IET Journals
The Institution of Engineering and Technology

ISSN 1751-8644
doi: 0000000000
www.ietdl.org

Bayesian Information Fusion and Multitarget Tracking for Maritime Situational Awareness

Domenico Gaglione [1,*], *Giovanni Soldi* [1,*], *Florian Meyer* [2], *Franz Hlawatsch* [3], *Paolo Braca* [1], *Alfonso Farina* [4], *Moe Z. Win* [5]

[1] *NATO Centre for Maritime Research and Experimentation (CMRE), La Spezia, Italy*
[2] *Scripps Institution of Oceanography and Electrical and Computer Engineering Department, University of California San Diego, La Jolla, CA, USA*
[3] *Institute of Telecommunications, TU Wien, Vienna, Austria*
[4] *Professional Consultant, Rome, Italy*
[5] *Laboratory for Information and Decision Systems (LIDS), Massachusetts Institute of Technology (MIT), Cambridge, MA, USA*
[*] *Co-first authors*
** E-mail: domenico.gaglione@cmre.nato.int, giovanni.soldi@cmre.nato.int*

Abstract:
The goal of maritime situational awareness (MSA) is to provide a seamless wide-area operational picture of ship traffic in coastal areas and the oceans in real time. Radar is a central sensing modality for MSA. In particular, oceanographic high-frequency surface-wave (HFSW) radars are attractive for surveying large sea areas at over-the-horizon distances, due to their low environmental footprint and low power requirements. However, their design is not optimal for the challenging conditions prevalent in MSA applications, thus calling for the development of dedicated information fusion and multisensor-multitarget tracking algorithms. In this paper, we show how the multisensor-multitarget tracking problem can be formulated in a Bayesian framework and efficiently solved by running the loopy sum-product algorithm on a suitably devised factor graph. Compared to previously proposed methods, our approach is advantageous in terms of estimation accuracy, computational complexity, implementation flexibility, and scalability. Moreover, its performance can be further enhanced by estimating unknown model parameters in an online fashion and by fusing automatic identification system (AIS) data and context-based information. The effectiveness of the proposed Bayesian multisensor-multitarget tracking and information fusion algorithms is demonstrated through experimental results based on simulated data as well as real HFSW radar data and real AIS data.

1 Introduction and Background

Maritime situational awareness (MSA) aims at providing a seamless wide-area operational picture of ship traffic in coastal areas and the oceans in real time [1–5]. Nowadays multiple heterogeneous sensors and information sources are available for MSA—all with their advantages and limitations—and a significant research effort has been made for their combination. A successful approach, also pursued in this article, is to perform Bayesian information fusion and multitarget tracking to estimate the positions and motion parameters of ships in the coverage area of multiple heterogenous sensors [6]. Within this general framework, we will consider the emerging application of a specific type of long-range radar sensors, and we will develop the use of an efficient and versatile message passing methodology for Bayesian information fusion and multitarget tracking.

1.1 Information Sources for MSA

A major information source for MSA is provided by pulse radar sensors installed along the coastline. However, their coverage area is limited by line-of-sight propagation (even though anomalous propagation increases coverage) [7], and they may impact the environment when they transmit with a high power. Another important information source is the automatic identification system (AIS), which is intended to identify and track all passenger ships and commercial ships whose gross tonnage exceeds 300 tons, using a network of base stations [8]. However, not all ships report their position and motion information. Further information sources include imaging satellites and surveillance cameras.

Recently, emerging long-range sensors such as high-frequency surface-wave (HFSW) radars have been considered for ship localization and tracking [5, 9–12]. Initially introduced for ocean remote sensing, HFSW radars promise to dramatically increase MSA coverage by their ability to detect targets at over-the-horizon (OTH) distances [5, 9–12]. HFSW radars can thus be expected to become an important element of future MSA systems. In particular, multiple HFSW radars [5, 9, 11–13], combined with other data sources such as the AIS, satellite images, and historical records of ship routes, have the potential to provide continuous-time coverage of large sea areas at OTH distances. However, to leverage this potential, several challenges have to be addressed. Because HFSW radars were originally conceived for oceanographic purposes [14, 15], their design is not optimal for the detection and tracking of an unknown number of targets [6, 16, 17]. Most notably, interference and sea/land clutter lead to a poor range-azimuth resolution and a high false-alarm rate [18]. These performance limitations call for the integration and fusion of information from other data sources. Such information fusion is difficult because of nonlinear system models, asynchronous sources, and sensors with imperfectly modeled parameters and time-varying performance.

1.2 Multitarget Tracking and Information Fusion for MSA

The development of future MSA systems combining multiple HFSW radars with other information sources requires dedicated algorithms for multitarget tracking and information fusion. These algorithms will have to cope with the heterogeneity of MSA information sources and with the measurement-origin uncertainty (data association problem) and, possibly, high false-alarm rate affecting HFSW radars.

Existing multitarget tracking algorithms can be classified as "vector-type" and "set-type" algorithms. Vector-type algorithms describe the target states and measurements by random vectors. In the joint probabilistic data association (JPDA) filter [16, Sec. 6.4], the target-measurement association variables are marginalized out under the constraint that each measurement is related to at most one target. Then, a Gaussian probability density function (pdf) is fitted to the posterior pdf associated with each target. The JPDA filter assumes that the number of targets is known in advance. This limitation is removed in the joint integrated probabilistic data association (JIPDA) filter [19, 20], which models the existence of each target by a binary Markov chain. New targets tracks are initialized using the measurements that are not in a neighborhood of an existing track. Multiple hypothesis tracking (MHT) methods [21] search for the most likely target-measurement associations over a moving window of consecutive time steps. An expanding tree of association hypotheses is built, with each leaf corresponding to a partitioning into subsets of all the measurements acquired until the present time step and believed to correspond to the same target. The high computational complexity of MHT algorithms can be significantly reduced by considering only the most likely hypotheses [22–24].

Set-type algorithms describe the target states and measurements by random finite sets (RFSs) [17]. This approach is suited to modeling target appearance and disappearance in a Bayesian framework, and to handling complex, hybrid continuous/discrete distributions. The probability hypothesis density (PHD) filter [17, 25, 26] and the cardinalized PHD (CPHD) filter [17, 27, 28] are popular examples. They both compute the posterior PHD of the target state in a sequential fashion, and the CPHD filter additionally propagates the cardinality distribution of the RFS. A multisensor extension of the PHD filter was proposed in [29, 30] but is computationally infeasible. Computationally feasible multisensor (C)PHD based algorithms include the iterated-corrector (C)PHD filter [31, 32] and the partition-based (C)PHD filter [33]. Multi-Bernoulli (MB) filters approximate the posterior multitarget state RFS by an MB RFS [17]. Two variants are the track-oriented marginal MB/Poisson (TOMB/P) filter and the measurement-oriented marginal MB/Poisson (MOMB/P) filter [34, 35]. These filters approximate the posterior multitarget state RFS by a mixture of MB RFSs, where each MB RFS component corresponds to one of the global association hypotheses in the MHT methods, and they also model target appearance and disappearance. A particle-based implementation of the TOMB/P filter was presented in [36]. A slightly modified version of that algorithm that does not use the prediction step has been applied to static source localization in [37]. Labeled RFS-based multitarget tracking methods, such as the δ-generalized labeled MB (δ-GLMB) filter [38] and the labeled MB (LMB) filter [39], use labels to explicitly maintain track continuity. The elements of a labeled RFS are the target state vectors augmented by distinct labels identifying the respective targets.

Algorithms of either type, i.e, vector-type and set-type, have been developed and evaluated using datasets from HFSW radars [11, 12]. However, these algorithms have several limitations: they do not permit an easy integration of additional information sources besides radar measurements [40], they do not adapt to time-varying model parameters, and their complexity usually does not scale well in relevant system parameters.

Recently, a new approach to the development of multitarget tracking and information fusion algorithms has been provided by a methodology that is based on factor graphs and the *sum-product algorithm* (SPA) [6, 41]. Factor graphs constitute a graphical representation of the underlying statistical model. The SPA relies on a factor graph and yields a principled and intuitive approximation to optimal Bayesian inference [42, 43]. An early SPA-based algorithm involving so-called mutual exclusion constraints for addressing the data association problem [44] was shown in [45, 46] to be outperformed by a bipartite SPA-based algorithm for probabilistic data association [46–48]. The latter algorithm, discussed in Section 2.4, has been successfully applied in set-type multitarget methods [6, 34], in vector-type multisensor-multitarget methods [6, 41, 49], and in vector-type methods for indoor localization [50, 51]. The underlying factor graph has also been used for group tracking [52] and multitarget tracking using batch filtering [53] (both based on the expectation-maximization algorithm), and for multitarget multipath tracking [54, 55].

An important feature of the SPA is its ability to exploit conditional independence properties of the involved random variables for a drastic reduction of complexity. Thereby, SPA-based inference algorithms can achieve an attractive performance-complexity compromise (see [6] and references therein), In addition, the SPA approach has significant advantages regarding accuracy, flexibility, and scalability as analyzed in [6, 41]; it is able to implicitly estimate unknown and time-varying model parameters; and it supports nonlinear and/or non-Gaussian system models. The SPA's high versatility and intuitiveness enable the establishment of a suite of Bayesian multisensor-multitarget tracking and information fusion algorithms where, similar to a construction kit system, algorithm parts can be combined, extended, or adapted to achieve desired functionalities and properties. In particular, existing SPA-based multisensor-multitarget tracking algorithms can be systematically extended to integrate additional data (e.g., AIS data) or to adapt to unknown time-varying model parameters (e.g., the detection probabilities of radar sensors or an index choosing from multiple target motion models). These features cause the SPA methodology for multitarget tracking and information fusion to be attractive for MSA applications, with a high potential for increasing the coverage, accuracy, and reliability of future MSA systems.

1.3 Contribution and Paper Organization

This paper unifies under a common SPA-based framework the work presented in [6, 41, 49, 56, 57], providing guidelines for the construction of Bayesian multisensor-multitarget tracking algorithms with inherent information fusion. We describe how the multisensor-multitarget tracking problem can be formulated in a Bayesian framework, represented by a factor graph, and efficiently solved by means of the SPA. We also show how SPA-based tracking algorithms can cope with challenges inherent to maritime radar, namely, adaptation to changing dynamic models [49] and incorporation of non-radar data such as AIS data and contextual information [56]. Finally, we demonstrate the effectiveness of the SPA approach for MSA through experimental results based on simulated data as well as real HFSW maritime radar data and real AIS data.

The paper is organized as follows. In Section 2, we review factor graphs and the SPA in the context of Bayesian probabilistic inference and probabilistic data association. In Section 3, we describe how factor graphs and the SPA can be used to develop an algorithm that tracks an unknown number of targets and fuses data from multiple radar sensors. Extensions of our factor graph/SPA framework for multitarget tracking to support multiple dynamic models and to fuse radar measurements with real-time AIS data are presented in Section 4 and Section 5, respectively. In Section 6, we describe how the knowledge of common sea lanes extracted from historical AIS data can be used as additional contextual information in SPA-based tracking algorithms. Section 7 and Section 8 present the results of numerical experiments based on, respectively, simulated and real radar sensor and AIS data. Section 9 concludes the paper.

2 Bayesian Inference Using The Sum-Product Algorithm

In this section, we review factor graphs and the SPA in the context of Bayesian probabilistic inference and probabilistic data association.

2.1 Factor Graphs and the SPA

Consider the estimation of K parameter vectors $\boldsymbol{x}_k$, $k \in \{1, 2, \ldots, K\}$ from an observed measurement vector $\boldsymbol{z}$. We assume a Bayesian framework, so that the estimation of $\boldsymbol{x}_k$ relies on the posterior pdf $f(\boldsymbol{x}_k|\boldsymbol{z})$ [58]. The K posterior pdfs $f(\boldsymbol{x}_k|\boldsymbol{z})$ are marginal densities of the joint posterior pdf $f(\boldsymbol{x}|\boldsymbol{z})$, where $\boldsymbol{x} = \left[\boldsymbol{x}_1^{\mathrm{T}} \cdots \boldsymbol{x}_K^{\mathrm{T}}\right]^{\mathrm{T}}$; these marginal densities can be computed efficiently

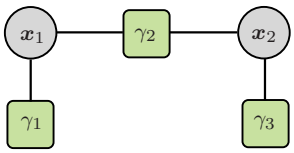

Fig. 1: Factor graph describing the factorization $f(\boldsymbol{x}|\boldsymbol{z}) \propto \gamma_1(\boldsymbol{x}_1)\gamma_2(\boldsymbol{x}_1,\boldsymbol{x}_2)\gamma_3(\boldsymbol{x}_2)$. Factor nodes and variable nodes are represented by squares and circles, respectively.

by means of the SPA. The underlying assumption is that

$$f(\boldsymbol{x}|\boldsymbol{z}) \propto \prod_{q=1}^{Q} \gamma_q\big(\boldsymbol{x}^{(q)}\big), \tag{1}$$

where $\propto$ denotes equality up to a constant normalization factor, Q is the number of factors, and each argument $\boldsymbol{x}^{(q)}$ comprises certain $\boldsymbol{x}_k$. Note that each $\boldsymbol{x}_k$ may be contained in several $\boldsymbol{x}^{(q)}$, and the factors $\gamma_q\big(\boldsymbol{x}^{(q)}\big)$ may also depend on $\boldsymbol{z}$. According to (1), the marginalization of $f(\boldsymbol{x}|\boldsymbol{z})$ amounts to the calculation of a *sum* (or integral) of a *product* of functions. A straightforward but powerful graphical representation of the factorization structure (1) is provided by a factor graph [42, 43]. Each factor $\gamma_q(\cdot)$ corresponds to a factor node in the factor graph and each variable $\boldsymbol{x}_k$ to a variable node. Furthermore, a factor node "γ_q" and a variable node "$\boldsymbol{x}_k$" are connected by an edge (i.e., they are *adjacent*) if $\boldsymbol{x}_k$ is part of the argument $\boldsymbol{x}^{(q)}$ of $\gamma_q(\cdot)$. Fig. 1 shows the factor graph for a simple example, where $\boldsymbol{x} = \big[\boldsymbol{x}_1^{\mathrm{T}}\ \boldsymbol{x}_2^{\mathrm{T}}\big]^{\mathrm{T}}$ and the joint posterior pdf factorizes as $f(\boldsymbol{x}|\boldsymbol{z}) \propto \gamma_1(\boldsymbol{x}_1)\gamma_2(\boldsymbol{x}_1,\boldsymbol{x}_2)\gamma_3(\boldsymbol{x}_2)$.

The SPA is a *message passing* algorithm, where "messages" are calculated for each node of the factor graph and passed to adjacent nodes. A message passed on an edge of the graph is a function of the variable whose associated node is connected to that edge. Let the set $\mathcal{V}_q$ comprise the indices k of all those variables $\boldsymbol{x}_k$ that are part of $\boldsymbol{x}^{(q)}$, or equivalently of all those variable nodes "$\boldsymbol{x}_k$" that are adjacent to factor node "γ_q". Then, the following message is passed from factor node "γ_q" to variable node "$\boldsymbol{x}_k$" with $k \in \mathcal{V}_q$:

$$\zeta_{\gamma_q \to \boldsymbol{x}_k}(\boldsymbol{x}_k) = \int \gamma_q\big(\boldsymbol{x}^{(q)}\big) \prod_{k' \in \mathcal{V}_q \setminus \{k\}} \eta_{\boldsymbol{x}_{k'} \to \gamma_q}(\boldsymbol{x}_{k'})\, \mathrm{d}\boldsymbol{x}_{\bar{k}}\,. \tag{2}$$

Here, $\int \ldots \mathrm{d}\boldsymbol{x}_{\bar{k}}$ is the integral with respect to all $\boldsymbol{x}_{k'}$, $k' \in \mathcal{V}_q$ except $\boldsymbol{x}_k$, and the messages $\eta_{\boldsymbol{x}_{k'} \to \gamma_q}(\boldsymbol{x}_{k'})$ are calculated as described presently. If the factorization (1) involves (also) discrete variables, then the corresponding integrations in (2) are replaced by summations. Furthermore, let the set $\mathcal{F}_k$ comprise the indices q of all those factors $\gamma_q\big(\boldsymbol{x}^{(q)}\big)$ whose argument $\boldsymbol{x}^{(q)}$ involves $\boldsymbol{x}_k$, or equivalently of all those factor nodes "γ_q" that are adjacent to variable node "$\boldsymbol{x}_k$". Then, the following message is passed from variable node "$\boldsymbol{x}_k$" to factor node "γ_q" with $q \in \mathcal{F}_k$:

$$\eta_{\boldsymbol{x}_k \to \gamma_q}(\boldsymbol{x}_k) = \prod_{q' \in \mathcal{F}_k \setminus \{q\}} \zeta_{\gamma_{q'} \to \boldsymbol{x}_k}(\boldsymbol{x}_k)\,. \tag{3}$$

Let us determine these messages for our example of Fig. 1. The message passed from factor node "γ_2" to variable node "$\boldsymbol{x}_2$" is $\zeta_{\gamma_2 \to \boldsymbol{x}_2}(\boldsymbol{x}_2) = \int \gamma_2(\boldsymbol{x}_1,\boldsymbol{x}_2)\, \eta_{\boldsymbol{x}_1 \to \gamma_2}(\boldsymbol{x}_1)\, \mathrm{d}\boldsymbol{x}_1$; note that $\boldsymbol{x}^{(2)} = \big[\boldsymbol{x}_1^{\mathrm{T}}\ \boldsymbol{x}_2^{\mathrm{T}}\big]^{\mathrm{T}}$. Furthermore, the message passed from variable node "$\boldsymbol{x}_1$" to factor node "γ_2" is $\eta_{\boldsymbol{x}_1 \to \gamma_2}(\boldsymbol{x}_1) = \zeta_{\gamma_1 \to \boldsymbol{x}_1}(\boldsymbol{x}_1)$. Message passing starts at variable nodes and/or factor nodes with only one edge; these pass a constant message or the corresponding factor, respectively. In our example of Fig. 1, the message passed from factor node "γ_1" to variable node "$\boldsymbol{x}_1$" is $\zeta_{\gamma_1 \to \boldsymbol{x}_1}(\boldsymbol{x}_1) = \gamma_1(\boldsymbol{x}_1)$.

If the factor graph does not contain loops, i.e., if it is a tree, then at each node, an outgoing message can be calculated as soon as all the required incoming messages are available. When all the messages have been obtained, then a *belief* $\tilde{f}(\boldsymbol{x}_k)$ is calculated for each variable node "$\boldsymbol{x}_k$" by multiplying all the incoming messages and normalizing the resulting function in the sense that $\int \tilde{f}(\boldsymbol{x}_k)\, \mathrm{d}\boldsymbol{x}_k = 1$. If the factor graph does not contain loops, then the belief $\tilde{f}(\boldsymbol{x}_k)$ is exactly equal to the marginal posterior pdf $f(\boldsymbol{x}_k|\boldsymbol{z})$. This is verified for our example in Fig. 1 within a sequential Bayesian estimation setting in Section 2.2. If the factor graph contains loops, then the SPA is usually executed in an iterative manner, in the sense that the calculation of the messages is performed multiple times ("iterative SPA"). The beliefs $\tilde{f}(\boldsymbol{x}_k)$ are calculated only after the final iteration. In the loopy case, the beliefs are approximations of the marginal posterior pdfs $f(\boldsymbol{x}_k|\boldsymbol{z})$. These approximations tend to be overconfident, which is also known as the data incest problem. Different orders (schedules) of message calculation are possible, which may lead to different beliefs in the loopy case. The choice of an order (schedule) introduces a certain flexibility in the design of iterative SPA algorithms. A closed-form implementation of the message and belief calculation rules is possible for linear-Gaussian system models [42]. In the general case, computationally feasible approximate implementations can be based on particle representations [41] or the unscented transform [59].

2.2 SPA Interpretation of Sequential Bayesian Estimation

In this section, we verify for our example in Fig. 1 the general fact that for a factor graph without loops, the belief $\tilde{f}(\boldsymbol{x}_k)$ provided by the SPA equals the marginal posterior pdf $f(\boldsymbol{x}_k|\boldsymbol{z})$. We adopt a sequential Bayesian estimation framework by setting $\boldsymbol{z} = \boldsymbol{z}_{1:2} \triangleq [\boldsymbol{z}_1^{\mathrm{T}}\ \boldsymbol{z}_2^{\mathrm{T}}]^{\mathrm{T}}$, $\gamma_1(\boldsymbol{x}_1) = f(\boldsymbol{x}_1|\boldsymbol{z}_1)$, $\gamma_2(\boldsymbol{x}_1,\boldsymbol{x}_2) = f(\boldsymbol{x}_2|\boldsymbol{x}_1)$, and $\gamma_3(\boldsymbol{x}_2) = f(\boldsymbol{z}_2|\boldsymbol{x}_2)$, so that the factorization of Fig. 1 reads

$$f(\boldsymbol{x}|\boldsymbol{z}_{1:2}) \propto f(\boldsymbol{x}_1|\boldsymbol{z}_1)\, f(\boldsymbol{x}_2|\boldsymbol{x}_1)\, f(\boldsymbol{z}_2|\boldsymbol{x}_2). \tag{4}$$

This factorization is based on the assumptions that given $\boldsymbol{x}_1$, $\boldsymbol{x}_2$ is conditionally independent of $\boldsymbol{z}_1$, and given $\boldsymbol{x}_2$, $\boldsymbol{z}_2$ is conditionally independent of all the other variables. Let us now calculate the marginal posterior pdf $f(\boldsymbol{x}_2|\boldsymbol{z}_{1:2}) = \int f(\boldsymbol{x}|\boldsymbol{z}_{1:2})\, \mathrm{d}\boldsymbol{x}_1$. We obtain

$$\begin{aligned} f(\boldsymbol{x}_2|\boldsymbol{z}_{1:2}) &\propto \int f(\boldsymbol{x}_1|\boldsymbol{z}_1)\, f(\boldsymbol{x}_2|\boldsymbol{x}_1)\, f(\boldsymbol{z}_2|\boldsymbol{x}_2)\, \mathrm{d}\boldsymbol{x}_1 \\ &= f(\boldsymbol{z}_2|\boldsymbol{x}_2) \underbrace{\int f(\boldsymbol{x}_2|\boldsymbol{x}_1)\, f(\boldsymbol{x}_1|\boldsymbol{z}_1)\, \mathrm{d}\boldsymbol{x}_1}_{= f(\boldsymbol{x}_2|\boldsymbol{z}_1)\ \text{(PREDICTION STEP)}} \qquad (5) \\ &= \underbrace{f(\boldsymbol{z}_2|\boldsymbol{x}_2)\, f(\boldsymbol{x}_2|\boldsymbol{z}_1)}_{\text{(MEASUREMENT UPDATE STEP)}}\,. \qquad (6) \end{aligned}$$

Our factorization of the joint posterior pdf $f(\boldsymbol{x}|\boldsymbol{z}_{1:2})$ in (4) corresponds to the sequential Bayesian estimation framework underlying, e.g., the Kalman filter [60]. Indeed, the above calculations correspond to the prediction step (see (5)) and the measurement update step (see (6)) of sequential Bayesian estimation. On the other hand, executing the SPA for the factor graph depicted in Fig. 1, we obtain the belief for $\boldsymbol{x}_2$ as

$$\begin{aligned} \tilde{f}(\boldsymbol{x}_2) &\propto \zeta_{\gamma_3 \to \boldsymbol{x}_2}(\boldsymbol{x}_2)\, \zeta_{\gamma_2 \to \boldsymbol{x}_2}(\boldsymbol{x}_2) \\ &= f(\boldsymbol{z}_2|\boldsymbol{x}_2)\, \zeta_{\gamma_2 \to \boldsymbol{x}_2}(\boldsymbol{x}_2), \end{aligned} \tag{7}$$

where

$$\begin{aligned} \zeta_{\gamma_2 \to \boldsymbol{x}_2}(\boldsymbol{x}_2) &= \int f(\boldsymbol{x}_2|\boldsymbol{x}_1)\, \eta_{\boldsymbol{x}_1 \to \gamma_2}(\boldsymbol{x}_1)\, \mathrm{d}\boldsymbol{x}_1 \\ &= \int f(\boldsymbol{x}_2|\boldsymbol{x}_1)\, f(\boldsymbol{x}_1|\boldsymbol{z}_1) \mathrm{d}\boldsymbol{x}_1. \end{aligned} \tag{8}$$

The expressions (7) and (8) are seen to be consistent with the measurement update step (6) and the prediction step in (5), respectively.

2.3 Factor Graph Representation of Probabilistic Data Association

A challenging issue in multitarget tracking is the measurement origin uncertainty (MOU) problem, i.e., the problem that the association between the targets and the measurements is unknown. Because of the MOU problem, Bayesian multitarget tracking methods include a probabilistic "data association" component. We assume that a target

can generate at most one measurement and a measurement can originate from at most one target [16, 17]; this will be referred to as the "data association assumption." To illustrate the MOU problem, consider two targets moving in close proximity and two measurements. Because of MOU, the association between targets and measurements is unknown: it is possible that the first target generates the first measurement and the second target the second measurement, or the first target generates the second measurement and the second target the first measurement, or a target does not generate any measurement (missed detection), or a measurement does not originate from a target (false alarm, clutter).

Let us generalize and formalize this data association problem. We consider K targets indexed by $k \in \mathcal{K} \triangleq \{1, 2, \ldots, K\}$, at—for now—a single time step. A sensor produces M measurements indexed by $m \in \mathcal{M} \triangleq \{1, 2, \ldots, M\}$, which result from a detection process. The target-measurement associations can then be characterized by the *target-oriented association variables* a_k for $k \in \mathcal{K}$, which are defined as $m \in \mathcal{M}$ if target k generates measurement m and 0 if target k does not generate any measurement. We assume that a priori the a_k for different targets k are statistically independent, and the distribution of a single a_k is described by some probability mass function (pmf) $p_{\mathrm{s}}(a_k)$, $a_k \in \{0, 1, \ldots, M\}$. Defining $\boldsymbol{a} \triangleq [a_1 \cdots a_K]^{\mathrm{T}}$, the joint pmf of all the a_k is then given by

$$p_{\mathrm{j}}(\boldsymbol{a}) \propto \Phi(\boldsymbol{a}) \prod_{k=1}^{K} p_{\mathrm{s}}(a_k), \tag{9}$$

where $\Phi(\boldsymbol{a}) = 0$ if $\boldsymbol{a}$ associates more than one target with a measurement (i.e., if there are $k, k' \in \mathcal{K}$ with $k \neq k'$ such that $a_k = a_{k'} \neq 0$) and $\Phi(\boldsymbol{a}) = 1$ otherwise. For example, for $K = 2$ targets and $M = 2$ measurements, seven associations are possible, i.e., $\Phi(\boldsymbol{a}) = 1$ for $\boldsymbol{a} \in \mathcal{A} \triangleq \{[0\ 0]^{\mathrm{T}}, [1\ 0]^{\mathrm{T}}, [2\ 0]^{\mathrm{T}}, [0\ 1]^{\mathrm{T}}, [0\ 2]^{\mathrm{T}}, [1\ 2]^{\mathrm{T}}, [2\ 1]^{\mathrm{T}}\}$. The other combinations are impossible, i.e., $\Phi(\boldsymbol{a}) = 0$ for $\boldsymbol{a} \in \{[1\ 1]^{\mathrm{T}}, [2\ 2]^{\mathrm{T}}\}$, since they would associate one measurement with two targets simultaneously.

The *marginal association pmf* $p_{\mathrm{m}}(a_k) \triangleq \sum_{\boldsymbol{a}_{\bar{k}}} p_{\mathrm{j}}(\boldsymbol{a})$, with $\boldsymbol{a}_{\bar{k}}$ denoting the vector $\boldsymbol{a}$ without the kth component, is the sum of all the probabilities that target k generated the various measurements $m \in \mathcal{M}$ or no measurement while taking into account the possible presence of other targets $k' \in \mathcal{K} \backslash \{k\}$. The marginal association pmfs are key elements of many multitarget tracking algorithms, and their calculation is referred to as *probabilistic data association*. However, a direct exact calculation is typically infeasible since it has a complexity of order $O\big((M+1)^K\big)$ and thus scales poorly with K and M.

An efficient and scalable approximate calculation of the marginal association pmfs $p_{\mathrm{m}}(a_k)$ can be obtained by introducing the *measurement-oriented association variables* b_m for $m \in \mathcal{M}$, which are defined as $k \in \mathcal{K}$ if measurement m originates from target k and 0 if measurement m does not originate from any target [6, 46]. We also define $\boldsymbol{b} \triangleq [b_1 \cdots b_M]^{\mathrm{T}}$. Describing valid measurement-target associations (i.e., those satisfying the data association assumption) in terms of both $\boldsymbol{a}$ and $\boldsymbol{b}$ is redundant in that $\boldsymbol{b}$ can be determined from $\boldsymbol{a}$ and also vice versa. However, since $\boldsymbol{a}$ describes only associations where at most one measurement is associated with a target and $\boldsymbol{b}$ describes only associations where at most one target is associated with a measurement, using $\boldsymbol{b}$ alongside with $\boldsymbol{a}$ automatically takes into account the data association assumption. The factor $\Phi(\boldsymbol{a})$ in (9) can now be replaced by $\Psi(\boldsymbol{a}, \boldsymbol{b}) = \prod_{k=1}^{K} \prod_{m=1}^{M} \psi_{km}(a_k, b_m)$, where the "association consistency indicator" $\psi_{km}(a_k, b_m)$ is defined as $\psi_{km}(a_k, b_m) = 0$ if either $a_k = m$ and $b_m \neq k$ or $b_m = k$ and $a_k \neq m$, and $\psi_{km}(a_k, b_m) = 1$ otherwise [6, 46]. In our example, $\Psi(\boldsymbol{a}, \boldsymbol{b}) = 1$ only if $\boldsymbol{a} \in \mathcal{A}$ and if $\boldsymbol{b}$ describes the same association as $\boldsymbol{a}$: for instance, if $\boldsymbol{a} = [2\ 0]^{\mathrm{T}}$ (the first target generates the second measurement and the second target does not generate a measurement), then $\Psi(\boldsymbol{a}, \boldsymbol{b}) = 1$ if and only if $\boldsymbol{b} = [0\ 1]^{\mathrm{T}}$ (the first measurement does not originate from a target and the second measurement originates from the first target). Using $\boldsymbol{b}$ alongside with $\boldsymbol{a}$, expression (9) is formally replaced by

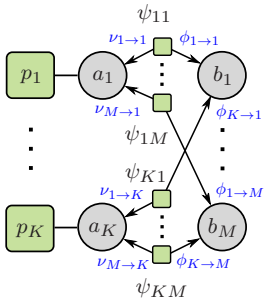

Fig. 2: Factor graph describing the factorization of $p_{\mathrm{j}}(\boldsymbol{a}, \boldsymbol{b})$ in (10). Messages are depicted in blue, and the short notations ψ_{km}, p_k, $\phi_{k\to m}$, and $\nu_{m\to k}$ are used for $\psi_{km}(a_k, b_m)$, $p_{\mathrm{s}}(a_k)$, $\phi^{[\ell]}_{k\to m}$, and $\nu^{[\ell]}_{m\to k}$, respectively.

$$p_{\mathrm{j}}(\boldsymbol{a}, \boldsymbol{b}) \propto \prod_{k=1}^{K} p_{\mathrm{s}}(a_k) \prod_{m=1}^{M} \psi_{km}(a_k, b_m). \tag{10}$$

This can be interpreted as a factorization of the type (1); the corresponding factor graph is depicted in Fig. 2. We note that $p_{\mathrm{j}}(\boldsymbol{a}, \boldsymbol{b})$ is consistent with $p_{\mathrm{j}}(\boldsymbol{a})$ in (9) in that $p_{\mathrm{j}}(\boldsymbol{a}) = \sum_{\boldsymbol{b} \in \{0,1,\ldots,K\}^M} p_{\mathrm{j}}(\boldsymbol{a}, \boldsymbol{b})$.

2.4 Sum-Product Algorithm for Probabilistic Data Association (SPADA)

Accurate approximations of the marginal association probabilities $p_{\mathrm{m}}(a_k)$ can be calculated by running the iterative SPA on the factor graph depicted in Fig. 2. In message passing iteration $\ell \in \{1, 2, \ldots, L\}$, messages $\zeta^{[\ell]}_{\psi_{km}\to b_m}(b_m)$ and $\zeta^{[\ell]}_{\psi_{km}\to a_k}(a_k)$ are passed from factor node "$\psi_{km}(a_k, b_m)$" to variable nodes "b_m" and "a_k", respectively. Using the discrete counterpart of (2) gives $\zeta^{[\ell]}_{\psi_{km}\to b_m}(b_m) = \sum_{a_k=0}^{M} \psi_{km}(a_k, b_m)\, \eta^{[\ell]}_{a_k\to\psi_{km}}(a_k)$. Here, $\eta^{[\ell]}_{a_k\to\psi_{km}}(a_k)$ is the message passed from variable node "a_k" to factor node "$\psi_{km}(a_k, b_m)$", which according to (3) is given by $\eta^{[\ell]}_{a_k\to\psi_{km}}(a_k) = p_{\mathrm{s}}(a_k) \prod_{\substack{m'=1 \\ m' \neq m}}^{M} \zeta^{[\ell]}_{\psi_{km'}\to a_k}(a_k)$. Thus, we obtain

$$\zeta^{[\ell]}_{\psi_{km}\to b_m}(b_m) = \sum_{a_k=0}^{M} p_{\mathrm{s}}(a_k)\, \psi_{km}(a_k, b_m) \prod_{\substack{m'=1 \\ m' \neq m}}^{M} \zeta^{[\ell]}_{\psi_{km'}\to a_k}(a_k), \tag{11}$$

for $k = 1, 2, \ldots, K$ and $m = 1, 2, \ldots, M$. In a similar manner, we obtain

$$\zeta^{[\ell]}_{\psi_{km}\to a_k}(a_k) = \sum_{b_m=0}^{K} \psi_{km}(a_k, b_m) \prod_{\substack{k'=1 \\ k' \neq k}}^{K} \zeta^{[\ell-1]}_{\psi_{k',m}\to b_m}(b_m).$$

These expressions of $\zeta^{[\ell]}_{\psi_{km}\to b_m}(b_m)$ and $\zeta^{[\ell]}_{\psi_{km}\to a_k}(a_k)$ establish a recursive calculation for $\ell = 1, 2, \ldots, L$, which is initialized by $\zeta^{[0]}_{\psi_{km}\to b_m}(b_m) = \sum_{a_k=0}^{M} p_{\mathrm{s}}(a_k)\, \psi_{km}(a_k, b_m)$. After the last iteration $\ell = L$, approximations $\tilde{p}_{\mathrm{m}}(a_k)$ of the marginal association pmfs $p_{\mathrm{m}}(a_k)$ are obtained by calculating the respective beliefs as explained in Section 2.1. This yields $\tilde{p}_{\mathrm{m}}(a_k) = A_k\, p_{\mathrm{s}}(a_k) \prod_{m=1}^{M} \zeta^{[L]}_{\psi_{km}\to a_k}(a_k)$, for $k = 1, 2, \ldots, K$, where the A_k are normalization factors.

An efficient reformulation of this recursion, which we term the *Sum-Product Algorithm for Data Association* (SPADA), is obtained as follows [6, 46]. Due to the binary consistency constraints expressed by $\psi_{km}(a_k, b_m)$, $\zeta^{[\ell]}_{\psi_{km}\to b_m}(b_m)$ in (11) takes on only two different values, namely, one value for $b_m = k$ and another value for all $b_m \neq k$, and similarly for $\zeta^{[\ell]}_{\psi_{km}\to a_k}(a_k)$. Thus, $\zeta^{[\ell]}_{\psi_{km}\to b_m}(b_m)$ and $\zeta^{[\ell]}_{\psi_{km}\to a_k}(a_k)$ can be represented in "normalized" form by the ratio of these two values, which we denote by

$\phi_{k\to m}^{[\ell]}$ and $\nu_{m\to k}^{[\ell]}$, respectively. These ratios can be calculated as [6, 41]

$$\phi_{k\to m}^{[\ell]} = \frac{p_{\mathrm{s}}(m)}{p_{\mathrm{s}}(0) + \sum_{\substack{m'=1 \\ m'\neq m}}^{M} p_{\mathrm{s}}(m')\,\nu_{m'\to k}^{[\ell]}}, \tag{12}$$

$$\nu_{m\to k}^{[\ell]} = \frac{1}{1 + \sum_{\substack{k'=1 \\ k'\neq k}}^{K} \phi_{k'\to m}^{[\ell-1]}}, \tag{13}$$

for $k=1,2,\ldots,K$ and $m=1,2,\ldots,M$. The recursion established by (12) and (13) is initialized by $\phi_{k\to m}^{[0]} = p_{\mathrm{s}}(m)\sum_{\substack{m'=0 \\ m'\neq m}}^{M} p_{\mathrm{s}}(m')$. After termination of the recursion, the beliefs are calculated from the final messages $\nu_{m\to k}^{[L]}$ as

$$\tilde{p}_{\mathrm{m}}(a_k = m) = \frac{p_{\mathrm{s}}(m)\,\nu_{m\to k}^{[L]}}{p_{\mathrm{s}}(0) + \sum_{m'=1}^{M} p_{\mathrm{s}}(m')\,\nu_{m'\to k}^{[L]}},$$

for $m = 0, 1, \ldots, M$. Here, $\nu_{0\to k}^{[L]} \triangleq 1$. The recursion is guaranteed to converge [46]. The complexity of one SPADA iteration (12), (13) scales only linearly in both K and M.

3 SPA-based Multisensor-Multitarget Tracking

Next, we describe how factor graphs and the SPA enable the development of an algorithm that tracks an unknown number of targets using fused data from multiple radar sensors [6, 41].

3.1 Basic System Model

Since the number of targets is unknown, we consider K *potential targets* (PTs) $k \in \mathcal{K} \triangleq \{1, 2, \ldots, K\}$ [6, 41]. Thus, the number of targets is at most K. The existence of PT k at a given (current) time step n is described by the binary variable $r_{k,n} \in \{0, 1\}$, i.e., PT k exists if $r_{k,n} = 1$ and does not exist if $r_{k,n} = 0$. The *state* of PT k at time n is denoted as $\boldsymbol{x}_{k,n}$; it is formally considered also if $r_{k,n} = 0$. The state consists of the PT's position and possibly other (continuous) parameters. We define the *augmented state* of PT k as $\boldsymbol{y}_{k,n} \triangleq [\boldsymbol{x}_{k,n}^{\mathrm{T}}\ r_{k,n}]^{\mathrm{T}}$, and we furthermore define $\boldsymbol{y}_n \triangleq [\boldsymbol{y}_{1,n}^{\mathrm{T}} \cdots \boldsymbol{y}_{K,n}^{\mathrm{T}}]^{\mathrm{T}}$ and $\boldsymbol{y}_{1:n} \triangleq [\boldsymbol{y}_1^{\mathrm{T}} \cdots \boldsymbol{y}_n^{\mathrm{T}}]^{\mathrm{T}}$. We assume that the temporal evolution of each augmented PT state $\boldsymbol{y}_{k,n}$ follows a first-order Markov model, which is described by a state-transition pdf $f(\boldsymbol{y}_{k,n}|\boldsymbol{y}_{k,n-1})$. Different PTs k evolve independently, so that $f(\boldsymbol{y}_n|\boldsymbol{y}_{n-1}) = \prod_{k=1}^{K} f(\boldsymbol{y}_{k,n}|\boldsymbol{y}_{k,n-1})$. An expression for $f(\boldsymbol{y}_{k,n}|\boldsymbol{y}_{k,n-1})$ is provided in [41]. This expression involves a state-transition pdf $f(\boldsymbol{x}_{k,n}|\boldsymbol{x}_{k,n-1})$, a survival probability $p_{k,n}^{(\mathrm{s})}$, a birth probability $p_{k,n}^{(\mathrm{b})}$, and a birth pdf $f^{(\mathrm{b})}(\boldsymbol{x}_{k,n})$.

We consider S sensors $s \in \{1, 2, \ldots, S\}$ producing "detected" radar measurements, i.e., the results of the detection stage of the radar signal processing chain [7]. Let $\boldsymbol{z}_{m,n}^{(s)}$ with $m \in \mathcal{M}_n^{(s)} \triangleq \{1, 2, \ldots, M_n^{(s)}\}$ denote the measurements observed by sensor s at time n. We also set $\boldsymbol{z}_n^{(s)} \triangleq \left[\boldsymbol{z}_{1,n}^{(s)\mathrm{T}} \cdots \boldsymbol{z}_{M_n^{(s)},n}^{(s)\mathrm{T}}\right]^{\mathrm{T}}$, and define $\boldsymbol{z}_{1:n}$ as the vector stacking the $\boldsymbol{z}_{n'}^{(s)}$ for all $s = 1, 2, \ldots, S$ and $n' = 1, 2, \ldots, n$. An existing PT k generates a measurement $\boldsymbol{z}_{m,n}^{(s)}$ at sensor s at time n—i.e., it is "detected" by sensor s—with probability $P_{\mathrm{d}}^{(s)}$. Furthermore, for each sensor s and time n, we consider target-oriented association variables $a_{k,n}^{(s)}$, $k \in \mathcal{K}$ and measurement-oriented association variables $b_{m,n}^{(s)}$, $m \in \mathcal{M}_n^{(s)}$ according to Section 2.3. The vectors of all the $a_{k,n}^{(s)}$, $k \in \mathcal{K}$ and all the $b_{m,n}^{(s)}$, $m \in \mathcal{M}_n^{(s)}$ for all the sensors s at time n are denoted by $\boldsymbol{a}_n$ and $\boldsymbol{b}_n$, respectively. If measurement $\boldsymbol{z}_{m,n}^{(s)}$ originates from PT k, i.e., $a_{k,n}^{(s)} = m \in \mathcal{M}_n^{(s)}$ and $b_{m,n}^{(s)} = k \in \mathcal{K}$, then its distribution given $\boldsymbol{x}_{k,n}$ is described by $f\big(\boldsymbol{z}_{m,n}^{(s)}|\boldsymbol{x}_{k,n}\big)$. Measurements can originate only from existing PTs. Some of the measurements $\boldsymbol{z}_{m,n}^{(s)}$ may be false alarms. Each false alarm is distributed according to some pdf $f_{\mathrm{FA}}\big(\boldsymbol{z}_{m,n}^{(s)}\big)$, and the number of false alarms is modeled by a Poisson distribution with mean $\mu^{(s)}$. The quantities $f\big(\boldsymbol{z}_{m,n}^{(s)}|\boldsymbol{x}_{k,n}\big)$, $\mu^{(s)}$, and $f_{\mathrm{FA}}\big(\boldsymbol{z}_{m,n}^{(s)}\big)$ characterize the measurement model of sensor s. It is assumed that given $\boldsymbol{y}_n$, all measurements $\boldsymbol{z}_{m,n}^{(s)}$ are statistically independent across s and m as well as statistically independent of all previous and future PT states. Furthermore, it is assumed that given $\boldsymbol{y}_{n-1}$, $\boldsymbol{y}_n$ is conditionally independent of all previous measurements.

3.2 Detection and Estimation

Our goal is to decide if PT $k \in \mathcal{K}$ exists at time n, or equivalently, to detect the event $r_{k,n} = 1$ (this should not be confused with the "detection" performed by the sensors), and to calculate estimates of the states $\boldsymbol{x}_{k,n}$ of the detected PTs. In our Bayesian framework, target detection and state estimation are based on the posterior existence probabilities $p(r_{k,n} = 1|\boldsymbol{z}_{1:n})$ and the posterior state pdfs $f(\boldsymbol{x}_{k,n}|r_{k,n} = 1, \boldsymbol{z}_{1:n})$. PT k is detected—i.e., its existence is declared—if $p(r_{k,n} = 1|\boldsymbol{z}_{1:n})$ exceeds a suitably chosen threshold P_{th} [58, Ch. 2]. In that case, an estimate of $\boldsymbol{x}_{k,n}$ is given, e.g., by $\hat{\boldsymbol{x}}_{k,n} \triangleq \int \boldsymbol{x}_{k,n} f(\boldsymbol{x}_{k,n}|r_{k,n} = 1, \boldsymbol{z}_{1:n})\,\mathrm{d}\boldsymbol{x}_{k,n}$ [58, Ch. 4]. The required statistics $p(r_{k,n} = 1|\boldsymbol{z}_{1:n})$ and $f(\boldsymbol{x}_{k,n}|r_{k,n} = 1, \boldsymbol{z}_{1:n})$ can be obtained from the posterior pdf $f(\boldsymbol{x}_{k,n}, r_{k,n}|\boldsymbol{z}_{1:n}) = f(\boldsymbol{y}_{k,n}|\boldsymbol{z}_{1:n})$ according to

$$p(r_{k,n} = 1|\boldsymbol{z}_{1:n}) = \int f(\boldsymbol{x}_{k,n}, r_{k,n} = 1|\boldsymbol{z}_{1:n})\,\mathrm{d}\boldsymbol{x}_{k,n}$$

and

$$f(\boldsymbol{x}_{k,n}|r_{k,n} = 1, \boldsymbol{z}_{1:n}) = \frac{f(\boldsymbol{x}_{k,n}, r_{k,n} = 1|\boldsymbol{z}_{1:n})}{p(r_{k,n} = 1|\boldsymbol{z}_{1:n})},$$

respectively. Hence, it remains to calculate the posterior pdfs $f(\boldsymbol{x}_{k,n}, r_{k,n}|\boldsymbol{z}_{1:n})$.

3.3 Factor Graph and SPA-based Tracking Algorithm

The posterior pdf $f(\boldsymbol{x}_{k,n}, r_{k,n}|\boldsymbol{z}_{1:n}) = f(\boldsymbol{y}_{k,n}|\boldsymbol{z}_{1:n})$ is a marginal density of the joint posterior pdf $f(\boldsymbol{y}_n, \boldsymbol{y}_{n-1}, \boldsymbol{a}_n, \boldsymbol{b}_n|\boldsymbol{z}_{1:n})$, which describes the joint distribution of all the current and previous augmented states and all the current association variables, conditioned on all the measurements observed so far. Following Section 2.1, we can use the SPA for an efficient and scalable approximate calculation of the K marginal posterior pdfs $f(\boldsymbol{y}_{k,n}|\boldsymbol{z}_{1:n})$, $k \in \mathcal{K}$.

3.3.1 Known Measurement Origins: As a preliminary step, we first consider the calculation of $f(\boldsymbol{y}_{k,n}|\boldsymbol{z}_{1:n})$ in the absence of MOU, missed detections, and false alarms. More specifically, we assume that it is known that each target k generates a corresponding measurement $\boldsymbol{z}_{k,n}^{(s)}$ at each sensor s, i.e., $m = k$ and $M_n^{(s)} = M_n$ for all $s \in \{1, 2, \ldots, S\}$. It follows that the number of targets is equal to M_n and hence known. We thus have $K = M_n$ and $r_{k,n} = 1$ for all $k = 1, \ldots, K$. In addition, at time $n = 0$ we initialize the prior pdf of the augmented state $\boldsymbol{y}_{k,0}$ as $f(\boldsymbol{x}_{k,0}, r_{k,0} = 1) = f(\boldsymbol{x}_{k,0})$, where $f(\boldsymbol{x}_{k,0})$ is an arbitrary prior pdf, and as $f(\boldsymbol{x}_{k,0}, r_{k,0} = 0) = 0$ because the event $r_{k,0} = 0$ is impossible.

In the absence of MOU, the joint posterior pdf is $f(\boldsymbol{y}_n, \boldsymbol{y}_{n-1}|\boldsymbol{z}_{1:n})$, without the association vectors $\boldsymbol{a}_n$ and $\boldsymbol{b}_n$. By applying Bayes' rule and using assumptions introduced in Section 3.1, the joint posterior pdf factorizes as

$$\begin{aligned}
&f(\boldsymbol{y}_n, \boldsymbol{y}_{n-1}|\boldsymbol{z}_{1:n}) \\
&\quad = f(\boldsymbol{y}_n, \boldsymbol{y}_{n-1}|\boldsymbol{z}_n, \boldsymbol{z}_{1:n-1}) \\
&\quad \propto f(\boldsymbol{z}_n|\boldsymbol{y}_n, \boldsymbol{y}_{n-1}, \boldsymbol{z}_{1:n-1})\, f(\boldsymbol{y}_n, \boldsymbol{y}_{n-1}|\boldsymbol{z}_{1:n-1}) \\
&\quad = f(\boldsymbol{z}_n|\boldsymbol{y}_n)\, f(\boldsymbol{y}_n, \boldsymbol{y}_{n-1}|\boldsymbol{z}_{1:n-1}) \\
&\quad = f(\boldsymbol{z}_n|\boldsymbol{y}_n)\, f(\boldsymbol{y}_n|\boldsymbol{y}_{n-1})\, f(\boldsymbol{y}_{n-1}|\boldsymbol{z}_{1:n-1}) \\
&\quad = f(\boldsymbol{y}_{n-1}|\boldsymbol{z}_{1:n-1})\, f(\boldsymbol{y}_n|\boldsymbol{y}_{n-1}) \prod_{s=1}^{S} f(\boldsymbol{z}_n^{(s)}|\boldsymbol{y}_n)
\end{aligned}$$

$$= \prod_{k=1}^{K} \underbrace{f(\boldsymbol{y}_{k,n-1}|\boldsymbol{z}_{1:n-1})}_{\substack{\text{PREVIOUS MARGINAL}\\ \text{POSTERIOR PDF}}} \underbrace{f(\boldsymbol{y}_{k,n}|\boldsymbol{y}_{k,n-1})}_{\substack{\text{STATE-TRANSITION}\\ \text{PDF}}} \times \prod_{s=1}^{S} \underbrace{f(\boldsymbol{z}_{k,n}^{(s)}|\boldsymbol{y}_{k,n})}_{\substack{\text{LIKELIHOOD}\\ \text{FUNCTION}}} .$$

Here, we defined $\boldsymbol{z}_n$ to be the vector stacking $\boldsymbol{z}_n^{(1)}, \ldots, \boldsymbol{z}_n^{(S)}$, and we introduced $f(\boldsymbol{z}_{k,n}^{(s)}|\boldsymbol{y}_{k,n}) \triangleq f(\boldsymbol{z}_{k,n}^{(s)}|\boldsymbol{x}_{k,n}, r_{k,n})$, with $f(\boldsymbol{z}_{k,n}^{(s)}|\boldsymbol{x}_{k,n}, r_{k,n}=1) = f(\boldsymbol{z}_{k,n}^{(s)}|\boldsymbol{x}_{k,n})$ and $f(\boldsymbol{z}_{k,n}^{(s)}|\boldsymbol{x}_{k,n}, r_{k,n}=0)$ arbitrary (this pdf is irrelevant because the event $r_{k,n}=0$ is impossible). Furthermore, we assumed that $f(\boldsymbol{y}_{n-1}|\ \boldsymbol{z}_{1:n-1}) = \prod_{k=1}^{K} f(\boldsymbol{y}_{k,n-1}|\boldsymbol{z}_{1:n-1})$. This is a standard assumption that underlies many multitarget tracking algorithms [16] and can be shown to be equivalent to sending messages only forward in time on a factor graph derived for all times n [6]. Next, marginalizing out $\boldsymbol{y}_{n-1}$ yields

$$\begin{aligned} f(\boldsymbol{y}_n|\boldsymbol{z}_{1:n}) &= \int f(\boldsymbol{y}_n, \boldsymbol{y}_{n-1}|\boldsymbol{z}_{1:n})\,\mathrm{d}\boldsymbol{y}_{n-1} \\ &= \prod_{k=1}^{K} \underbrace{\int f(\boldsymbol{y}_{k,n}|\boldsymbol{y}_{k,n-1})\, f(\boldsymbol{y}_{k,n-1}|\boldsymbol{z}_{1:n-1})\,\mathrm{d}\boldsymbol{y}_{k,n-1}}_{= f(\boldsymbol{y}_{k,n}|\boldsymbol{z}_{1:n-1}) \text{ (PREDICTION STEP)}} \\ &\quad \times \prod_{s=1}^{S} f(\boldsymbol{z}_{k,n}^{(s)}|\boldsymbol{y}_{k,n}) \\ &= \prod_{k=1}^{K} \underbrace{f(\boldsymbol{y}_{k,n}|\boldsymbol{z}_{1:n-1}) \prod_{s=1}^{S} f(\boldsymbol{z}_{k,n}^{(s)}|\boldsymbol{y}_{k,n})}_{\propto f(\boldsymbol{y}_{k,n}|\boldsymbol{z}_{1:n})} . \end{aligned}$$

(MULTISENSOR MEASUREMENT UPDATE STEP)

This shows that the multitarget tracking problem for our K targets simplifies to K independent sequential Bayesian estimation problems, where the kth problem consists of a prediction step ($f(\boldsymbol{y}_{k,n}|\boldsymbol{z}_{1:n-1}) = \int f(\boldsymbol{y}_{k,n}|\boldsymbol{y}_{k,n-1})\, f(\boldsymbol{y}_{k,n-1}|\boldsymbol{z}_{1:n-1})\,\mathrm{d}\boldsymbol{y}_{k,n-1}$, cf. (5)) and a multisensor measurement update step ($f(\boldsymbol{y}_{k,n}|\boldsymbol{z}_{1:n}) \propto f(\boldsymbol{y}_{k,n}|\boldsymbol{z}_{1:n-1}) \prod_{s=1}^{S} f(\boldsymbol{z}_{k,n}^{(s)}|\boldsymbol{y}_{k,n})$, cf. (6)). At time $n=1$, the pdf $f(\boldsymbol{y}_{k,n-1}|\boldsymbol{z}_{1:n-1}) = f(\boldsymbol{y}_{k,0}|\boldsymbol{z}_{1:0})$ is formally replaced by the prior pdf $f(\boldsymbol{y}_{k,0}) = f(\boldsymbol{x}_{k,0}, r_{k,0})$. Since $f(\boldsymbol{x}_{k,0}, r_{k,0}=0) = 0$, it is straightforward to see that after the prediction step, we obtain the pdf $f(\boldsymbol{y}_{k,1}) = f(\boldsymbol{x}_{k,1}, r_{k,1})$ which also satisfies $f(\boldsymbol{x}_{k,1}, r_{k,1}=0) = 0$. Similarly, also after the update step the obtained posterior pdf $f(\boldsymbol{y}_{k,1}|\boldsymbol{z}_{k,1}) = f(\boldsymbol{x}_{k,1}, r_{k,1}|\boldsymbol{z}_{k,1})$ still satisfies $f(\boldsymbol{x}_{k,1}, r_{k,1}=0|\boldsymbol{z}_{k,1}) = 0$. Continuing this way, it is seen that $f(\boldsymbol{x}_{k,n}, r_{k,n}=0|\boldsymbol{z}_{1:n-1}) = 0$ and $f(\boldsymbol{x}_{k,n}, r_{k,n}=0|\boldsymbol{z}_{1:n}) = 0$ also for all subsequent times n.

Furthermore, the marginalization of $f(\boldsymbol{y}_{k,n}|\boldsymbol{z}_{1:n}) = f(\boldsymbol{x}_{k,n}, r_{k,n}|\boldsymbol{z}_{1:n})$ for $r_{k,n}=1$ yields

$$p(r_{k,n}=1|\boldsymbol{z}_{1:n}) = \int f(\boldsymbol{x}_{k,n}, r_{k,n}=1|\boldsymbol{z}_{1:n})\,\mathrm{d}\boldsymbol{x}_{k,n} = 1$$

for all $k \in \mathcal{K}$, where the final result follows from the fact that $f(\boldsymbol{x}_{k,n}, r_{k,n}=0|\boldsymbol{z}_{1:n}) = 0$.

3.3.2 Unknown Measurement Origins: We now revert to the general case with MOU, missed detections, and false alarms (and, thus, unknown $\boldsymbol{a}_n$ and $\boldsymbol{b}_n$). Here, the K target tracking problems are coupled and the existence of the PTs is not certain. Using the assumptions from Section 3.1, one can show [41] that $f(\boldsymbol{y}_n, \boldsymbol{y}_{n-1}, \boldsymbol{a}_n, \boldsymbol{b}_n|\boldsymbol{z}_{1:n})$ factorizes as

$$f(\boldsymbol{y}_n, \boldsymbol{y}_{n-1}, \boldsymbol{a}_n, \boldsymbol{b}_n|\boldsymbol{z}_{1:n}) \propto \prod_{k=1}^{K} \underbrace{f(\boldsymbol{y}_{k,n-1}|\boldsymbol{z}_{1:n-1})}_{\substack{\text{PREVIOUS MARGINAL}\\ \text{POSTERIOR PDF}}} \underbrace{f(\boldsymbol{y}_{k,n}|\boldsymbol{y}_{k,n-1})}_{\substack{\text{STATE-TRANSITION}\\ \text{PDF}}}$$

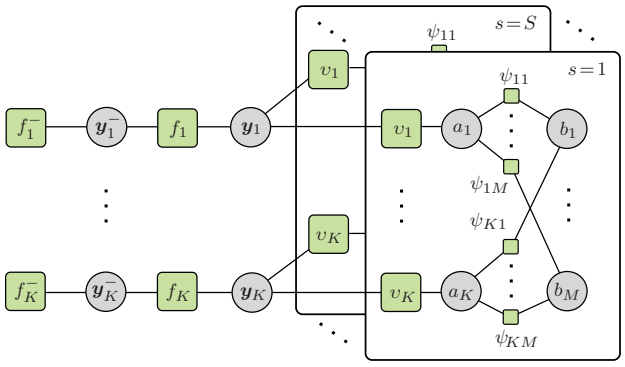

Fig. 3: Factor graph describing the factorization of $f(\boldsymbol{y}_n, \boldsymbol{y}_{n-1}, \boldsymbol{a}_n, \boldsymbol{b}_n|\boldsymbol{z}_{1:n})$ in (14). The time index n and the sensor index s are omitted. Furthermore, we use the following short notations: $f_k^- \triangleq f(\boldsymbol{y}_{k,n-1}|\boldsymbol{z}_{1:n-1})$, $\boldsymbol{y}_k^- \triangleq \boldsymbol{y}_{k,n-1}$, $f_k \triangleq f(\boldsymbol{y}_{k,n}|\boldsymbol{y}_{k,n-1})$, $\boldsymbol{y}_k \triangleq \boldsymbol{y}_{k,n}$, $\upsilon_k \triangleq \upsilon^{(s)}(\boldsymbol{y}_{k,n}, a_{k,n}^{(s)}; \boldsymbol{z}_n^{(s)})$, $a_k \triangleq a_{k,n}^{(s)}$, $b_m \triangleq b_{m,n}^{(s)}$, $\psi_{km} \triangleq \psi_{km}(a_{k,n}^{(s)}, b_{m,n}^{(s)})$, and $M \triangleq M_n^{(s)}$.

$$\times \prod_{s=1}^{S} \underbrace{\upsilon^{(s)}(\boldsymbol{y}_{k,n}, a_{k,n}^{(s)}; \boldsymbol{z}_n^{(s)})}_{\substack{\text{PSEUDO-LIKELIHOOD}\\ \text{FUNCTION}}} \prod_{m=1}^{M_n^{(s)}} \underbrace{\psi_{km}(a_{k,n}^{(s)}, b_{m,n}^{(s)})}_{\substack{\text{ASSOCIATION}\\ \text{CONSISTENCY INDICATOR}}} . \tag{14}$$

Here, the association consistency indicators $\psi_{km}(a_{k,n}^{(s)}, b_{m,n}^{(s)})$ were introduced in Section 2.3, and $\upsilon^{(s)}(\boldsymbol{y}_{k,n}, a_{k,n}^{(s)}; \boldsymbol{z}_n^{(s)}) = \upsilon^{(s)}(\boldsymbol{x}_{k,n}, r_{k,n}, a_{k,n}^{(s)}; \boldsymbol{z}_n^{(s)})$ is given by

$$\upsilon^{(s)}(\boldsymbol{x}_{k,n}, 1, a_{k,n}^{(s)}; \boldsymbol{z}_n^{(s)}) = \begin{cases} \dfrac{P_{\mathrm{d}}^{(s)} f(\boldsymbol{z}_{m,n}^{(s)}|\boldsymbol{x}_{k,n})}{\mu^{(s)} f_{\mathrm{FA}}(\boldsymbol{z}_{m,n}^{(s)})}, & a_{k,n}^{(s)} = m \in \mathcal{M}_n^{(s)} \\ 1 - P_{\mathrm{d}}^{(s)}, & a_{k,n}^{(s)} = 0 \end{cases} \tag{15}$$

and

$$\upsilon^{(s)}(\boldsymbol{x}_{k,n}, 0, a_{k,n}^{(s)}; \boldsymbol{z}_n^{(s)}) = \delta_{a_{k,n}^{(s)},0}, \tag{16}$$

where $\delta_{a,0}$ is defined to be 1 if $a=0$ and 0 otherwise. According to (14), the joint posterior pdf $f(\boldsymbol{y}_n, \boldsymbol{y}_{n-1}, \boldsymbol{a}_n, \boldsymbol{b}_n|\boldsymbol{z}_{1:n})$ is nonzero only if $\boldsymbol{y}_n$, $\boldsymbol{a}_n$, and $\boldsymbol{b}_n$ are consistent in that (i) $\boldsymbol{a}_n$ and $\boldsymbol{b}_n$ describe the same data association event (otherwise $\psi_{km}(a_{k,n}^{(s)}, b_{k,n}^{(s)})$ is zero for some $s \in \mathcal{S}$, $k \in \mathcal{K}$, and $m \in \mathcal{M}_n^{(s)}$) and (ii) for all $k \in \mathcal{K}$, $r_{k,n}=0$ implies $a_{k,n}^{(s)}=0$ for all $s \in \mathcal{S}$, i.e., PTs that do not exist cannot generate measurements (otherwise $\upsilon^{(s)}(\boldsymbol{x}_{k,n}, 0, a_{k,n}^{(s)}; \boldsymbol{z}_n^{(s)})$ is zero for some $s \in \mathcal{S}$).

The complexity of calculating the marginal posterior pdfs $f(\boldsymbol{x}_{k,n}, r_{k,n}|\boldsymbol{z}_{1:n}) = f(\boldsymbol{y}_{k,n}|\boldsymbol{z}_{1:n})$, $k \in \mathcal{K}$ by direct marginalization of the joint posterior pdf $f(\boldsymbol{y}_n, \boldsymbol{y}_{n-1}, \boldsymbol{a}_n, \boldsymbol{b}_n|\boldsymbol{z}_{1:n})$ in (14) scales exponentially in the number of PTs K, the number of measurements $M_n^{(s)}$, and the number of sensors S. However, running the SPA on the loopy factor graph representing the factorization (14) makes it possible to calculate approximations of $f(\boldsymbol{y}_{k,n}|\boldsymbol{z}_{1:n})$ in an efficient and scalable manner. This factor graph is shown in Fig. 3; it is seen to contain one instance of the SPADA factor graph (shown in Fig. 2) for each sensor s. The beliefs $\tilde{f}(\boldsymbol{y}_{k,n})$ produced by the SPA are used instead of $f(\boldsymbol{y}_{k,n}|\boldsymbol{z}_n)$ for detection and estimation according to Section 3.2, and also in the SPA for the next time step. Explicit expressions of the messages and beliefs can be obtained by applying the general SPA message calculation rules from Section 2.1 to the factor graph in Fig. 3 [41]. To reduce the complexity as

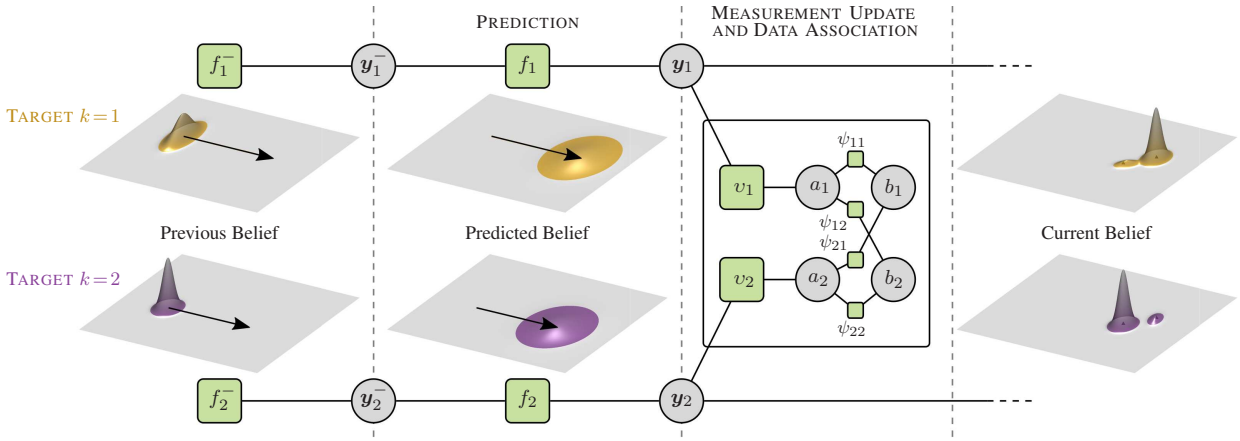

Fig. 4: Factor graph describing the factorization of $f(\boldsymbol{y}_n, \boldsymbol{y}_{n-1}, \boldsymbol{a}_n, \boldsymbol{b}_n|\boldsymbol{z}_{1:n})$ in (14) for the case of two targets and two measurements produced by a single sensor. In addition, the previous, predicted, and current beliefs for the two targets are shown.

well as detrimental effects of loops in the factor graph, we employ a schedule of message calculation such that *iterative* message passing is performed only for probabilistic data association (i.e., within the individual SPADA subgraphs) but not across the sensors [41]. Furthermore, to accommodate nonlinear and/or non-Gaussian state-transition and measurement models, we represent messages related to continuous state variables by particles [41].

For an illustration of the SPA-based tracking algorithm, we reconsider our simple example with $K=2$ targets and $M_n^{(1)}=2$ measurements produced by a single sensor. The factor graph representing the factorization (14) for this case is shown in Fig. 4. It is structured into four sections that refer to various distributions (messages, beliefs) and/or operations within the SPA-based tracking algorithm. The *previous belief* for target $k\in\{1,2\}$ approximates the marginal posterior pdf of the previous augmented state $\boldsymbol{y}_{k,n-1}$, i.e., $f(\boldsymbol{y}_{k,n-1}|\boldsymbol{z}_{1:n-1})$, which corresponds to the factor node "f_k^-" in Fig. 4. In the *prediction* step, the previous belief is converted into the *predicted belief*, which is the message passed from factor node "f_k" (representing the state-transition pdf $f(\boldsymbol{y}_{k,n}|\boldsymbol{y}_{k,n-1})$) to variable node "$\boldsymbol{y}_k$" (representing $\boldsymbol{y}_{k,n}$). In the *measurement update and data association* step, the measurement $\boldsymbol{z}_n^{(1)}$ is incorporated into the function $\upsilon^{(1)}(\boldsymbol{y}_{k,n}, a_{k,n}^{(1)}; \boldsymbol{z}_n^{(1)})$ (represented by the factor node "υ_k"), and the SPADA (cf. Section 2.4) is employed for probabilistic data association. This results in the *current belief*, which approximates the current marginal posterior pdf $f(\boldsymbol{y}_{k,n}|\boldsymbol{z}_{1:n})$. The current belief is centered around the respective measurement. It is used for target detection and for state estimation as discussed in Section 3.2, and also as the previous belief at the next time step. In addition to the factor graph, Fig. 4 also visualizes the previous, predicted, and current beliefs for the first and second target in the upper and lower three-dimensional plots, respectively. The arrows in the left-most plots represent the trajectories of the targets. Because of the uncertainty of target-measurement association and the proximity of the two targets, the current beliefs are bimodal. The smaller of the two modes is centered roughly at the position of the respective other target.

4 Adapting the Dynamic Model

In many tracking scenarios, an accurate description of the dynamics of a target requires the use of different dynamic models (DMs) in different time periods. A prominent example is the case of maneuvering targets [61]. Therefore, adopting the multiple model approach [61, Ch. 11], we describe the state evolution for an existing PT $k\in\mathcal{K}$ as

$$\boldsymbol{x}_{k,n} = \boldsymbol{\xi}_{\ell_{k,n}}\big(\boldsymbol{x}_{k,n-1}, \boldsymbol{u}_{k,n}^{(\ell_{k,n})}\big).$$

Here, the driving process $\boldsymbol{u}_{k,n}^{(\ell_{k,n})}$ is independent and identically distributed (iid) across k and n [16, 61]. The state-transition function $\boldsymbol{\xi}_{\ell_{k,n}}(\cdot,\cdot)$ and the statistics of $\boldsymbol{u}_{k,n}^{(\ell_{k,n})}$ then define the DM $\mathcal{D}_{\ell_{k,n}}$, which is chosen from a set of possible DMs, $\{\mathcal{D}_j\}_{j=1}^J$, by the DM index $\ell_{k,n}\in\{1,2,\ldots,J\}$. Note that the state-transition pdf $f_j(\boldsymbol{x}_{k,n}|\boldsymbol{x}_{k,n-1})$ follows from $\mathcal{D}_j$. The most common DM type is the nearly-constant velocity model, where the PT state consists of two-dimensional (2D) position and 2D velocity, i.e., $\boldsymbol{x}_{k,n} = [x_{1,k,n}\ \ x_{2,k,n}\ \ \dot{x}_{1,k,n}\ \ \dot{x}_{2,k,n}]^{\mathrm{T}}$, and

$$\boldsymbol{x}_{k,n} = \boldsymbol{\xi}_j\big(\boldsymbol{x}_{k,n-1}, \boldsymbol{u}_{k,n}^{(j)}\big) = \boldsymbol{A}\boldsymbol{x}_{k,n-1} + \boldsymbol{W}\boldsymbol{u}_{k,n}^{(j)}. \tag{17}$$

Here, the matrices $\boldsymbol{A}\in\mathbb{R}^{4\times4}$ and $\boldsymbol{W}\in\mathbb{R}^{4\times2}$ are as specified in [61, Sec. 6.3.2] (they involve the time step duration ΔT as a parameter), and $\boldsymbol{u}_{k,n}^{(j)}\sim\mathcal{N}\big(\boldsymbol{0},\sigma_j^2\,\mathbf{I}_2\big)$ is a sequence of 2D Gaussian random vectors that is iid across PT k and time n. Note that different DMs $\mathcal{D}_j$ differ solely in the driving process variance σ_j^2. Since σ_j is a mean measure of the change of target speed in a time step of duration ΔT, higher values of σ_j^2 are appropriate during target maneuvers, i.e., when a target changes its speed or course.

The DM indices $\ell_{k,n}$ are independent across k; their temporal evolution follows a Markov chain model, i.e.,

$$p(\ell_{k,n}=j|\ell_{k,n-1}=i) = [\boldsymbol{L}]_{i,j}, \quad i,j\in\{1,2,\ldots,J\}, \tag{18}$$

with some transition matrix $\boldsymbol{L}\in[0,1]^{J\times J}$ [49]. We furthermore redefine the augmented state of PT k as $\boldsymbol{y}_{k,n}\triangleq[\boldsymbol{x}_{k,n}^T\ r_{k,n}\ \ell_{k,n}]^{\mathrm{T}}$, i.e., it now includes also the DM index $\ell_{k,n}$. The transition pdf of the augmented state, $f(\boldsymbol{y}_{k,n}|\boldsymbol{y}_{k,n-1})$, is then expressed in terms of a DM-dependent birth pdf $f_{\ell_{k,n}}^{(\mathrm{b})}(\boldsymbol{x}_{k,n})$, a birth probability $p_{k,n}^{(\mathrm{b})}$, the DM-dependent state-transition pdf $f_{\ell_{k,n}}(\boldsymbol{x}_{k,n}|\boldsymbol{x}_{k,n-1})$, and a survival probability $p_{k,n}^{(\mathrm{s})}$ as explained in [49]. Note that the birth pdf $f_{\ell_{k,n}}^{(\mathrm{b})}(\boldsymbol{x}_{k,n})$ may be DM-dependent; in particular, one can choose a broader birth pdf (expressing a larger variability of the state) in the case of fast maneuvering targets. We can then run iterative SPA message passing on a factor graph that differs from the factor graph of Fig. 3 only by the modified definition of the variable nodes "$\boldsymbol{y}_{k,n}$". The result is a "self-tuning" multisensor-multitarget tracking algorithm that performs a probabilistic adaptation of its DM model at each time n [49].

5 Integration of AIS Data

International regulations require commercial ships whose gross tonnage exceeds 300 tons and passenger ships to report their identity, position, and motion information via an AIS transponder. The AIS data format includes the maritime mobile service identity (MMSI) [8, 62]. We next extend our factor graph/SPA framework for multitarget tracking to include a fusion of real-time AIS data with radar measurements [56].

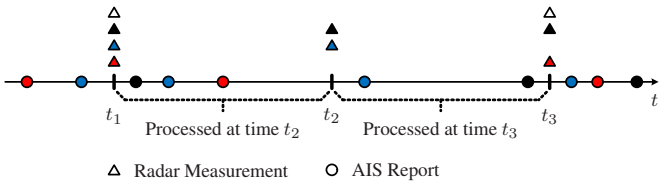

Fig. 5: Example of radar measurements and AIS reports from three ships. The radar measurements and AIS reports are available to the tracking algorithm during three consecutive time steps ending at times t_1, t_2, and t_3, and they are represented by triangles and circles, respectively. The ships are identified by black, blue, and red color.

5.1 A Model for Radar-AIS Information Fusion

The fusion of AIS data with radar measurements needs to cope with several challenges: (i) The rate at which AIS reports are transmitted by a ship depends on the ship's speed and is anything between one report every few seconds and one report every several minutes. (ii) Because of the wireless communication channel, it is possible that some of the AIS reports are not received by a base station. (iii) The reports transmitted by different ships are received by a given base station asynchronously, sparsely in time, and without a predefined order. (iv) The lack of reports from a ship does not imply that that ship is not present in the region of interest. (v) The MMSI may be set to a default value due to a malfunction of the transponder.

Fig. 5 shows an example of radar measurements and AIS reports that are available to the tracking algorithm during three consecutive time steps ending at times t_1, t_2, and t_3. Three ships are in the region of interest and generate measurements at a radar sensor $s \in \{1, 2, \ldots, S\}$ with a certain probability of detection. In addition, AIS reports are received asynchronously (both across the transmitting ships and relative to the radar measurements).

We will formally consider all the available AIS data to originate from an AIS receiver ("AIS sensor"), which is indexed by $s = 0$. The AIS sensor comes in addition to the radar sensors $s \in \{1, 2, \ldots, S\}$, and its data are included in the total measurement vector $\boldsymbol{z}_{1:n}$. The mth AIS report at time n, $\boldsymbol{z}_{m,n}^{(0)}$ (note that the sensor index is $s = 0$), consists of the position information $\boldsymbol{d}_{m,n}$ and the MMSI $\zeta_{m,n}$, i.e., $\boldsymbol{z}_{m,n}^{(0)} = [\boldsymbol{d}_{m,n}^{\mathrm{T}} \; \zeta_{m,n}]^{\mathrm{T}}$. Here, $\zeta_{m,n} \in \{0, 1, \ldots, L\}$; more specifically, $\zeta_{m,n} \in \{1, 2, \ldots, L\}$ is a valid MMSI identifying a ship and $\zeta_{m,n} = 0$ is the default MMSI value mentioned above.

The association between PTs and AIS reports is nontrivial, as the MMSI in an AIS report may be set to a default value or it may have been observed for the first time; in the latter case, no prior information about the ship generating the AIS report is available. Assuming that each ship transmits at most one AIS report in one time step, the data association assumption made in Section 2.3 is still valid, i.e., a PT k is associated with no more than one AIS report m and, conversely, an AIS report m is associated with no more than one PT k. Furthermore, we assume that an AIS report can only originate from a target, and thus it cannot be a false alarm. Therefore, the joint PT-report association pmf $p_{\mathrm{j}}(\boldsymbol{a}, \boldsymbol{b})$ can be expressed as in (10), with $\psi_{km}(a_k, b_m)$ replaced by a modified function $\psi_{km}^{\star}(a_k, b_m)$ that takes into account both the data association assumption and the no-false-alarms assumption, i.e., $\psi_{km}^{\star}(a_k, b_m) = 0$ if either $a_k = m$ and $b_m \neq k$, or $b_m = k$ and $a_k \neq m$, or $b_m = 0$, and $\psi_{km}^{\star}(a_k, b_m) = 1$ otherwise. It follows that the factor graph describing the factorization of $p_{\mathrm{j}}(\boldsymbol{a}, \boldsymbol{b})$ has the same structure as the factor graph depicted in Fig. 2, and hence a SPADA for the AIS sensor can be developed as described in Section 2.4.

5.2 Factor Graph and SPA-based Tracking Algorithm

We now present a framework for multitarget tracking from both radar measurements and AIS reports, which constitute the total measurement vector $\boldsymbol{z}_{1:n}$ [56]. We first redefine the augmented state as $\boldsymbol{y}_{k,n} \triangleq [\boldsymbol{x}_{k,n}^{\mathrm{T}} \; r_{k,n} \; \tau_{k,n}]^{\mathrm{T}}$, i.e., it now includes the *MMSI label* $\tau_{k,n} \in \{0, 1, \ldots, L\}$. Here, $\tau_{k,n} = l \in \{1, 2, \ldots, L\}$ means that at time n, PT k transmits an AIS report with MMSI l, and $\tau_{k,n} = 0$ means that at time n, PT k transmits an AIS report with a default label. Then, the joint posterior pdf $f(\boldsymbol{y}_n, \boldsymbol{y}_{1:n-1}, \boldsymbol{a}_n, \boldsymbol{b}_n | \boldsymbol{z}_{1:n})$ factorizes according to (14) with the following modifications: the product $\prod_{s=1}^{S}$ is replaced by $\prod_{s=0}^{S}$; the function $\psi_{km}(a_{k,n}^{(s)}, b_{m,n}^{(s)})$ is replaced by a modified function $\psi_{km}^{(s)}(a_{k,n}^{(s)}, b_{m,n}^{(s)})$ that equals $\psi_{km}(a_k, b_m)$ (introduced in Section 2.3) for $s \in \{1, 2, \ldots, S\}$ and $\psi_{km}^{\star}(a_k, b_m)$ for $s = 0$; and the function $\upsilon^{(s)}(\boldsymbol{y}_{k,n}, a_{k,n}^{(s)}; \boldsymbol{z}_n^{(s)})$ is redefined to equal $\upsilon^{(s)}(\boldsymbol{x}_{k,n}, r_{k,n}, a_{k,n}^{(s)}; \boldsymbol{z}_n^{(s)})$ (defined in (15) and (16)) for $s \in \{1, 2, \ldots, S\}$ and $\upsilon^{(0)}(\boldsymbol{x}_{k,n}, r_{k,n}, \tau_{k,n}, a_{k,n}^{(0)}; \boldsymbol{z}_n^{(0)})$ for $s = 0$, where

$$\upsilon^{(0)}(\boldsymbol{x}_{k,n}, 1, \tau_{k,n}, a_{k,n}^{(0)}; \boldsymbol{z}_n^{(0)}) = \begin{cases} f(\boldsymbol{z}_{m,n}^{(0)} | \boldsymbol{x}_{k,n}, \tau_{k,n}), & a_{k,n}^{(0)} = m \in \mathcal{M}_n^{(0)} \\ 1, & a_{k,n}^{(0)} = 0 \end{cases} \tag{19}$$

and

$$\upsilon^{(0)}(\boldsymbol{x}_{k,n}, 0, \tau_{k,n}, a_{k,n}^{(0)}; \boldsymbol{z}_n^{(0)}) = \delta_{a_{k,n}^{(0)}, 0},$$

with $\mathcal{M}_n^{(0)}$ denoting the set of indices m corresponding to the AIS reports at time n. Finally, the factor $f(\boldsymbol{y}_{k,n} | \boldsymbol{y}_{k,n-1})$ in (14) now accounts also for the evolution of the MMSI label $\tau_{k,n}$, besides the evolution of $\boldsymbol{x}_{k,n}$ and $r_{k,n}$.

The pdf $f(\boldsymbol{z}_{m,n}^{(0)} | \boldsymbol{x}_{k,n}, \tau_{k,n})$ involved in (19) describes the statistical dependence of the mth AIS report $\boldsymbol{z}_{m,n}^{(0)} = [\boldsymbol{d}_{m,n}^{\mathrm{T}} \; \zeta_{m,n}]^{\mathrm{T}}$ on the state $\boldsymbol{x}_{k,n}$ and label $\tau_{k,n}$ of PT k at time n. Assuming that (i) $\boldsymbol{d}_{m,n}$ and $\zeta_{m,n}$ are conditionally independent given $\boldsymbol{x}_{k,n}$ and $\tau_{k,n}$, (ii) $\boldsymbol{d}_{m,n}$ and $\tau_{k,n}$ are conditionally independent given $\boldsymbol{x}_{k,n}$, and (iii) $\zeta_{m,n}$ and $\boldsymbol{x}_{k,n}$ are conditionally independent given $\tau_{k,n}$, this pdf factorizes as $f(\boldsymbol{z}_{m,n}^{(0)} | \boldsymbol{x}_{k,n}, \tau_{k,n}) = f(\boldsymbol{d}_{m,n} | \boldsymbol{x}_{k,n}) p(\zeta_{m,n} | \tau_{k,n})$. Here, $f(\boldsymbol{d}_{m,n} | \boldsymbol{x}_{k,n})$ describes the AIS "measurement model," which takes into account the asynchronous nature of the AIS sensor (see [56] for details). The second factor, $p(\zeta_{m,n} | \tau_{k,n})$, equals p_{c} if $\zeta_{m,n} = \tau_{k,n}$ and $(1 - p_{\mathrm{c}})/L$ if $\zeta_{m,n} \neq \tau_{k,n}$, where p_{c} is the probability that the received MMSI $\zeta_{m,n}$ matches the MMSI label $\tau_{k,n}$ of its associated PT.

The factorization of $f(\boldsymbol{y}_n, \boldsymbol{y}_{1:n-1}, \boldsymbol{a}_n, \boldsymbol{b}_n | \boldsymbol{z}_{1:n})$ described above differs from (14) essentially by an additional factor for $s = 0$, which accounts for the AIS data. Thus, the associated factor graph equals that in Fig. 3 with an additional SPADA subgraph (cf. Section 2.4) for $s = 0$, which describes the processing of the AIS data. Based on this factor graph, multisensor-multitarget tracking with an inherent fusion of radar measurements and AIS data can be accomplished by means of iterative SPA message passing as discussed in Section 3.3. The resulting beliefs $\tilde{f}(\boldsymbol{x}_{k,n}, r_{k,n}, \tau_{k,n})$ approximate the marginal posterior pdfs $f(\boldsymbol{x}_{k,n}, r_{k,n}, \tau_{k,n} | \boldsymbol{z}_{1:n})$; they are used for target detection and for state estimation according to Section 3.2. We note that the SPA message passing algorithm can be easily extended to multiple DMs along the lines of Section 4.

The above-described SPA-based multitarget tracking algorithm with radar-AIS fusion has been implemented in the context of the EU-funded RANGER project. The algorithm has been successfully used to fuse information from heterogeneous sensors, such as an HFSW radar, a photonic-enhanced multiple-input multiple-output radar, AIS data, and existing coastal radars, in four sea trials (two in France and two in Greece) [63, 64].

6 Integration of Contextual Information

A ship either moves in certain geographical regions known as sea lanes, or it moves more irregularly (e.g., during fishing operations). The knowledge of common sea lanes extracted from historical AIS data can be used as additional contextual information in a tracking algorithm. Consider J_{S} sea lanes indexed by $j \in \{1, 2, \ldots, J_{\mathrm{S}}\}$. The motion of a ship along the jth sea lane can be described by a nearly-constant velocity DM $\mathcal{D}_j$ that is defined as in (17) except that it

uses a *directional* driving process $\boldsymbol{u}_{k,n}^{(j)} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{Q}^{(j)})$ whose standard deviation in the direction of the sea lane, $\sigma_{\parallel}$, is larger than that in the direction orthogonal to the sea lane, $\sigma_{\perp}$ [65]. The covariance matrix of $\boldsymbol{u}_{k,n}^{(j)}$ is $\boldsymbol{Q}^{(j)} = \boldsymbol{G}^{(j)}\boldsymbol{D}\boldsymbol{G}^{(j)\mathrm{T}}$, with $\boldsymbol{D} = \mathrm{diag}\{\sigma_{\perp}^2, \sigma_{\parallel}^2\}$ and $\boldsymbol{G}^{(j)} = \begin{bmatrix} -\cos\delta_j & \sin\delta_j \\ \sin\delta_j & \cos\delta_j \end{bmatrix}$, where δ_j is the angle of the current sea lane direction. In addition, we describe ships that are freely moving off the sea lanes by a nondirectional nearly-constant velocity DM $\mathcal{D}_0$ with equal standard deviation σ_0 for both directions, corresponding to $\boldsymbol{Q}^{(0)} = \sigma_0^2 \mathbf{I}_2$.

Similarly to Section 4, the random variable $\ell_{k,n} \in \{0, 1, \ldots, J_{\mathrm{S}}\}$ selects the DM index j of ship (PT) k at time n; this DM corresponds to some sea lane ($\ell_{k,n} \in \{1, 2, \ldots, J_{\mathrm{S}}\}$) or to free off-lane motion ($\ell_{k,n} = 0$). The evolution of $\ell_{k,n}$ follows a Markov chain model similar to (18). More specifically, we employ a reduced Markov chain as proposed in [65] since the number of sea lanes J_{S} is usually quite large. We define the *active DM index set* of PT k at time n, $\mathcal{J}_{k,n} \subseteq \{0, 1, \ldots, J_{\mathrm{S}}\}$, as the set containing, besides the off-lane index $j = 0$, the indices of those sea lanes that can be reached at time n by PT k. These latter indices depend on the previous state $\boldsymbol{x}_{k,n-1}$, and thus $\mathcal{J}_{k,n}$ depends on $\boldsymbol{x}_{k,n-1}$. We then allow index transitions $\ell_{k,n-1} \to \ell_{k,n}$ only between active indices, i.e., for $\ell_{k,n-1} \in \mathcal{J}_{k,n-1}$ and $\ell_{k,n} \in \mathcal{J}_{k,n}$. Using index transformations $\mathcal{J}_{k,n-1} \to \{0, 1, \ldots, |\mathcal{J}_{k,n-1}| - 1\}$ and $\mathcal{J}_{k,n} \to \{0, 1, \ldots, |\mathcal{J}_{k,n}| - 1\}$, the index transition pmf can be formulated as (cf. (18)) $p(\ell_{k,n} = j | \ell_{k,n-1} = i) = [\boldsymbol{L}_{k,n}]_{i,j}$ for $i \in \mathcal{J}_{k,n-1}$ and $j \in \mathcal{J}_{k,n}$, with a transition matrix $\boldsymbol{L}_{k,n} \in [0, 1]^{|\mathcal{J}_{k,n-1}| \times |\mathcal{J}_{k,n}|}$. This transition matrix depends on k, n, and $\boldsymbol{x}_{k,n-1}$.

Typically, the restriction of the DM index transitions to active sea lanes results in a strong reduction of complexity and also in an improvement of tracking accuracy. The reduced Markov chain is still consistent with our formulation in Section 4. Thus, iterative SPA message passing can be used for multisensor-multitarget tracking aided by a knowledge of sea lanes; the only difference from Section 4 is that the active index set $\mathcal{J}_{k,n}$ and the corresponding transition matrix $\boldsymbol{L}_{k,n}$ have to be determined at each time n and for each PT k. The factor graph/SPA formulation also allows a straightforward integration of online AIS data as discussed in Section 5.

7 Simulated-Data Experiments

In this section, we provide a quantitative performance analysis of the SPA-based multisensor-multitarget tracking algorithm applied to simulated data. More specifically, we present a comparison with multisensor (C)PHD filters (Section 7.1) and an analysis of the effect of radar-AIS information fusion (Section 7.2).

7.1 Comparison with Multisensor (C)PHD Filters

We compare the performance of the SPA-based multisensor-multitarget tracking algorithm described in Section 3 with that of particle implementations of the iterated-corrector (C)PHD (IC-(C)PHD) filter and the partition-based (C)PHD (PB-(C)PHD) filter [66]. The IC-PHD and IC-CPHD filters [31, 32] are straightforward multisensor extensions of the PHD filter [17, 25] and CPHD filter [17, 27, 28], respectively, in which the single-sensor (C)PHD update step is performed sequentially for each sensor. The PB-(C)PHD filter [33] approximates the exact multisensor (C)PHD filter [29, 30]. The update step of the exact multisensor (C)PHD filter involves a summation over all the partitions of the sets of measurements generated by all the sensors, and is therefore computationally infeasible. The PB-(C)PHD filter achieves a significant reduction of complexity by extracting only the most relevant partitions — according to a score function — before performing the update step. Since the partition extraction procedure presupposes a Gaussian-mixture implementation, it is here adapted to a particle-based implementation as done in [41].

We consider a region of interest (ROI) given by $[-5\,\mathrm{km}, 5\,\mathrm{km}] \times [-5\,\mathrm{km}, 5\,\mathrm{km}]$. The ROI is covered by $S = 2$ radar sensors located at

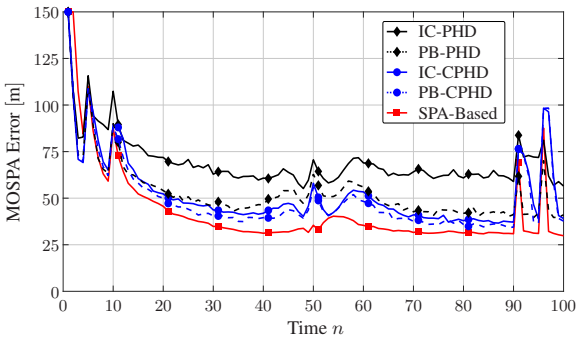

Fig. 6: MOSPA error obtained with the IC-PHD and PB-PHD filters, the IC-CPHD and PB-CPHD filters, and the SPA-based tracking algorithm.

$\boldsymbol{p}^{(1)} = [0, -5\,\mathrm{km}]$ and $\boldsymbol{p}^{(2)} = [5\,\mathrm{km}, 0]$. Three targets, initially uniformly located on a circle of radius 1 km, move toward the center of the ROI with constant velocity 4 m/s. We simulate 100 time steps with time step duration $\Delta T = 5$ s. The three targets appear at times $n = 1$, $n = 5$, and $n = 10$, disappear at times $n = 90$, $n = 95$, and $n = 100$, and cross their trajectories at time $n = 50$. The two radar sensors produce range-bearing measurements modeled as

$$\boldsymbol{z}_{m,n}^{(s)} = \begin{bmatrix} \|\tilde{\boldsymbol{x}}_{k,n} - \boldsymbol{p}^{(s)}\| \\ \phi(\tilde{\boldsymbol{x}}_{k,n}, \boldsymbol{p}^{(s)}) \end{bmatrix} + \boldsymbol{v}_{m,n}^{(s)}, \quad s \in \{1, 2\},$$

where $\tilde{\boldsymbol{x}}_{k,n} \triangleq [x_{1,k,n} \;\; x_{2,k,n}]^{\mathrm{T}}$ is the position of target k, $\phi(\tilde{\boldsymbol{x}}_{k,n}, \boldsymbol{p}^{(s)})$ is the angle of the vector $\tilde{\boldsymbol{x}}_{k,n} - \boldsymbol{p}^{(s)}$, and $\boldsymbol{v}_{m,n}^{(s)} \sim \mathcal{N}\big(\boldsymbol{0}, \mathrm{diag}\{\sigma_{\mathrm{r}}^2, \sigma_{\mathrm{b}}^2\}\big)$ is a sequence of 2D Gaussian random vectors that is iid across m, n, and s. The range and bearing standard deviations are set to $\sigma_{\mathrm{r}} = 80$ m and $\sigma_{\mathrm{b}} = 1^{\circ}$, respectively. The targets are detected by the radar sensors with probability $P_{\mathrm{d}}^{(1)} = P_{\mathrm{d}}^{(2)} = 0.8$. The mean number of false alarms is $\mu^{(1)} = \mu^{(2)} = 2$, and the distribution of false alarms, $f_{\mathrm{FA}}(\boldsymbol{z}_{m,n}^{(s)})$ for $s \in \{1, 2\}$, is uniform on the ROI.

In all the considered tracking algorithms, the evolution of the targets is modeled by a single nearly-constant velocity DM (cf. (17)) with driving process variance $\sigma_{\mathrm{l}}^2 = 0.05^2\,\mathrm{m}^2/\mathrm{s}^4$, and the survival probability — introduced in Section 3.1 for the SPA-based tracking algorithm and also used within the (C)PHD filters — is set to $p_{k,n}^{(\mathrm{s})} = 0.999$. For the SPA-based tracking algorithm, the birth probability is $p_{k,n}^{(\mathrm{b})} = 0.001$, the number of PTs is $K = 15$, the number of particles per PT is 6000, and the detection threshold used to declare the existence of a PT (cf. Section 3.2) is $P_{\mathrm{th}} = 0.8$. For the IC-(C)PHD and PB-(C)PHD filters, the mean number of newborn targets is 0.01 and the number of particles used to represent the PHD of the target states is 30000 [66]. For the PB-(C)PHD filter, the maximum numbers of subsets and partitions used in the partition extraction procedure are 100 and 500, respectively [33].

Fig. 6 presents for each tracking algorithm the Euclidean distance based mean optimal sub-pattern assignment (MOSPA) error with cutoff parameter $c = 150$ m and order $p = 1$ [67], averaged over 500 simulation runs, as a function of time. The MOSPA error is a performance metric that takes into consideration both the effects of incorrectly detected targets and the estimation errors obtained for correctly detected targets. The results in Fig. 6 show that the SPA-based tracking algorithm outperforms both the IC-(C)PHD filter and the PB-(C)PHD filter. This is mainly due to the clustering step in the particle implementations of the IC-(C)PHD and PB-(C)PHD filters, which is needed to extract the target state estimates [66]. Note that the peaks observed at various times are due to target appearance and disappearance.

These results are confirmed by Table 1, which reports the time-on-target (ToT) and track fragmentation (TF) values — averaged over the three targets and the 500 simulation runs — obtained with the different tracking algorithms. The ToT is defined as the percentage of time during which a target is successfully tracked in the sense that the Euclidean distance of its estimated position from its true position

Tracking Algorithm	ToT	TF
IC-PHD	83.4%	11.1
PB-PHD	90.1%	7.3
IC-CPHD	92.8%	5.0
PB-CPHD	94.2%	4.1
SPA-Based	97.2%	1.2

Table 1 Average ToT and TF obtained with the IC-PHD and PB-PHD filters, the IC-CPHD and PB-CPHD filters, and the SPA-based tracking algorithm.

does not exceed the MOSPA cutoff parameter $c = 150$ m. The TF is defined as the number of different partial tracks that are associated with a target within its entire lifetime.

Further simulation results presented in [41] compare the performance of the SPA-based tracking algorithm, of the IC-(C)PHD filter, and of the PB-(C)PHD filter as a function of the number of radar sensors, the probability of detection, and the mean number of false alarms.

7.2 Radar-AIS Information Fusion

In this section, we study the performance of the AIS fusion extension of the SPA-based tracking algorithm presented in Section 5. The simulated scenario is as described in Section 7.1, with the difference that the ROI is now covered also by the AIS. The targets transmit AIS reports containing their ID and noisy position, i.e., the position information is modeled as

$$\boldsymbol{d}_{m,n} = \tilde{\boldsymbol{x}}_{k,n} + \boldsymbol{w}_{m,n},$$

where $\boldsymbol{w}_{m,n} \sim \mathcal{N}(\boldsymbol{0}, \sigma_w^2 \mathbf{I}_2)$ is a sequence of 2D Gaussian random vectors that is iid across m and n, with $\sigma_w = 10$ m. Nominally, each target transmits on average 0.5 reports per time step, but one target does not transmit any reports in the time interval $[70, 90]$.

The performance of the SPA-based tracking-and-fusion algorithm is evaluated for the following fusion strategies: strategy "1-S" only uses the measurements of one radar sensor, strategy "2-S" fuses the measurements of both radar sensors, strategy "1-S + AIS" fuses the measurements of one radar sensor and the AIS reports, and strategy "2-S + AIS" fuses the measurements of both radar sensors and the AIS reports. In Fig. 7, we show for each fusion strategy the Euclidean distance based MOSPA error with cutoff parameter $c = 150$ m and order $p = 1$ [67], averaged over 500 simulation runs, as a function of time. The results in Fig. 7 clearly demonstrate the advantage of fusing multiple sources and the ability of the proposed algorithm to correctly deal with the asynchronous and intermittent AIS reports. The larger MOSPA error in the time interval $[70, 90]$ obtained with the algorithm versions that fuse AIS reports, i.e., 1-S + AIS and 2-S + AIS, is due to the temporary lack of reports from one of the targets. Nevertheless, fusion strategies 1-S + AIS and 2-S + AIS always lead to an overall better tracking accuracy than the strategies that fuse only radar measurements, i.e., 1-S and 2-S.

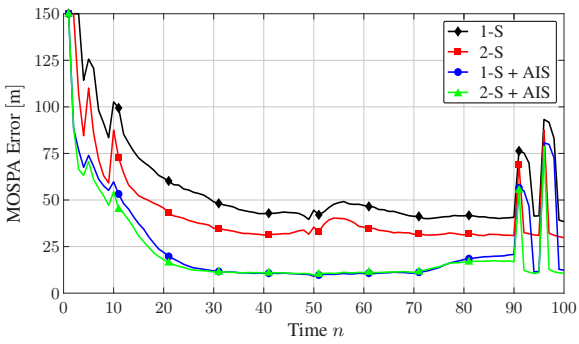

Fig. 7: MOSPA error obtained with the SPA-based tracking-and-fusion algorithm using four different fusion strategies.

Fusion Strategy	ToT	TF
1-S	94.9%	1.5
2-S	97.2%	1.2
1-S + AIS	98.5%	1.1
2-S + AIS	98.6%	1.1

Table 2 Average ToT and TF obtained with the SPA-based tracking-and-fusion algorithm using four different fusion strategies.

Furthermore, it is interesting to note that 1-S + AIS consistently outperforms 2-S. These results are confirmed by Table 2 which reports the ToT and TF values — averaged over the three targets and the 500 simulation runs — obtained with the different fusion strategies.

8 Real-Data Experiments

In this section, we report the results of experiments in which real measurements provided by two HFSW radars — of the WEllen RAdar (WERA) type [68] — and real AIS data are fused by the SPA-based tracking algorithm and its extensions described in Section 4 and Section 5. The two HFSW radars were located on the island of Palmaria (IP) and in San Rossore Park (SRP), on the coast of the Ligurian Sea in Italy. The intersection of the fields-of-view of the two radar sensors defines the ROI. The radar measurements were obtained with a three-dimensional order statistics constant false alarm rate (CFAR) algorithm [7, Section 16.6]. The measurements comprise range, bearing, and range rate. They are modeled in the tracking algorithm as

$$\boldsymbol{z}_{m,n}^{(s)} = \begin{bmatrix} \|\tilde{\boldsymbol{x}}_{k,n} - \boldsymbol{p}^{(s)}\| \\ \phi(\tilde{\boldsymbol{x}}_{k,n}, \boldsymbol{p}^{(s)}) \\ \frac{(\tilde{\boldsymbol{x}}_{k,n} - \boldsymbol{p}^{(s)})^{\mathrm{T}} \dot{\tilde{\boldsymbol{x}}}_{k,n}}{\|\tilde{\boldsymbol{x}}_{k,n} - \boldsymbol{p}^{(s)}\|} \end{bmatrix} + \boldsymbol{v}_{m,n}^{(s)}, \quad s \in \{1, 2\},$$

where $\tilde{\boldsymbol{x}}_{k,n}$, $\boldsymbol{p}^{(s)}$, and $\phi(\tilde{\boldsymbol{x}}_{k,n}, \boldsymbol{p}^{(s)})$ are defined as in Section 7, $\dot{\tilde{\boldsymbol{x}}}_{k,n} \triangleq [\dot{x}_{1,k,n} \;\; \dot{x}_{2,k,n}]^{\mathrm{T}}$, and $\boldsymbol{v}_{m,n}^{(s)} \sim \mathcal{N}\big(\boldsymbol{0}, \mathrm{diag}\{\sigma_{\mathrm{r}}^2, \sigma_{\mathrm{b}}^2, \sigma_{\dot{\mathrm{r}}}^2\}\big)$ is a sequence of three-dimensional Gaussian random vectors that is iid across m, n, and s, with $\sigma_{\mathrm{r}} = 100$ m, $\sigma_{\mathrm{b}} = 1.5^\circ$, and $\sigma_{\dot{\mathrm{r}}}^2 = 0.1$ m/s. The time step duration is $\Delta T = 16.64$ s. The false-alarm pdf $f_{\mathrm{FA}}\big(\boldsymbol{z}_{m,n}^{(s)}\big)$ is assumed uniform on the ROI. The mean number of false alarms is $\mu^{(s)} = 10$. The SPA-based tracking algorithm and its extensions use $K = 80$ PTs, 5000 particles per PT, survival probability $p_{k,n}^{(\mathrm{s})} = 0.999$, birth probability $p_{k,n}^{(\mathrm{b})} = 0.08$, and detection threshold $P_{\mathrm{th}} = 0.65$.

8.1 Radar-AIS Information Fusion

The results shown in Fig. 8 were produced by the SPA-based tracking-and-fusion algorithm that fuses the measurements of the two HFSW radars with the AIS data as described in Section 5. The algorithm uses detection probability $P_{\mathrm{d}}^{(s)} = 0.65$ and a single nearly-constant velocity DM with driving process variance $\sigma_1^2 = 0.005^2$ m^2/s^4. Fig. 8 allows the distinction between estimated trajectories corresponding to a nonzero MMSI label (and, thereby, to the detection of a specific ship identity) and estimated trajectories corresponding to unidentified ships (i.e., ships that do not transmit AIS reports) or corresponding to false tracks. One can see that fusing radar measurements with AIS data makes it possible to successfully identify the respective ship for many of the trajectories and, also, to successfully detect and track ships that do not transmit AIS data.

8.2 DM-Adaptive Multisensor-Multitarget Tracking

Next, to model both nonmaneuvering and maneuvering targets, we consider the DM-adaptive tracking algorithm of Section 4. The algorithm uses two DMs $\mathcal{D}_1$ and $\mathcal{D}_2$ of the nearly-constant velocity

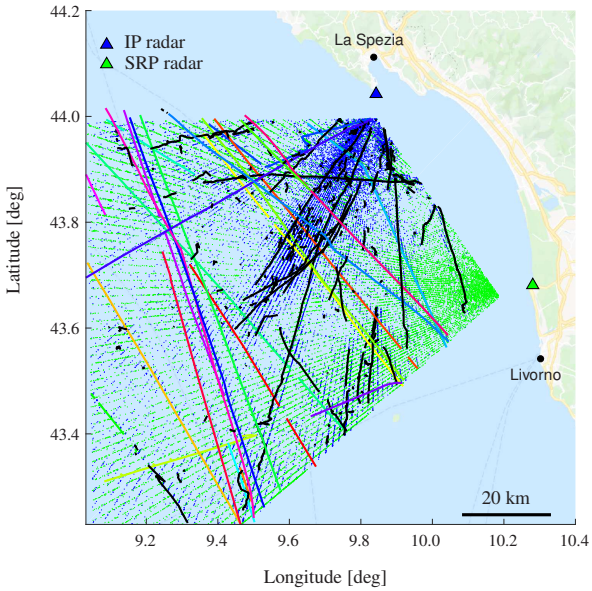

Fig. 8: Estimated ship trajectories obtained with the SPA-based tracking-and-fusion algorithm performing radar-AIS fusion. Colored lines represent estimated trajectories corresponding to ships identified by an MMSI, whereas black lines represent estimated trajectories corresponding to ships not transmitting AIS reports or to false tracks. The measurements provided by the IP and SRP radars are indicated by blue and green dots, respectively, and the positions of the IP and SRP radars by a blue and a green triangle, respectively. (Map courtesy of Google)

type (see (17)) with driving process variances $\sigma_1^2 = 0.001^2$ m^2/s^4 and $\sigma_2^2 = 0.01^2$ m^2/s^4, respectively. This choice of the variances is motivated by the long scan time of the radars: as a rule of thumb, the driving noise variance for maneuvering targets should be larger by two orders of magnitude than that of nonmaneuvering targets. As discussed in Section 4, the evolution of the unknown DM indices $\ell_k \in \{1, 2\}$ is modeled by a Markov chain; the transition probabilities are chosen as $[\boldsymbol{L}]_{1,1} = [\boldsymbol{L}]_{2,2} = 0.985$ and $[\boldsymbol{L}]_{1,2} = [\boldsymbol{L}]_{2,1} = 0.015$. The diagonal values $[\boldsymbol{L}]_{1,1}$ and $[\boldsymbol{L}]_{2,2}$ represent the prior probability of remaining in the same DM, i.e., $\mathcal{D}_1$ or $\mathcal{D}_2$, respectively. The DM-adaptive algorithm is not very sensitive to these values, e.g., $[\boldsymbol{L}]_{1,1}$ and $[\boldsymbol{L}]_{2,2}$ can be chosen between 0.8 and 0.99. Choosing them much higher than the off-diagonal values $[\boldsymbol{L}]_{1,2}$ and $[\boldsymbol{L}]_{2,1}$ increases the accuracy of the state estimates during straight trajectories while still being able to track targets at the beginning of a sharp maneuver or a sudden acceleration (e.g., when $\mathcal{D}_1$ changes to $\mathcal{D}_2$). In addition to the adaptive choice of the DM, the tracking algorithm also adaptively chooses the unknown detection probabilities of the two radar sensors as proposed in [49].

Fig. 9(a) shows the ship trajectories estimated during approximately nine hours. Also shown are "ground truth" trajectories that were determined through linear interpolation of the positions provided by AIS reports. One can observe that the estimated trajectories are generally close to the AIS trajectories. In particular, the detail shown in Fig. 9(b) demonstrates a high estimation accuracy also during sharp maneuvers, even when there are only few radar measurements in the relevant part of the surveillance region (as indicated by the low density of the blue dots in the vicinity of the trajectory). This is an important advantage of the online adaptation of the DM indices and detection probabilities.

9 Conclusions

Among the sensors that can be used for maritime situational awareness, high-frequency surface-wave (HFSW) maritime radars are attractive because of their over-the-horizon field of view. Improving the target detection and tracking performance of HFSW radars

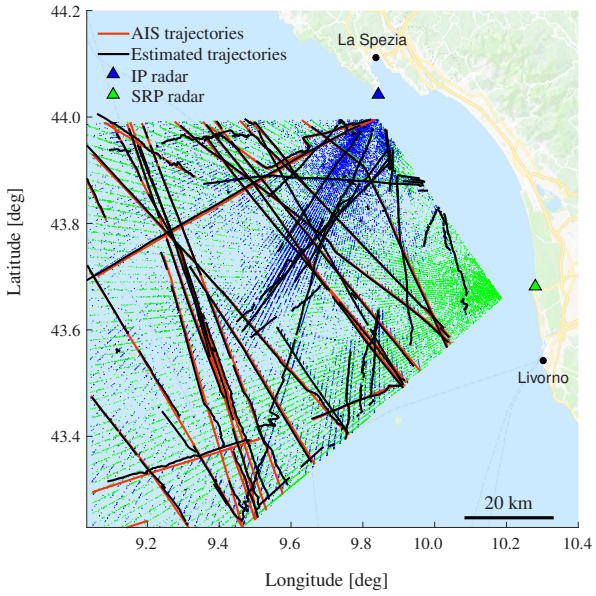

(a)

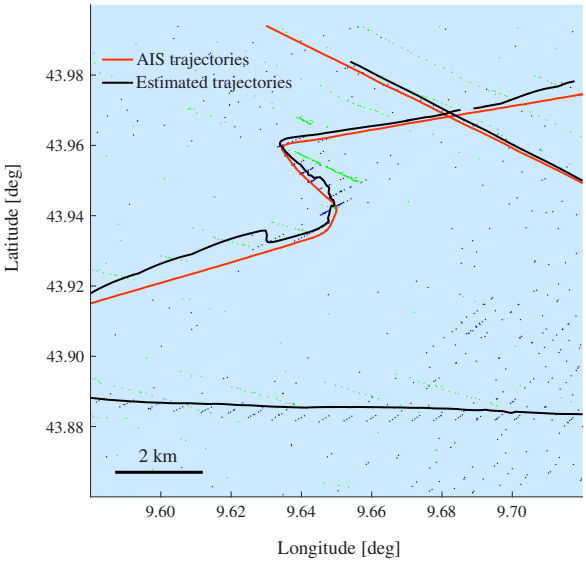

(b)

Fig. 9: Estimated ship trajectories (black lines) obtained with the SPA-based tracking algorithm using measurements provided by the IP and SRP radars (indicated by blue and green dots, respectively): (a) Complete tracking scenario, (b) detail depicting a target maneuver. The positions of the IP and SRP radars are indicated by a blue and green triangle, respectively. Red lines represent "ground truth" ship trajectories determined from AIS reports. (Map courtesy of Google)

requires the development of advanced algorithms for multisensor-multitarget tracking with inherent information fusion. In this article, we surveyed the use of factor graphs and of the iterative sum-product algorithm (SPA) for that purpose. A major advantage of the SPA methodology is its versatility. In particular, SPA-based algorithms facilitate the integration of additional maritime data, such as real-time automatic identification system (AIS) data and contextual information in the form of sea lanes, as well as an online adaptation to dynamic models and other unknown time-varying model parameters. We demonstrated the effectiveness of the proposed SPA-based algorithms through experiments using simulated data as well as real measurements obtained from HFSW radars and real AIS reports.

Suggestions for future work include the development of extensions of the proposed framework and methodology to incorporate

synthetic aperture radar images [69] and ship features (such as size) as additional sources of information, as well as applications of our methodologies to indoor localization [51, 70] and the tracking of aircrafts using over-the-horizon radar systems [71, 72].

10 Acknowledgments

This work was supported in part by the NATO Allied Command Transformation (ACT) under the DKOE project, by the European Research Council (ERC) under grant 700478 (project RANGER) within the Horizon 2020 program, by the Austrian Science Fund (FWF) under grants J 3886-N31 and P 32055-N31, by the Czech Science Foundation (GAČR) under grant 17-19638S, and by the Army Research Office through the MIT Institute for Soldier Nanotechnologies under contract W911NF-13-D-0001.

11 References

1 Gad, A., Farooq, M. 'Data fusion architecture for maritime surveillance'. In: Proc. FUSION-02. (Annapolis, MD, USA, 2002. pp. 448–455
2 Carthel, C., Coraluppi, S., Grignan, P. 'Multisensor tracking and fusion for maritime surveillance'. In: Proc. FUSION-07. (Quebec, Canada, 2007.
3 Nilsson, M., van Laere, J., Ziemke, T., Edlund, J. 'Extracting rules from expert operators to support situation awareness in maritime surveillance'. In: Proc. FUSION-08. (Cologne, Germany, 2008.
4 Perera, L.P., Oliveira, P., Soares, C.G.: 'Maritime traffic monitoring based on vessel detection, tracking, state estimation, and trajectory prediction', *IEEE Trans Intell Transp Syst*, 2012, **13**, (3), pp. 1188–1200
5 Braca, P., Maresca, S., Grasso, R., Bryan, K., Horstmann, J.: 'Maritime surveillance with multiple over-the-horizon HFSW radars: An overview of recent experimentation', *IEEE Trans Aerosp Electron Syst*, 2015, **30**, (12), pp. 4–18
6 Meyer, F., Kropfreiter, T., Williams, J.L., Lau, R.A., Hlawatsch, F., Braca, P., et al.: 'Message passing algorithms for scalable multitarget tracking', *Proc IEEE*, 2018, **106**, (2), pp. 221–259
7 Richards, M.A., Scheer, J.A., Holm, W.A.: 'Principles of Modern Radar: Basic Principles'. (Raleigh, NC, USA: Scitech Publishing, 2010)
8 Tetreault, B.J. 'Use of the automatic identification system (AIS) for maritime domain awareness (MDA)'. In: Proc. MTS/IEEE OCEANS-05. vol. 2. (Washington, D.C., USA, 2005. pp. 1590–1594
9 Fabrizio, G.A.: 'High Frequency Over-the-Horizon Radar: Fundamental Principles, Signal Processing, and Practical Applications'. (New York, NY, USA: McGraw Hill, 2013)
10 Li, M.Q., Xu, J., Zhou, X., Qian, L.C., Long, T., Bian, M.M. 'OTHR highly maneuvering target detection via generalized Radon-Fourier transform'. In: Proc. CIE RADAR-16. (Guangzhou, China, 2016.
11 Ponsford, A.M., Wang, J.: 'A review of high frequency surface wave radar for detection and tracking of ships', *Turk J Elec Eng Comp Sci*, 2010, **18**, (3), pp. 409–428
12 Maresca, S., Braca, P., Horstmann, J., Grasso, R.: 'Maritime surveillance using multiple high-frequency surface-wave radars', *IEEE Trans Geosci Remote Sens*, 2014, **52**, (8), pp. 5056–5071
13 Li, J., Stoica, P.: 'MIMO Radar Signal Processing'. (Hoboken, NJ, USA: Wiley-IEEE Press, 2008)
14 Martin, R.J., Kearney, M.J.: 'Remote sea current sensing using HF radar: An autoregressive approach', *IEEE J Ocean Eng*, 1997, **22**, (1), pp. 151–155
15 Ohlmann, C., White, P., Washburn, L., Emery, B., Terrill, E., Otero, M.: 'Interpretation of coastal HF radar–derived surface currents with high-resolution drifter data', *J Atmospheric Ocean Technol*, 2007, **24**, (4), pp. 666–680
16 Bar.Shalom, Y., Willett, P.K., Tian, X.: 'Tracking and Data Fusion: A Handbook of Algorithms'. (Storrs, CT, USA: Yaakov Bar-Shalom, 2011)
17 Mahler, R.: 'Statistical Multisource-Multitarget Information Fusion'. (Norwood, MA, USA: Artech House, 2007)
18 Greco, M., Stinco, P., Gini, F., Rangaswamy, M.: 'Impact of sea clutter nonstationarity on disturbance covariance matrix estimation and CFAR detector performance', *IEEE Trans Aerosp Electron Syst*, 2010, **46**, (3), pp. 1502–1513
19 Challa, S., Morelande, M.R., Mušicki, D., Evans, R.J.: 'Fundamentals of Object Tracking'. (Cambridge, UK: Cambridge University Press, 2011)
20 Mušicki, D., Evans, R.: 'Joint integrated probabilistic data association: JIPDA', *IEEE Trans Aerosp Electron Syst*, 2004, **40**, (3), pp. 1093–1099
21 Reid, D.B.: 'An algorithm for tracking multiple targets', *IEEE Trans Autom Control*, 1979, **24**, (6), pp. 843–854
22 Danchick, R., Newnam, G.E.: 'A fast method for finding the exact N-best hypotheses for multitarget tracking', *IEEE Trans Aerosp Electron Syst*, 1993, **29**, (2), pp. 555–560
23 Cox, I.J., Hingorani, S.L.: '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*, 1996, **18**, (2), pp. 138–150
24 Danchick, R., Newnam, G.E.: 'Reformulating reid's MHT method with generalised Murty K-best ranked linear assignment algorithm', *IEE Radar Sonar Navig*, 2006, **153**, (1), pp. 13–22
25 Mahler, R.P.S.: 'Multitarget Bayes filtering via first-order multitarget moments', *IEEE Trans Aerosp Electron Syst*, 2003, **39**, (4), pp. 1152–1178
26 Braca, P., Marano, S., Matta, V., Willett, P.: 'Asymptotic efficiency of the PHD in multitarget/multisensor estimation', *IEEE J Sel Topics Signal Process*, 2013, **7**, (3), pp. 553–564
27 Mahler, R.: 'PHD filters of higher order in target number', *IEEE Trans Aerosp Electron Syst*, 2007, **43**, (4), pp. 1523–1543
28 Vo, B.T., Vo, B.N., Cantoni, A.: 'Analytic implementations of the cardinalized probability hypothesis density filter', *IEEE Trans Signal Process*, 2007, **55**, (7), pp. 3553–3567
29 Mahler, R. 'The multisensor PHD filter: I. General solution via multitarget calculus'. In: Proc. SPIE-09. vol. 7336. (Orlando, FL, USA, 2009. pp. 194–205
30 Delande, E., Duflos, E., Heurguier, D., Vanheeghe, P. 'Multi-target PHD filtering: proposition of extensions to the multi-sensor case'. (INRIA, 2010. RR-7337
31 Mahler, R. 'The multisensor PHD filter: II. Erroneous solution via "Poisson magic"'. In: Proc. SPIE-09. vol. 7336. (Orlando, FL, USA, 2009. pp. 182–193
32 Nagappa, S., Clark, D.E. 'On the ordering of the sensors in the iterated-corrector probability hypothesis density (PHD) filter'. In: Proc. SPIE-11. vol. 8050. (Orlando, FL, USA, 2011. pp. 26–28
33 Nannuru, S., Blouin, S., Coates, M., Rabbat, M.: 'Multisensor CPHD filter', *IEEE Trans Aerosp Electron Syst*, 2016, **52**, (4), pp. 1834–1854
34 Williams, J.L.: 'Marginal multi-Bernoulli filters: RFS derivation of MHT, JIPDA and association-based MeMBer', *IEEE Trans Aerosp Electron Syst*, 2015, **51**, (3), pp. 1664–1687
35 Williams, J.L. 'Experiments with graphical model implementations of multiple target multiple-Bernoulli filters'. In: Proc. ISSNIP-11. (Adelaide, SA, Australia, 2011. pp. 532–537
36 Kropfreiter, T., Meyer, F., Hlawatsch, F. 'Sequential Monte Carlo implementation of the track-oriented marginal multi-Bernoulli/Poisson filter'. In: Proc. FUSION-16. (Heidelberg, Germany, 2016. pp. 972–979
37 Meyer, F., Tesei, A., Win, M.Z. 'Localization of multiple sources using time-difference arrival measurements'. In: IEEE ICASSP-17. (New Orleans, LA, USA, 2017. pp. 3151–3155
38 Vo, B.T., Vo, B.N.: 'Labeled random finite sets and multi-object conjugate priors', *IEEE Trans Signal Process*, 2013, **61**, (13), pp. 3460–3475
39 Reuter, S., Vo, B., Vo, B., Dietmayer, K.: 'The labeled multi-Bernoulli filter', *IEEE Trans Signal Process*, 2014, **62**, (12), pp. 3246–3260
40 Benavoli, A., Chisci, L., Farina, A., Immediata, S., Timmoneri, L., Zappa, G.: 'Knowledge-based system for multi-target tracking in a littoral environment', *IEEE Trans Aerosp Electron Syst*, 2006, **42**, (3), pp. 1100–1119
41 Meyer, F., Braca, P., Willett, P., Hlawatsch, F.: 'A scalable algorithm for tracking an unknown number of targets using multiple sensors', *IEEE Trans Signal Process*, 2017, **65**, (13), pp. 3478–3493
42 Kschischang, F.R., Frey, B.J., Loeliger, H.A.: 'Factor graphs and the sum-product algorithm', *IEEE Trans Inf Theory*, 2001, **47**, (2), pp. 498–519
43 Koller, D., Friedman, N.: 'Probabilistic Graphical Models: Principles and Techniques'. (Cambridge, MA, USA: MIT Press, 2009)
44 Frank, A., Smyth, P., Ihler, A.: 'Beyond MAP estimation with the track-oriented multiple hypothesis tracker', *IEEE Trans Signal Process*, 2014, **62**, (9), pp. 2413–2423
45 Williams, J.L., Lau, R.A. 'Data association by loopy belief propagation'. In: Proc. FUSION-10. (Edinburgh, UK, 2010.
46 Williams, J.L., Lau, R.: 'Approximate evaluation of marginal association probabilities with belief propagation', *IEEE Trans Aerosp Electron Syst*, 2014, **50**, (4), pp. 2942–2959
47 Chertkov, M., Kroc, L., Krzakala, F., Vergassola, M., Zdeborová, L.: 'Inference in particle tracking experiments by passing messages between images', *PNAS*, 2010, **107**, (17), pp. 7663–7668
48 Vontobel, P.O.: 'The Bethe permanent of a nonnegative matrix', *IEEE Trans Inf Theory*, 2013, **59**, (3), pp. 1866–1901
49 Soldi, G., Meyer, F., Braca, P., Hlawatsch, F.: 'Self-tuning algorithms for multisensor-multitarget tracking using belief propagation', *IEEE Trans Signal Process*, 2019, **67**, (15), pp. 3922–3937
50 Leitinger, E., Meyer, F., Meissner, P., Witrisal, K., Hlawatsch, F. 'Belief propagation based joint probabilistic data association for multipath-assisted indoor navigation and tracking'. In: Proc. ICL-GNSS-16. (Barcelona, Spain, 2016.
51 Leitinger, E., Meyer, F., Hlawatsch, F., Witrisal, K., Tufvesson, F., Win, M.Z.: 'A belief propagation algorithm for multipath-based SLAM', *IEEE Trans Wireless Commun*, 2019, **18**, (12), pp. 5613–5629
52 Lau, R.A., Williams, J.L. 'Tracking a coordinated group using expectation maximisation'. In: Proc. ISSNIP-13. (Melbourne, Vic, Australia, 2013. pp. 282–287
53 Rahmathullah, A.S., Selvan, R., Svensson, L.: 'A batch algorithm for estimating trajectories of point targets using expectation maximization', *IEEE Trans Signal Process*, 2016, **64**, (18), pp. 4792–4804
54 Lan, H., Pan, Q., Yang, F., Sun, S., Li, L. 'Variational Bayesian approach for joint multitarget tracking of multiple detection systems'. In: Proc. FUSION-16. (Heidelberg, Germany, 2016. pp. 1260–1267
55 Sun, S., Lan, H., Wang, Z., Pan, Q., Zhang, H. 'The application of sum-product algorithm for data association'. In: Proc. FUSION-16. (Heidelberg, Germany, 2016. pp. 416–423
56 Gaglione, D., Braca, P., Soldi, G. 'Belief propagation based AIS/radar data fusion for multi-target tracking'. In: Proc. FUSION-18. (Cambridge, UK, 2018. pp. 2143–2150
57 Soldi, G., Gaglione, D., Meyer, F., Hlawatsch, F., Braca, P., Farina, A., et al. 'Heterogeneous information fusion for multitarget tracking using the sum-product algorithm'. In: Proc. IEEE ICASSP-19. (Brighton, UK, 2019. pp. 5471–5475
58 Poor, H.V.: 'An Introduction to Signal Detection and Estimation'. 2nd ed. (New York, NY, USA: Springer, 1994)
59 Meyer, F., Hlinka, O., Hlawatsch, F.: 'Sigma point belief propagation', *IEEE Signal Process Lett*, 2014, **21**, (2), pp. 145–149
60 Anderson, B.D., Moore, J.B.: 'Optimal Filtering'. (Englewood Cliffs, NJ, USA: Prentice-Hall, 1979)

61 Bar.Shalom, Y., Li, X.R., Kirubarajan, T.: 'Estimation with Applications to Tracking and Navigation'. (New York, NY, USA: Wiley, 2001)
62 ITU Radiocommunication Sector (ITU-R). 'Assignment and use of identities in the maritime mobile service' [Recommendation]. (, 2015
63 Karagiannidis, L., Dres, D., Protopapadakis, E., Lamole, F., Jacquin, F., Rigal, G., et al. 'RANGER: Radars and early warning technologies for long distance maritime surveillance'. In: Proc. 2019 Maritime Situational Awareness Workshop. (Lerici, Italy, 2019.
64 Consortium, E.R.P.. 'RANGER project completes its lifecycle'. (, 2020. Press Release. Available from: `https://ranger-project.eu/information-centre/press-releases-press-clippings/`
65 Vivone, G., Braca, P., Horstmann, J.: 'Knowledge-based multitarget ship tracking for HF surface wave radar systems', *IEEE Trans Geosci Remote Sens*, 2015, **53**, (7), pp. 3931–3949
66 Vo, B.N., Singh, S., Doucet, A.: 'Sequential Monte Carlo methods for multitarget filtering with random finite sets', *IEEE Trans Aerosp Electron Syst*, 2005, **41**, (4), pp. 1224–1245
67 Schuhmacher, D., Vo, B.T., Vo, B.N.: 'A consistent metric for performance evaluation of multi-object filters', *IEEE Trans Signal Process*, 2008, **56**, (8), pp. 3447–3457
68 Gurgel, K.W., Antonischki, G., Essen, H.H., Schlick, T.: 'Wellen Radar (WERA): A new ground-wave HF radar for ocean remote sensing', *Coast Eng*, 1999, **37**, (3), pp. 219–234
69 Franceschetti, G., Lanari, R.: 'Synthetic Aperture Radar Processing'. (Boca Raton, FL, USA: CRC Press, 2016)
70 Win, M.Z., Meyer, F., Liu, Z., Dai, W., Bartoletti, S., Conti, A.: 'Efficient multi-sensor localization for the Internet-of-Things', *IEEE Signal Process Mag*, 2018, **35**, (5), pp. 153–167
71 Papazoglou, M., Krolik, J.L.: 'Matched-field estimation of aircraft altitude from multiple over-the-horizon radar revisits', *IEEE Trans Signal Process*, 1999, **47**, (4), pp. 966–976
72 Anderson, R.H., Krolik, J.L.: 'Track association for over-the-horizon radar with a statistical ionospheric model', *IEEE Trans Signal Process*, 2002, **50**, (11), pp. 2632–2643