of GAUSS-BP and PART-BP are close to the ground truth, whereas those of GAUSS-EX and PART-EX tend to merge and exhibit a delay in separating again. This behavior is also reflected by the GOSPA curves in Fig. 5(c): the track coalescence effect causes an increase of the GOSPA error of GAUSS-EX and PART-EX for $k = 200, \dots, 250$. This increase is smaller for PART-EX than for GAUSS-EX. Regarding the D-Center curves in Fig. 5(b), all four methods exhibit a similar behavior. In summary, the results in Fig. 5(a) and (c) demonstrate that both particle representation and BP-based marginalization tend to lead to reduced track coalescence effects, thereby confirming our reasoning of Section VII.

Finally, we demonstrate that the particle representation preserves multimodality of the marginal posterior pdfs. For scenario S1, we plot in Fig. 6 example realizations of the two particle sets employed by PART-EX to represent the marginal posterior pdfs $f(\mathbf{x}_k^{(1)}|\mathbf{z}_{1:k})$ and $f(\mathbf{x}_k^{(2)}|\mathbf{z}_{1:k})$ at time $k = 220$, as well as the corresponding state estimates $\hat{\mathbf{x}}_{220}^{(1)}$ and $\hat{\mathbf{x}}_{220}^{(2)}$. It can be seen that the particle sets express a bimodality of the marginal posterior pdfs. Furthermore, a comparison of the spatial separation between the state estimates $\hat{\mathbf{x}}_{220}^{(1)}$ and $\hat{\mathbf{x}}_{220}^{(2)}$ with the true separation of the targets (which is 10m) suggests that the track coalescence effect is not significant. We note that for a reduction of track coalescence, the particle representation of the marginal posterior pdfs is especially beneficial in PART-EX, whereas in PART-BP, the overconfident nature of the BP messages is predominant.

C. ANALYSIS OF TRACK COALESCENCE AND REPULSION FOR THE JPDA, JPDA*, SJPD, MHT, AND BP METHODS

Next, we analyze the track coalescence and repulsion effects potentially exhibited by the JPDA, JPDA*, SJPD, (track-oriented) MHT, and BP methods for our four scenarios S1 through S4. The number of targets and the target birth times and positions are now unknown to all methods. JPDA, JPDA*, and SJPD use gating with a gate validation threshold of 13.82, corresponding to an in-gate probability (i.e., the probability that a target-generated measurement is within the gate) of 0.999. We remark that a gate validation threshold of 9.21, corresponding to an in-gate probability of 0.99, led to similar results. MHT is a reference implementation provided by *Systems & Technology Research*, Woburn, MA, USA; this implementation uses a Gaussian representation of the individual tracks [52] and a hypothesis depth of five scans. The track confirmation logics [1] of JPDA, JPDA*, and SJPD are set to 12/24, those of MHT to 8/16 (these values were chosen such that each method achieves its best performance across all scenarios). Finally, BP generates a new track for each




FIGURE 5. Performance of PART-BP, GAUSS-BP, PART-EX, and GAUSS-EX for scenario S1: (a) D-Tracks, (b) D-Center, and (c) GOSPA.




FIGURE 6. Example realizations of particle sets representing the marginal posterior pdfs (a) $f(x_k^{(1)}|z_{1:k})$ and (b) $f(x_k^{(2)}|z_{1:k})$ in PART-EX at time $k = 220$, for scenario S1. The black dots indicate the corresponding state estimates $\hat{x}_{220}^{(1)}$ and $\hat{x}_{220}^{(2)}$. The dashed straight line corresponds to $x_2 = (\hat{x}_{220}^{(1)} + \hat{x}_{220}^{(2)})/2$.

measurement, sets the corresponding existence probability to 10^{-4} , and prunes existing tracks with existence probabilities below 10^{-4} .

Fig. 7 shows the D-Tracks, D-Center, and GOSPA curves for scenario S1. One can see in Fig. 7(a) that when the two targets move in close proximity, D-Tracks is increased for MHT and decreased for JPDA compared to the ground truth tracks; this indicates track repulsion and track coalescence, respectively. The low D-Tracks curve of JPDA for $k = 200, \dots, 240$ demonstrates that, after they coalesced, the tracks estimated by JPDA separate again only with a delay. In JPDA*, SJPDA, and BP, track coalescence is strongly reduced. In the case of JPDA* and SJPDA, this is due to the special design of these methods; in the case of BP, it is due to the particle representation of the marginal posterior state pdfs and the overconfident nature of the DA-related messages in BP as argued in Section VII-B and verified experimentally in Section VIII-B. We note that JPDA*, as BP, performs approximate averaging of the DA vectors, which leads to overconfident marginal posterior DA pmfs [45].

Furthermore, in Fig. 7(b), the D-Center curves of JPDA, JPDA*, SJPDA, and BP are generally close to the true curves. An exception is JPDA for $k = 200, \dots, 240$, which again reflects the fact that, after coalescence, the tracks estimated by

JPDA separate again only with a delay. The GOSPA curves shown in Fig. 7(c) confirm the results of Figs. 7(a) and (b). In particular, the increased GOSPA error of MHT for $k = 100, \dots, 200$ is due to track repulsion, and that of JPDA for $k = 200, \dots, 240$ to track coalescence. JPDA*, SJPDA, and BP perform similarly well, with a slight performance advantage for BP.

Still for scenario S1, Fig. 8(b)–(d) show the individual GOSPA error components, i.e., the localization error, the error due to missed targets (“missed error”), and the error due to false targets (“false error”) [74]. In addition, the total GOSPA curves from Fig. 7(c) are replicated in Fig. 8(a) for easy reference. The localization error curves in Fig. 8(b) are seen to be similar to the total GOSPA curves; this reflects the fact that missed and false targets contribute only little to the total GOSPA error.

Fig. 9 shows the GOSPA error and its components for scenario S2. In S2, differently from S1, the two targets are in close proximity right from the beginning (see Fig. 4(b)). This poses a challenge for the track initiation stage of the JPDA filter methods. Indeed, JPDA, JPDA*, and SJPDA generate only one track as long as the targets remain in close proximity. This means that one of the targets is missed, resulting in a large missed error component of JPDA, JPDA*, and SJPDA in the time range $k = 1, \dots, 200$, as shown in Fig. 9(c). By contrast, as also shown in Fig. 9(c), MHT and BP initialize both tracks correctly. Fig. 9(b) indicates an increased localization error of MHT for $k = 1, \dots, 150$; this is due to track repulsion as in S1. It is furthermore seen that the localization error of JPDA, JPDA*, and SJPDA is reduced; this can be explained by the fact that these methods estimate just a single track whereas the other methods estimate both tracks, combined with the fact that GOSPA is an “unnormalized” error metric, which implies that tracking fewer targets generally results in a lower localization error. Note, however, that the total GOSPA error is lowest for BP. Finally, Fig. 9(d) shows that the false error component is small for all methods.

Fig. 10 presents the GOSPA error and its components for scenario S3. Similarly to scenarios S1 and S2, it can be seen that when the two targets are close to each other, i.e., for $k = 125, \dots, 175$, MHT suffers from track repulsion, and




FIGURE 7. Performance of JPDA, JPDA*, SJPD, MHT, and BP for scenario S1: (a) D-Tracks, (b) D-Center, and (c) GOSPA.

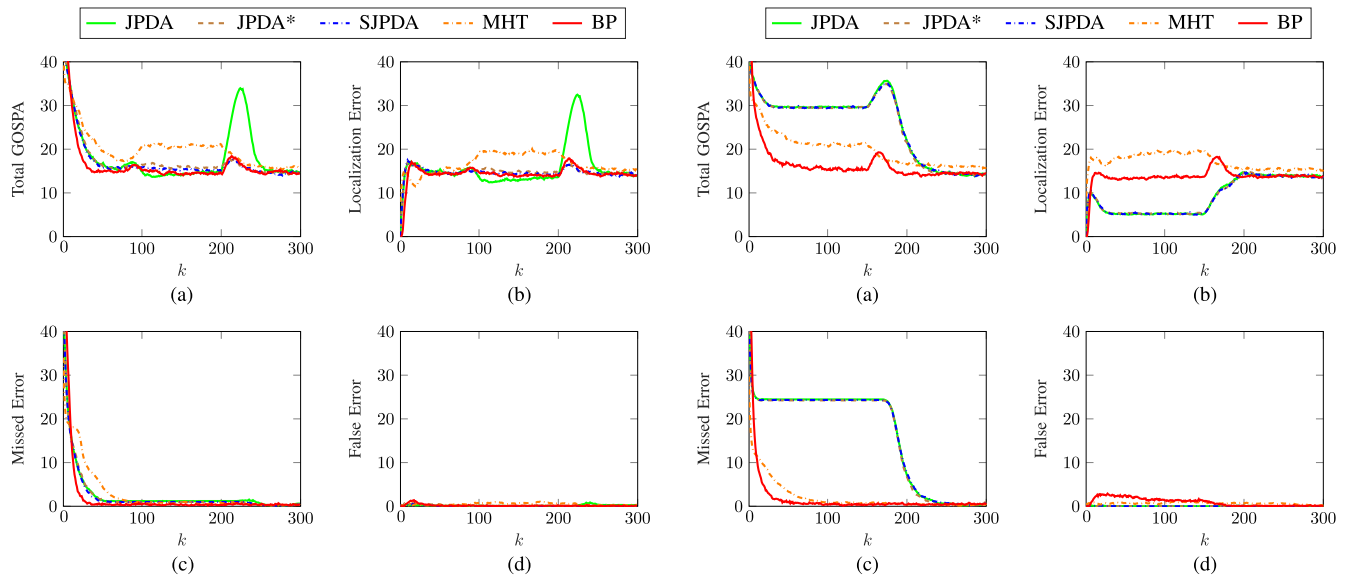


FIGURE 8. GOSPA error decomposition for scenario S1: (a) Total GOSPA error (same as Fig. 7(c)), (b) localization error component, (c) missed error component, and (d) false error component.

after the targets separate again, i.e., for $k = 175, \dots, 225$, JPDA suffers from track coalescence. By contrast, track repulsion and track coalescence are nonexistent or significantly reduced in JPDA*, SJPD, and BP.

D. FURTHER ANALYSIS OF TRACK COALESCENCE FOR THE JPDA, JPDA*, SJPD, AND BP METHODS

The results reported so far for scenarios S1 through S3 demonstrate that JPDA, JPDA*, SJPD, and BP do not exhibit track repulsion and MHT does not exhibit track coalescence. In addition, JPDA*, SJPD, and BP exhibit reduced track coalescence compared to conventional JPDA. Since JPDA*, SJPD, and BP performed equally well, we now further analyze the track coalescence behavior of these methods and, for comparison, of conventional JPDA in the even more challenging scenario S4. Since MHT is not susceptible to




FIGURE 9. Total GOSPA error (a) and its components (b)–(d) for scenario S2.

track coalescence, it is not considered in this analysis. S4 features six targets, which come in close proximity around time $k = 150$ (see Fig. 4(d)). We also consider more challenging variants of S4 with seven, eight, and nine targets. The GOSPA error for these scenarios is shown in Fig. 11. It can be seen that all four methods exhibit only small track coalescence effects for the cases of six, seven, and eight targets. For the case of nine targets, JPDA and JPDA* exhibit an increased GOSPA error compared to BP. In particular, the GOSPA error of JPDA remains high for a long time after the targets separate again.

Still considering scenario S4 and its variants, we report in Table 1 the minimum, median, mean, and maximum total (for all 300 time steps) filter runtimes of Matlab implementations of JPDA, JPDA*, SJPD, and BP on an Intel Xeon Gold 5222 CPU averaged over 1000 simulation runs. Our JPDA, JPDA*, SJPD, and BP implementations use gating and clustering to reduce computational complexity. However, in Scenario S4 and its variants, this has only a limited effect during the time




FIGURE 10. Total GOSPA error (a) and its components (b)–(d) for scenario S3.




FIGURE 11. GOSPA error for scenario S4 (a) and for three variants of S4 with seven, eight, and nine targets (b)–(d).

the targets are in close proximity because the targets are within the same gate and thus clustered together. As Table 1 shows, the runtimes of JPDA, JPDA*, and SJPDA increase rapidly with the number of targets. This is consistent with the exponential scaling of the complexity of these filters in the number of tracks; note that the number of tracks may be higher than the number of targets because of false tracks generated due to clutter measurements. Because of the exponential scaling, JPDA and JPDA* are infeasible for more than nine targets and SJPDA is infeasible for more than eight targets. Table 1

TABLE 1. Minimum, Median, Mean, and Maximum Runtimes of JPDA, JPDA*, SJPDA, and BP for Scenario S4 (Six Targets) and its Variants With Seven, Eight, and Nine Targets

		6 Targets	7 Targets	8 Targets	9 Targets
JPDA	min	0.54 s	0.90 s	2.05 s	18.76 s
	median	1.42 s	4.54 s	17.45 s	84.14 s
	mean	1.44 s	4.59 s	33.28 s	2952.90 s
	max	3.43 s	18.34 s	2758.00 s	9458.90 s
JPDA*	min	0.52 s	0.90 s	2.42 s	16.31 s
	median	1.22 s	4.89 s	34.66 s	256.08 s
	mean	1.29 s	5.35 s	39.70 s	309.18 s
	max	3.62 s	22.26 s	340.30 s	1653.30 s
SJPDA	min	1.28 s	16.21 s	132.57 s	
	median	10.30 s	70.63 s	531.93 s	
	mean	10.46 s	74.90 s	640.120 s	
	max	36.73 s	251.72 s	4351.40 s	
BP	min	18.62 s	19.33 s	21.03 s	21.76 s
	median	24.04 s	24.84 s	27.78 s	31.18 s
	mean	24.05 s	24.86 s	27.86 s	31.26 s
	max	28.44 s	29.88 s	31.80 s	36.48 s

also shows that for JPDA, JPDA*, and SJPDA, there is a large difference between the minimum and maximum filter runtimes. The high filter runtimes in large scenarios are due to simulation runs where the number of tracks is considerably higher than the actual number of targets. Furthermore, because there are significantly more simulation runs with a low filter runtime than with a high filter runtime, the median runtimes differ significantly from the mean runtimes. On the other hand, the runtime of BP increases only slowly with the number of targets. This is consistent with the fact that the complexity of BP scales only linearly in the number of targets.

IX. CONCLUSION

We reviewed and analyzed three major methodologies for multitarget tracking: the classical joint probabilistic data association (JPDA) filter and multiple hypothesis tracking (MHT) methodologies and the recently introduced belief propagation (BP) methodology. The focus of our study was on track coalescence and track repulsion effects, which are well known to compromise the performance of, respectively, JPDA filter and MHT methods when targets are in close proximity. In particular, we investigated the potential occurrence of track coalescence and track repulsion effects in the BP method. We argued that the BP method does not suffer from track repulsion because it performs soft data association similarly to the JPDA filter. Moreover, track coalescence effects in the BP method are significantly smaller than in the JPDA filter, because certain properties of the BP messages related to data association encourage separation of target state estimates. These theoretical arguments were confirmed by simulation experiments. Our numerical results demonstrated excellent performance of the BP method compared to the JPDA, JPDA*, SJPDA, and MHT methods in scenarios with

targets in close proximity. Future work will include an analysis of the quality of the BP state estimates in terms of the normalized estimation error squared (NEES) [76].

DEDICATION AND ACKNOWLEDGMENT

This article is dedicated to the memory of Dr. Craig A. Carthel, a brilliant applied mathematician and software architect who passed away in July 2022. Dr. Carthel contributed to this work while employed with Systems & Technology Research, Woburn, MA, USA.

REFERENCES

- [1] Y. Bar-Shalom, P. K. Willett, and X. Tian, *Tracking and Data Fusion: A Handbook of Algorithms*. Storrs, CT, USA: Yaakov Bar-Shalom, 2011.
- [2] D. B. Reid, "An algorithm for tracking multiple targets," *IEEE Trans. Autom. Control*, vol. AC-24, no. 6, pp. 843–854, Dec. 1979.
- [3] S. Challa, M. R. Morelande, D. Mušicki, and R. J. Evans, *Fundamentals of Object Tracking*. Cambridge, U.K.: Cambridge Univ. Press, 2011.
- [4] R. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Norwood, MA, USA: Artech House, 2007.
- [5] W. Koch, *Tracking and Sensor Data Fusion: Methodological Framework and Selected Applications*. Berlin, Germany: Springer, 2014.
- [6] F. Meyer et al., "Message passing algorithms for scalable multitarget tracking," *Proc. IEEE*, vol. 106, no. 2, pp. 221–259, Feb. 2018.
- [7] Y. Bar-Shalom, "Extension of the probabilistic data association filter to multi-target tracking," in *Proc. 5th Symp. Nonlinear Estimation*, Sep. 1974, pp. 16–21.
- [8] Y. Bar-Shalom, F. Daum, and J. Huang, "The probabilistic data association filter," *IEEE Control Syst. Mag.*, vol. 29, no. 6, pp. 82–100, Dec. 2009.
- [9] T. Fortmann, Y. Bar-Shalom, and M. Scheffe, "Sonar tracking of multiple targets using joint probabilistic data association," *IEEE J. Ocean. Eng.*, vol. JOE-8, no. 3, pp. 173–184, Jul. 1983.
- [10] G. Ferri et al., "Cooperative robotic networks for underwater surveillance: An overview," *IET Radar Sonar Navig.*, vol. 11, no. 12, pp. 1740–1761, Dec. 2017.
- [11] D. Y. Kim, B. Ristic, R. Guan, and L. Rosenberg, "A Bernoulli track-before-detect filter for interacting targets in maritime radar," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 57, no. 3, pp. 1981–1991, Jun. 2021.
- [12] D. Gaglione et al., "Bayesian information fusion and multitarget tracking for maritime situational awareness," *IET Radar, Sonar Navig.*, vol. 14, no. 12, pp. 1845–1857, Oct. 2020.
- [13] J. Levinson et al., "Towards fully autonomous driving: Systems and algorithms," in *Proc. 2011 IEEE Intell. Veh. Symp.*, Jun. 2011, pp. 163–168.
- [14] C. Urmson et al., "Autonomous driving in urban environments: Boss and the urban challenge," *J. Field Robot., Special Issue 2007 DARPA Urban Challenge, Part I*, vol. 25, no. 8, pp. 425–466, Jun. 2008.
- [15] F. Meyer and J. L. Williams, "Scalable detection and tracking of geometric extended objects," *IEEE Trans. Signal Process.*, vol. 69, pp. 6283–6298, 2021.
- [16] Y. Bar-Shalom, *Multitarget-Multisensor Tracking: Advanced Applications*. Norwood, MA, USA: Artech House, 1990.
- [17] D. Musicki and R. Evans, "Joint integrated probabilistic data association: JIPDA," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 40, no. 3, pp. 1093–1099, Jul. 2004.
- [18] J. Vermaak, S. Maskell, and M. Briers, "A unifying framework for multi-target tracking and existence," in *Proc. 2005 7th Int. Conf. Inf. Fusion*, Jul. 2005, pp. 250–258.
- [19] D. Musicki and B. La Scala, "Multi-target tracking in clutter without measurement assignment," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 44, no. 3, pp. 877–896, Jul. 2008.
- [20] P. Horridge and S. Maskell, "Real-time tracking of hundreds of targets with efficient exact JPDAF implementation," in *Proc. 2006 9th Int. Conf. Inf. Fusion*, Jul. 2006, pp. 1–8.
- [21] K. Pattipati, R. L. Popp, and T. Kirubarajan, "Survey of assignment techniques for multitarget tracking," in *Multitarget-Multisensor Tracking: Applications and Advances*, vol. 3, Y. Bar-Shalom and W. D. Blair, Eds. Norwood, MA, USA: Artech House, 2000, ch. 2, pp. 77–159.
- [22] I. J. Cox and S. L. Hingorani, "An efficient implementation of Reid's multiple hypothesis tracking algorithm and its evaluation for the purpose of visual tracking," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 18, no. 2, pp. 138–150, Feb. 1996.
- [23] R. Danchick and G. E. Newnam, "Reformulating Reid's MHT method with generalised Murty K-best ranked linear assignment algorithm," *IEEE Radar Sonar Navigation*, vol. 153, no. 1, pp. 13–22, Feb. 2006.
- [24] S. S. Blackman, "Multiple hypothesis tracking for multiple target tracking," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 19, no. 1, pp. 5–18, Jan. 2004.
- [25] C. Morefield, "Application of 0-1 integer programming to multitarget tracking problems," *IEEE Trans. Autom. Control*, vol. AC-22, no. 3, pp. 302–312, Jun. 1977.
- [26] S. Deb, M. Yeddanapudi, K. Pattipati, and Y. Bar-Shalom, "A generalized S-D assignment algorithm for multisensor-multitarget state estimation," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 33, no. 2, pp. 523–538, Apr. 1997.
- [27] W. Zhang and F. Meyer, "Multisensor multiobject tracking with improved sampling efficiency," *IEEE Trans. Signal Process.*, vol. 72, pp. 2036–2053, 2024.
- [28] M. Liang and F. Meyer, "Neural enhanced belief propagation for multi-object tracking," *IEEE Trans. Signal Process.*, vol. 72, pp. 15–30, 2024.
- [29] R. P. S. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, no. 4, pp. 1152–1178, Oct. 2003.
- [30] B.-N. Vo, S. Singh, and A. Doucet, "Sequential Monte Carlo methods for multitarget filtering with random finite sets," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 41, no. 4, pp. 1224–1245, Oct. 2005.
- [31] B.-T. Vo, B.-N. Vo, and A. Cantoni, "Analytic implementations of the cardinalized probability hypothesis density filter," *IEEE Trans. Signal Process.*, vol. 55, no. 7, pp. 3553–3567, Jul. 2007.
- [32] J. L. 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.
- [33] S. Nannuru, S. Blouin, M. Coates, and M. Rabbat, "Multisensor CPHD filter," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 52, no. 4, pp. 1834–1854, Aug. 2016.
- [34] Á. F. García-Fernández, L. Svensson, and M. R. Morelande, "Multiple target tracking based on sets of trajectories," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 56, no. 3, pp. 1685–1707, Jun. 2020.
- [35] K. Granström, L. Svensson, Y. Xia, J. Williams, and Á. F. García-Fernández, "Poisson multi-Bernoulli mixture trackers: Continuity through random finite sets of trajectories," in *Proc. 2018 21st Int. Conf. Inf. Fusion*, Jul. 2018, pp. 973–981.
- [36] T. Kropfreiter, F. Meyer, and F. Hlawatsch, "A fast labeled multi-Bernoulli filter using belief propagation," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 56, no. 3, pp. 2478–2488, Jun. 2020.
- [37] T. Kropfreiter, F. Meyer, and F. Hlawatsch, "An efficient labeled/unlabeled random finite set algorithm for multiobject tracking," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 58, no. 6, pp. 5256–5275, Dec. 2022.
- [38] D. P. Bertsekas, *Linear Network Optimization: Algorithms and Codes*. Cambridge, MA, USA: MIT Press, 1991.
- [39] D. Salmond, "Mixture reduction algorithms for target tracking in clutter," in *Proc. SPIE-90*, Apr. 1990, pp. 434–445.
- [40] L. Y. Pao, "Multisensor multitarget mixture reduction algorithms for tracking," *J. Guid. Control Dyn.*, vol. 17, no. 6, pp. 1205–1211, Nov. 1994.
- [41] J. Vermaak, S. J. Godsill, and P. Perez, "Monte Carlo filtering for multi target tracking and data association," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 41, no. 1, pp. 309–332, Jan. 2005.
- [42] F. Meyer, P. Braca, P. Willett, and F. Hlawatsch, "A scalable algorithm for tracking an unknown number of targets using multiple sensors," *IEEE Trans. Signal Process.*, vol. 65, no. 13, pp. 3478–3493, Jul. 2017.
- [43] R. J. Fitzgerald, "Track biases and coalescence with probabilistic data association," *IEEE Trans. Signal Process.*, vol. SP-21, no. 6, pp. 822–825, Nov. 1985.
- [44] S. Coraluppi, C. Carthel, P. Willett, M. Dingboe, O. O'Neill, and T. Luginbuhl, "The track repulsion effect in automatic tracking," in *Proc. 2009 12th Int. Conf. Inf. Fusion*, Jul. 2009, pp. 2225–2230.
- [45] H. A. P. Blom and E. A. Bloem, "Probabilistic data association avoiding track coalescence," *IEEE Trans. Autom. Control*, vol. 45, no. 2, pp. 247–259, Feb. 2000.

- [46] H. A. P. Blom, E. A. Bloem, and D. Musicki, "JIPDA: Automatic target tracking avoiding track coalescence," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 51, no. 2, pp. 962–974, Apr. 2015.
- [47] S. Coraluppi and C. Carthel, "An equivalence-class approach to multiple-hypothesis tracking," in *Proc. 2012 IEEE Aerosp. Conf.*, Mar. 2012, pp. 1–8.
- [48] L. Svensson, D. Svensson, M. Guerriero, and P. Willett, "Set JPDA filter for multitarget tracking," *IEEE Trans. Signal Process.*, vol. 59, no. 10, pp. 4677–4691, Oct. 2011.
- [49] J. L. Williams, "An efficient, variational approximation of the best fitting multi-Bernoulli filter," *IEEE Trans. Signal Process.*, vol. 63, no. 1, pp. 258–273, Jan. 2015.
- [50] R. A. Lau and J. L. Williams, "A structured mean field approach for existence-based multiple target tracking," in *Proc. 2016 19th Int. Conf. Inf. Fusion*, Jul. 2016, pp. 1111–1118.
- [51] S. Reuter, B.-T. Vo, B.-N. Vo, and K. Dietmayer, "The labeled multi-Bernoulli filter," *IEEE Trans. Signal Process.*, vol. 62, no. 12, pp. 3246–3260, Jun. 2014.
- [52] S. Coraluppi and C. Carthel, "Track management in multiple-hypothesis tracking," in *Proc. 2018 IEEE 10th Sensor Array Multichannel Signal Process. Workshop*, Jul. 2018, pp. 3840–3843.
- [53] D. F. Crouse, P. Willett, M. Guerriero, and L. Svensson, "An approximate minimum MOSPA estimator," in *Proc. 2011 IEEE Int. Conf. Acoust., Speech Signal Process.*, Jul. 2011, pp. 3644–3647.
- [54] J. L. Williams and R. Lau, "Approximate evaluation of marginal association probabilities with belief propagation," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 50, no. 4, pp. 2942–2959, Oct. 2014.
- [55] G. Soldi, F. Meyer, P. Braca, and F. Hlawatsch, "Self-tuning algorithms for multisensor-multitarget tracking using belief propagation," *IEEE Trans. Signal Process.*, vol. 67, no. 15, pp. 3922–3937, Aug. 2019.
- [56] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 498–519, Feb. 2001.
- [57] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*. Cambridge, MA, USA: MIT Press, 2009.
- [58] T. Kropfreiter, F. Meyer, S. Coraluppi, C. Carthel, R. Mendrzik, and P. Willett, "Track coalescence and repulsion: MHT, JPDA, and BP," in *Proc. 2021 24th Int. Conf. Inf. Fusion*, Nov. 2021, pp. 1–8.
- [59] S. Mori, C.-Y. Chong, and K. C. Chang, "Evaluation of data association hypotheses: Non-Poisson I.I.D. cases," in *Proc. 7th Int. Conf. Inf. Fusion*, Jun. 2004, pp. 1133–1140.
- [60] K. R. Pattipati, S. Deb, Y. Bar-Shalom, and R. B. J. Washburn, "A new relaxation algorithm and passive sensor data association," *IEEE Trans. Autom. Control*, vol. 37, no. 2, pp. 198–213, Feb. 1992.
- [61] A. P. 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.
- [62] A. B. Poore and S. Gadaleta, "Some assignment problems arising from multiple target tracking," *Math. Comput. Model.*, vol. 43, no. 9/10, pp. 1074–1091, 2006.
- [63] S. Coraluppi, D. Grimmert, and P. de Theije, "Benchmark evaluation of multistatic trackers," in *Proc. 2006 9th Int. Conf. Inf. Fusion*, Jul. 2006, pp. 1–7.
- [64] P. Willett, T. Luginbuhl, and E. Giannopoulos, "MHT tracking for crossing sonar targets," in *Proc. SPIE*, Sep. 2007, pp. 469–480.
- [65] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," *IEEE Trans. Signal Process.*, vol. 56, no. 8, pp. 3447–3457, Aug. 2008.
- [66] D. F. Crouse, "Advances in displaying uncertain estimates of multiple targets," in *Proc. SPIE*, May 2013, vol. 8745, pp. 15–45.
- [67] D. F. Crouse, P. Willett, and Y. Bar-Shalom, "Developing a real-time track display that operators do not hate," *IEEE Trans. Signal Process.*, vol. 59, no. 7, pp. 3441–3447, Jul. 2011.
- [68] R. Fitzgerald, "Development of practical PDA logic for multitarget tracking by microprocessor," in *Multitarget-Multisensor Tracking: Advanced Applications*, Y. Bar-Shalom, Ed. Norwood, MA, USA: Artech House, 1990.
- [69] F. Meyer and K. L. Gemba, "Probabilistic focalization for shallow water localization," *J. Acoust. Soc. Amer.*, vol. 150, no. 2, pp. 1057–1066, Aug. 2021.
- [70] D. Gaglione, P. Braca, G. Soldi, F. Meyer, F. Hlawatsch, and M. Z. Win, "Fusion of sensor measurements and target-provided information in multitarget tracking," *IEEE Trans. Signal Process.*, vol. 70, pp. 322–336, 2022.
- [71] Y. Weiss and W. T. Freeman, "Correctness of belief propagation in Gaussian graphical models of arbitrary topology," *Neural Comput.*, vol. 13, no. 10, pp. 2173–2200, 2001.
- [72] M. Wainwright and M. Jordan, "Graphical models, exponential families, and variational inference," *Foundations Trends Mach. Learn.*, vol. 1, no. 1/2, pp. 1–305, 2008.
- [73] Y. Bar-Shalom, T. Kirubarajan, and X.-R. Li, *Estimation With Applications to Tracking and Navigation*. New York, NY, USA: Wiley, 2002.
- [74] A. S. Rahmathullah, Á. F. García-Fernández, and L. Svensson, "Generalized optimal sub-pattern assignment metric," in *Proc. 2017 20th Int. Conf. Inf. Fusion*, Jul. 2017, pp. 1–8.
- [75] T. Vu, "A complete optimal subpattern assignment (COSPA) metric," in *Proc. 2020 23rd Int. Conf. Inf. Fusion*, Jul. 2020, pp. 1–8.
- [76] Y. Bar-Shalom and X.-R. Li, *Estimation and Tracking: Principles, Techniques, and Software*. Norwood, MA, USA: Artech House, 1998.