



CHALMERS

Chalmers Publication Library

Distributed Localization and Tracking of Mobile Networks Including Noncooperative Objects

This document has been downloaded from Chalmers Publication Library (CPL). It is the author's version of a work that was accepted for publication in:

IEEE Transactions on Signal and Information Processing over Networks (ISSN: 2373-776X)

Citation for the published paper:

Meyer, F. ; Hlinka, O. ; Wymeersch, H. et al. (2016) "Distributed Localization and Tracking of Mobile Networks Including Noncooperative Objects". IEEE Transactions on Signal and Information Processing over Networks, vol. 2(1), pp. 57-71.

<http://dx.doi.org/10.1109/TSIPN.2015.2511920>

Downloaded from: <http://publications.lib.chalmers.se/publication/245086>

Notice: Changes introduced as a result of publishing processes such as copy-editing and formatting may not be reflected in this document. For a definitive version of this work, please refer to the published source. Please note that access to the published version might require a subscription.

Chalmers Publication Library (CPL) offers the possibility of retrieving research publications produced at Chalmers University of Technology. It covers all types of publications: articles, dissertations, licentiate theses, masters theses, conference papers, reports etc. Since 2006 it is the official tool for Chalmers official publication statistics. To ensure that Chalmers research results are disseminated as widely as possible, an Open Access Policy has been adopted. The CPL service is administrated and maintained by Chalmers Library.

(article starts on next page)

Distributed Localization and Tracking of Mobile Networks Including Noncooperative Objects

Florian Meyer, *Member, IEEE*, Ondrej Hlinka, Henk Wymeersch, *Member, IEEE*, Erwin Riegler, *Member, IEEE*, and Franz Hlawatsch, *Fellow, IEEE*

Abstract—We propose a Bayesian method for distributed sequential localization of mobile networks composed of both cooperative agents and noncooperative objects. Our method provides a consistent combination of cooperative self-localization (CS) and distributed tracking (DT). Multiple mobile agents and objects are localized and tracked using measurements between agents and objects and between agents. For a distributed operation and low complexity, we combine particle-based belief propagation with a consensus or gossip scheme. High localization accuracy is achieved through a probabilistic information transfer between the CS and DT parts of the underlying factor graph. Simulation results demonstrate significant improvements in both agent self-localization and object localization performance compared to separate CS and DT, and very good scaling properties with respect to the numbers of agents and objects.

Index Terms—Agent network, belief propagation, consensus, cooperative localization, distributed estimation, distributed tracking, factor graph, gossip, message passing, sensor network.

I. INTRODUCTION

A. Background and State of the Art

Cooperative self-localization (CS) [1], [2] and distributed tracking (DT) [3] are key signal processing tasks in decentralized agent networks. Applications include surveillance [4], environmental and agricultural monitoring [5], robotics [6], and pollution source localization [7]. In CS, each agent measures quantities related to the location of neighboring agents relative to its own location. By cooperating with other agents, it is able to estimate its own location. In DT, the measurements performed by the agents are related to the locations (or, more generally, states) of noncooperative objects to be tracked. At each agent, estimates of the object states are cooperatively calculated from all agent measurements. CS and DT are related since, ideally, an agent needs to know its own location to be able to contribute to DT. This relation

Final manuscript, December 2, 2015. To appear in the IEEE Transactions on Signal and Information Processing over Networks.

F. Meyer and F. Hlawatsch are with the Institute of Telecommunications, Vienna University of Technology, 1040 Vienna, Austria (email: {fmeyer, hlawatsch}@nt.tuwien.ac.at). O. Hlinka is with RobArt GmbH, 4020 Linz, Austria (e-mail: ondrej.hlinka@robart.cc). H. Wymeersch is with the Department of Signals and Systems, Chalmers University of Technology, Gothenburg 41296, Sweden (email: henk.wymeersch@ieec.org). E. Riegler is with the Department of Information Technology and Electrical Engineering, ETH Zurich, 8092 Zurich, Switzerland (email: eriegler@nari.ee.ethz.ch). This work was supported by the FWF under Grants S10603-N13 and P27370-N30, by the WWTF under Grant ICT10-066 (NOWIRE), and by the European Commission under ERC Grant No. 258418 (COOPNET), the Newcom# Network of Excellence in Wireless Communications, and the National Sustainability Program (Grant LO1401). This work was partly presented at the 46th Asilomar Conference on Signals, Systems and Computers, Pacific Grove, CA, Nov. 2012.

motivates the combined CS-DT method proposed in this paper, which achieves improved performance through a probabilistic information transfer between CS and DT.

For CS of static agents (hereafter termed “static CS”), the nonparametric belief propagation (BP) algorithm has been proposed in [8]. BP schemes are well suited to CS because their complexity scales only linearly with the number of agents and, under mild assumptions, a distributed implementation is easily obtained. In [2], a distributed BP message passing algorithm for CS of mobile agents (hereafter termed “dynamic CS”) is proposed. A message passing algorithm based on the mean field approximation is presented for static CS in [9]. In [10], [11], nonparametric BP is extended to dynamic CS and combined with a parametric message representation. In [12], a particle-based BP method using a Gaussian belief approximation is proposed. The low-complexity method for dynamic CS presented in [13] is based on the Bayesian filter and a linearized measurement equation. Another low-complexity CS method with low communication requirements is sigma point BP [14]. In [15], a censoring scheme for sigma point BP is proposed to further reduce communications.

For DT, various distributed recursive estimation methods are available [16]–[20]. Distributed particle filters [18] are especially attractive since they are suited to nonlinear, non-Gaussian systems. In particular, in the distributed particle filters proposed in [19], [20], consensus algorithms are used to compute global particle weights reflecting the measurements of all agents. For DT of an unknown, possibly time-varying number of objects in the presence of object-to-measurement association uncertainty, methods based on random finite sets are proposed in [21]–[23].

In the framework of *simultaneous localization and tracking* (SLAT), static agents track a noncooperative object and localize themselves, using measurements of the distances between each agent and the object [24]. In contrast to dynamic CS, measurements between agents are only used for initialization. A centralized particle-based SLAT method using BP is proposed in [25]. Distributed SLAT methods include a technique using a Bayesian filter and communication via a junction tree [26], iterative maximum likelihood methods [27], variational filtering [28], and particle-based BP [29].

B. Contributions and Paper Organization

We propose a method for distributed localization and tracking of cooperative agents and noncooperative objects in wireless networks, using measurements between agents and objects and between agents. This method, for the first time, provides

a consistent combination of CS and DT in decentralized agent networks where the agents and objects may be mobile. To the best of our knowledge, it is the first method for simultaneous CS and DT in a dynamic setting. It is different from SLAT methods in that the agents may be mobile and measurements between the agents are used also during runtime. A key feature of our method is a probabilistic information transfer between the CS and DT stages, which allows uncertainties in one stage to be taken into account by the other stage and thereby can improve the performance of both stages.

Contrary to the multitarget tracking literature [30], [31], we assume that the number of objects is known and the objects can be identified by the agents. Even with this assumption, the fact that the agents may be mobile and their states are unknown causes the object localization problem to be more challenging than in the setting of static agents with known states. This is because the posterior distributions of the object and agent states are coupled through recurrent pairwise measurements, and thus all these states should be estimated *jointly* and *sequentially*. This joint, sequential estimation is performed quite naturally through our factor graph formulation of the entire estimation problem and the use of BP message passing. In addition, BP message passing facilitates a distributed implementation and exhibits very good scalability in the numbers of agents and objects. We also present a new particle-based implementation of BP message passing that is less complex than conventional nonparametric BP [8], [10].

Our method is an extension of BP-based CS [2] to include noncooperative objects. This extension is nontrivial because, contrary to pure CS, the communication and measurement topologies (graphs) do not match. Indeed, because the objects do not communicate, certain messages needed to calculate the object beliefs are not available at the agents. The proposed method employs a consensus scheme [19] for a distributed calculation of these messages. The resulting combination of BP and consensus may also be useful in other distributed inference problems involving noncooperative objects.

This paper is organized as follows. The system model is described in Section II. A BP message passing scheme for joint CS and DT is developed in Section III, and a particle-based implementation of this scheme in Section IV. A distributed localization-and-tracking algorithm that combines particle-based BP and consensus is presented in Section V. The algorithm's communication requirements and delay are analyzed in Section VI. Finally, simulation results are presented in Section VII. This paper extends our previous work in [32], [33] in the following respects: we consider multiple objects, introduce a new low-complexity message multiplication scheme, provide an analysis of communication requirements and delay, present a numerical comparison with particle filtering and a numerical analysis of scaling properties, and demonstrate performance gains over separate CS and DT in a static scenario (in addition to two dynamic scenarios).




Fig. 1. Network with cooperative agents and noncooperative objects. Also shown are the sets $\mathcal{M}_{l,n}$, $\mathcal{M}_{l,n}^A$, and $\mathcal{M}_{l,n}^O$ for a specific agent l , the set $\mathcal{C}_{l,n}$ for a specific agent l' , and the set $\mathcal{A}_{m,n}$ for a specific object m .

II. SYSTEM MODEL AND PROBLEM STATEMENT

A. System Model

We consider a decentralized network of cooperative agents and noncooperative objects as shown in Fig. 1. We denote by $\mathcal{A} \subseteq \mathbb{N}$ the set of agents, by $\mathcal{O} \subseteq \mathbb{N}$ the set of objects, and by $\mathcal{E} \triangleq \mathcal{A} \cup \mathcal{O}$ the set of all entities (agents or objects). We use the indices $k \in \mathcal{E}$, $l \in \mathcal{A}$, and $m \in \mathcal{O}$ to denote a generic entity, an agent, and an object, respectively. The numbers of agents and objects are assumed known. The objects are noncooperative in that they do not communicate, do not perform computations, and do not actively perform any measurements. The *state* of entity $k \in \mathcal{E}$ at time $n \in \{0, 1, \dots\}$, denoted $\mathbf{x}_{k,n}$, consists of the current location and, possibly, motion parameters such as velocity [34]. The states evolve according to

$$\mathbf{x}_{k,n} = g(\mathbf{x}_{k,n-1}, \mathbf{u}_{k,n}), \quad k \in \mathcal{E}, \quad (1)$$

where $\mathbf{u}_{k,n}$ denotes driving noise with probability density function (pdf) $f(\mathbf{u}_{k,n})$. The statistical relation between $\mathbf{x}_{k,n-1}$ and $\mathbf{x}_{k,n}$ defined by (1) can also be described by the state-transition pdf $f(\mathbf{x}_{k,n} | \mathbf{x}_{k,n-1})$.

The communication and measurement topologies are described by sets $\mathcal{C}_{l,n}$ and $\mathcal{M}_{l,n}$ as follows. Agent $l \in \mathcal{A}$ is able to communicate with agent $l' \in \mathcal{A}$ if $l' \in \mathcal{C}_{l,n} \subseteq \mathcal{A} \setminus \{l\}$. Communication is symmetric, i.e., $l' \in \mathcal{C}_{l,n}$ implies $l \in \mathcal{C}_{l',n}$. Furthermore, agent $l \in \mathcal{A}$ acquires a measurement $\mathbf{y}_{l,k;n}$ relative to agent or object $k \in \mathcal{E}$ if $k \in \mathcal{M}_{l,n} \subseteq \mathcal{E} \setminus \{l\}$. Note that $\mathcal{C}_{l,n}$ consists of all agents that can communicate with agent $l \in \mathcal{A}$, and $\mathcal{M}_{l,n}$ consists of all entities measured by agent $l \in \mathcal{A}$. Thus, there are actually two networks: one is defined by the *communication graph*, i.e., by $\mathcal{C}_{l,n}$ for $l \in \mathcal{A}$, and involves only agents; the other is defined by the *measurement graph*, i.e., by $\mathcal{M}_{l,n}$ for $l \in \mathcal{A}$, and involves agents and objects. The communication graph is assumed to be connected. We also define $\mathcal{M}_{l,n}^A \triangleq \mathcal{M}_{l,n} \cap \mathcal{A}$ and $\mathcal{M}_{l,n}^O \triangleq \mathcal{M}_{l,n} \cap \mathcal{O}$, i.e., the subsets of $\mathcal{M}_{l,n}$ containing only agents and only objects, respectively, and $\mathcal{A}_{m,n} \triangleq \{l \in \mathcal{A} | m \in \mathcal{M}_{l,n}^O\}$, i.e., the set of agents that acquire measurements of object m . Note that $m \in \mathcal{M}_{l,n}^O$ if and only if $l \in \mathcal{A}_{m,n}$. We assume that $\mathcal{M}_{l,n}^A \subseteq \mathcal{C}_{l,n}$, i.e., if agent l acquires a measurement relative to agent l' , it is able to communicate with agent l' . The sets $\mathcal{C}_{l,n}$ etc. may be time-dependent. An example of communication and measurement topologies is given in Fig. 1.

We consider “pairwise” measurements $\mathbf{y}_{l,k;n}$ that depend on the states $\mathbf{x}_{l,n}$ and $\mathbf{x}_{k,n}$ according to

$$\mathbf{y}_{l,k;n} = h(\mathbf{x}_{l,n}, \mathbf{x}_{k,n}, \mathbf{v}_{l,k;n}), \quad l \in \mathcal{A}, k \in \mathcal{M}_{l,n}. \quad (2)$$

Here, $\mathbf{v}_{l,k;n}$ is measurement noise with pdf $f(\mathbf{v}_{l,k;n})$. An example is the scalar “noisy distance” measurement

$$y_{l,k;n} = \|\tilde{\mathbf{x}}_{l,n} - \tilde{\mathbf{x}}_{k,n}\| + v_{l,k;n}, \quad (3)$$

where $\tilde{\mathbf{x}}_{k,n}$ represents the location of entity k (this is part of the state $\mathbf{x}_{k,n}$). The statistical dependence of $\mathbf{y}_{l,k;n}$ on $\mathbf{x}_{l,n}$ and $\mathbf{x}_{k,n}$ is described by the local likelihood function $f(\mathbf{y}_{l,k;n}|\mathbf{x}_{l,n}, \mathbf{x}_{k,n})$. We denote by $\mathbf{x}_n \triangleq (\mathbf{x}_{k,n})_{k \in \mathcal{E}}$ and $\mathbf{y}_n \triangleq (\mathbf{y}_{l,k;n})_{l \in \mathcal{A}, k \in \mathcal{M}_{l,n}}$ the vectors of, respectively, all states and measurements at time n . Furthermore, we define $\mathbf{x}_{1:n} \triangleq (\mathbf{x}_1^T \cdots \mathbf{x}_n^T)^T$ and $\mathbf{y}_{1:n} \triangleq (\mathbf{y}_1^T \cdots \mathbf{y}_n^T)^T$.

B. Assumptions

We will make the following commonly used assumptions, which are reasonable in many practical scenarios [2].

(A1) All agent and object states are independent *a priori* at time $n=0$, i.e., $f(\mathbf{x}_0) = \prod_{k \in \mathcal{E}} f(\mathbf{x}_{k,0})$.

(A2) All agents and objects move according to a memory-less walk, i.e., $f(\mathbf{x}_{1:n}) = f(\mathbf{x}_0) \prod_{n'=1}^n f(\mathbf{x}_{n'}|\mathbf{x}_{n'-1})$.

(A3) The state transitions of the various agents and objects are independent, i.e., $f(\mathbf{x}_n|\mathbf{x}_{n-1}) = \prod_{k \in \mathcal{E}} f(\mathbf{x}_{k,n}|\mathbf{x}_{k,n-1})$.

(A4) The current measurements \mathbf{y}_n are conditionally independent, given the current states \mathbf{x}_n , of all the other states and of all past and future measurements, i.e., $f(\mathbf{y}_n|\mathbf{x}_{0:\infty}, \mathbf{y}_{1:n-1}, \mathbf{y}_{n+1:\infty}) = f(\mathbf{y}_n|\mathbf{x}_n)$.

(A5) The current states \mathbf{x}_n are conditionally independent of all past measurements, $\mathbf{y}_{1:n-1}$, given the previous states \mathbf{x}_{n-1} , i.e., $f(\mathbf{x}_n|\mathbf{x}_{n-1}, \mathbf{y}_{1:n-1}) = f(\mathbf{x}_n|\mathbf{x}_{n-1})$.

(A6) The measurements $\mathbf{y}_{l,k;n}$ and $\mathbf{y}_{l',k';n}$ are conditionally independent given \mathbf{x}_n unless $(l, k) = (l', k')$, and each measurement $\mathbf{y}_{l,k;n}$ depends only on the states $\mathbf{x}_{l,n}$ and $\mathbf{x}_{k,n}$. Together with (A4), this leads to the following factorization of the “total” likelihood function: $f(\mathbf{y}_{1:n}|\mathbf{x}_{1:n}) = \prod_{n'=1}^n \prod_{l \in \mathcal{A}} \prod_{k \in \mathcal{M}_{l,n'}} f(\mathbf{y}_{l,k;n'}|\mathbf{x}_{l,n'}, \mathbf{x}_{k,n'})$.

We also assume that the objects can be identified by the agents, i.e., object-to-measurement associations are known. This requires a certain degree of coordination: for example, in the case of received signal strength measurements, the objects have to use different frequency bands or different time slots; if ultra-wideband (UWB) signals are used, the identities of the UWB radios mounted on the objects have to be known. (We note that BP-based methods for multitarget tracking in the presence of object-to-measurement association uncertainty were recently proposed in [35] and [36]; however, these methods are not distributed.) Furthermore, we assume that each agent $l \in \mathcal{A}$ knows the functional forms of its own state-transition pdf and initial state pdf as well as of those of all objects, i.e., $f(\mathbf{x}_{k,n}|\mathbf{x}_{k,n-1})$ and $f(\mathbf{x}_{k,0})$ for $k \in \{l\} \cup \mathcal{O}$. Finally, all prior location and motion information is available in one global reference frame, and the internal clocks of all agents are synchronous (see [37]–[39] for distributed clock synchronization algorithms).

C. Problem Statement

The task we consider is as follows: Each agent $l \in \mathcal{A}$ estimates its own state $\mathbf{x}_{l,n}$ and all object states $\mathbf{x}_{m,n}$, $m \in \mathcal{O}$ from the *entire measurement vector* $\mathbf{y}_{1:n} = (\mathbf{y}_{l',k;n'})_{l' \in \mathcal{A}, k \in \mathcal{M}_{l',n'}, n' \in \{1, \dots, n\}}$, i.e., from the pairwise measurements between the agents and between the agents and objects up to time n . This is to be achieved using only communication with the “neighbor agents” as defined by $\mathcal{C}_{l,n}$, and without transmitting measurements between agents.

In this formulation of the estimation task, compared to pure CS of cooperative agents (e.g., [2]) or pure DT of noncooperative objects (e.g., [20]), the measurement set is extended in that it includes also the respective other measurements—i.e., the measurements between agents and objects for agent state estimation and those between agents in object state estimation. This explains why the proposed algorithm is able to outperform separate CS and DT. In fact, by using all present and past measurements available throughout the entire network, the inherent coupling between the CS and DT tasks can be exploited for improved performance.

III. BP MESSAGE PASSING SCHEME

In this section, we describe a BP message passing scheme for the joint CS-DT problem. A particle-based implementation of this scheme will be presented in Section IV, and the final distributed algorithm will be developed in Section V.

For estimating the agent or object state $\mathbf{x}_{k,n}$, $k \in \mathcal{E}$ from $\mathbf{y}_{1:n}$, a popular Bayesian estimator is the minimum mean-square error (MMSE) estimator given by [40]

$$\hat{\mathbf{x}}_{k,n}^{\text{MMSE}} \triangleq \int \mathbf{x}_{k,n} f(\mathbf{x}_{k,n}|\mathbf{y}_{1:n}) d\mathbf{x}_{k,n}. \quad (4)$$

This estimator involves the “marginal” posterior pdf $f(\mathbf{x}_{k,n}|\mathbf{y}_{1:n})$, $k \in \mathcal{E}$, which can be obtained by marginalizing the “joint” posterior pdf $f(\mathbf{x}_{1:n}|\mathbf{y}_{1:n})$. However, direct marginalization of $f(\mathbf{x}_{1:n}|\mathbf{y}_{1:n})$ is infeasible because it relies on nonlocal information and involves integration in spaces whose dimension grows with time and network size. This problem can be addressed by using a particle-based, distributed BP scheme that takes advantage of the temporal and spatial independence structure of $f(\mathbf{x}_{1:n}|\mathbf{y}_{1:n})$ and avoids explicit integration. The independence structure corresponds to the following factorization of $f(\mathbf{x}_{1:n}|\mathbf{y}_{1:n})$, which is obtained by using Bayes’ rule and assumptions (A1)–(A6):

$$\begin{aligned} f(\mathbf{x}_{1:n}|\mathbf{y}_{1:n}) &\propto \left(\prod_{k \in \mathcal{E}} f(\mathbf{x}_{k,0}) \right) \\ &\times \prod_{n'=1}^n \left(\prod_{k_1 \in \mathcal{E}} f(\mathbf{x}_{k_1,n'}|\mathbf{x}_{k_1,n'-1}) \right) \\ &\times \prod_{l \in \mathcal{A}} \prod_{k_2 \in \mathcal{M}_{l,n'}} f(\mathbf{y}_{l,k_2;n'}|\mathbf{x}_{l,n'}, \mathbf{x}_{k_2,n'}). \end{aligned} \quad (5)$$

Here, \propto denotes equality up to a constant normalization factor.

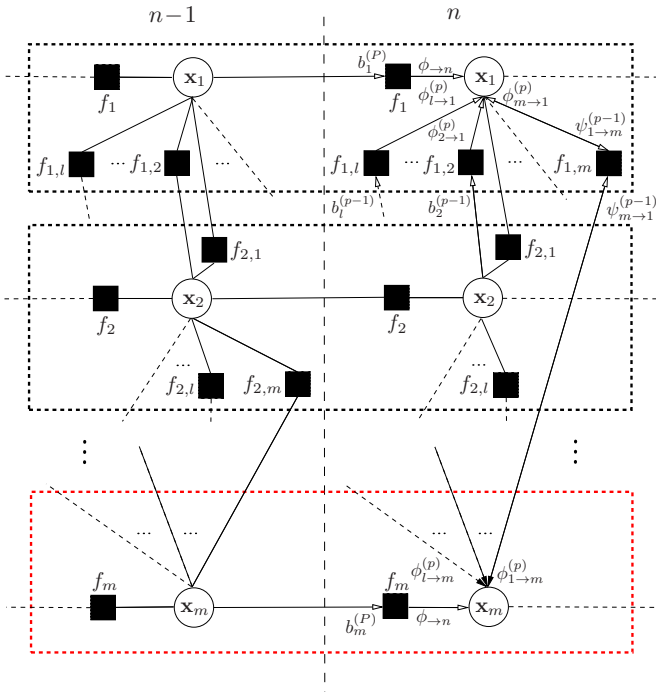


Fig. 2. Factor graph showing the states of agents $l = 1, 2$ and of an object m at time instants $n-1$ and n , assuming $2 \in \mathcal{C}_{1,n}$ and $m \in \mathcal{M}_{1,n}^{\mathcal{O}} \cup \mathcal{M}_{2,n}^{\mathcal{O}}$. Time indices are omitted for simplicity. The short notation $f_k \triangleq f(\mathbf{x}_{k,n'} | \mathbf{x}_{k,n'-1})$, $f_{l,k} \triangleq f(\mathbf{y}_{l,k;n'} | \mathbf{x}_{l,n'}, \mathbf{x}_{k,n'})$, $b_k^{(p)} \triangleq b^{(p)}(\mathbf{x}_{k,n'})$, $\phi_{k \rightarrow l}^{(p)} \triangleq \phi_{k \rightarrow l}^{(p)}(\mathbf{x}_{k,n'})$, etc. (for $n' \in \{1, \dots, n\}$) is used. The upper two (black) dotted boxes correspond to the CS part; the bottom (red) dotted box corresponds to the DT part. Edges between black dotted boxes imply communication between agents. Only messages and beliefs involved in the computation of $b^{(p)}(\mathbf{x}_{1,n})$ and $b^{(p)}(\mathbf{x}_{m,n})$ are shown. Edges with non-filled arrowheads depict particle-based messages and beliefs, while edges with filled arrowheads depict messages involved in the consensus scheme.

A. The Message Passing Scheme

The proposed BP message passing scheme uses the sum-product algorithm [41] to produce approximate marginal posterior pdfs (“beliefs”) $b(\mathbf{x}_{k,n}) \approx f(\mathbf{x}_{k,n} | \mathbf{y}_{1:n})$, $k \in \mathcal{E}$. This is based on the factor graph [42] corresponding to the factorization of $f(\mathbf{x}_{1:n} | \mathbf{y}_{1:n})$ in (5), which is shown in Fig. 2. The factor graph contains variable nodes and factor nodes, which are depicted in Fig. 2 as circles and squares, respectively. Because the factor graph contains loops, BP schemes provide only an approximate marginalization. However, the resulting beliefs have been observed to be quite accurate in many applications [2], [41], [43]. In factor graphs with loops, the BP scheme becomes iterative, and there exist different orders in which messages can be computed [2], [41], [43]. Here, we choose an order that enables real-time processing and facilitates a distributed implementation. More specifically, the belief of agent or object node $l \in \mathcal{A}$ or $m \in \mathcal{O}$ at time n and message passing iteration $p \in \{1, \dots, P\}$ is given by

$$b^{(p)}(\mathbf{x}_{l,n}) \propto \phi_{\rightarrow n}(\mathbf{x}_{l,n}) \prod_{k \in \mathcal{M}_{l,n}} \phi_{k \rightarrow l}^{(p)}(\mathbf{x}_{l,n}), \quad l \in \mathcal{A} \quad (6)$$

$$b^{(p)}(\mathbf{x}_{m,n}) \propto \phi_{\rightarrow n}(\mathbf{x}_{m,n}) \prod_{l \in \mathcal{A}_{m,n}} \phi_{l \rightarrow m}^{(p)}(\mathbf{x}_{m,n}), \quad m \in \mathcal{O}, \quad (7)$$

with the prediction message

$$\phi_{\rightarrow n}(\mathbf{x}_{k,n}) = \int f(\mathbf{x}_{k,n} | \mathbf{x}_{k,n-1}) b^{(P)}(\mathbf{x}_{k,n-1}) d\mathbf{x}_{k,n-1}, \quad k \in \mathcal{E} \quad (8)$$

and the measurement messages

$$\phi_{k \rightarrow l}^{(p)}(\mathbf{x}_{l,n}) = \begin{cases} \int f(\mathbf{y}_{l,k;n} | \mathbf{x}_{l,n}, \mathbf{x}_{k,n}) b^{(p-1)}(\mathbf{x}_{k,n}) d\mathbf{x}_{k,n}, & k \in \mathcal{M}_{l,n}^{\mathcal{A}}, l \in \mathcal{A} \\ \int f(\mathbf{y}_{l,k;n} | \mathbf{x}_{l,n}, \mathbf{x}_{k,n}) \psi_{k \rightarrow l}^{(p-1)}(\mathbf{x}_{k,n}) d\mathbf{x}_{k,n}, & k \in \mathcal{M}_{l,n}^{\mathcal{O}}, l \in \mathcal{A} \end{cases} \quad (9)$$

and

$$\phi_{l \rightarrow m}^{(p)}(\mathbf{x}_{m,n}) = \int f(\mathbf{y}_{l,m;n} | \mathbf{x}_{l,n}, \mathbf{x}_{m,n}) \psi_{l \rightarrow m}^{(p-1)}(\mathbf{x}_{l,n}) d\mathbf{x}_{l,n}, \quad l \in \mathcal{A}_{m,n}, m \in \mathcal{O}. \quad (10)$$

Here, $\psi_{m \rightarrow l}^{(p-1)}(\mathbf{x}_{m,n})$ and $\psi_{l \rightarrow m}^{(p-1)}(\mathbf{x}_{l,n})$ (constituting the “extrinsic information”) are given by

$$\psi_{m \rightarrow l}^{(p-1)}(\mathbf{x}_{m,n}) = \phi_{\rightarrow n}(\mathbf{x}_{m,n}) \prod_{l' \in \mathcal{A}_{m,n} \setminus \{l\}} \phi_{l' \rightarrow m}^{(p-1)}(\mathbf{x}_{m,n}) \quad (11)$$

$$\psi_{l \rightarrow m}^{(p-1)}(\mathbf{x}_{l,n}) = \phi_{\rightarrow n}(\mathbf{x}_{l,n}) \prod_{k \in \mathcal{M}_{l,n} \setminus \{m\}} \phi_{k \rightarrow l}^{(p-1)}(\mathbf{x}_{l,n}). \quad (12)$$

The messages and beliefs involved in calculating $b^{(p)}(\mathbf{x}_{l,n})$ and $b^{(p)}(\mathbf{x}_{m,n})$ are shown in Fig. 2. We note that messages entering or leaving an object variable node in Fig. 2 do not imply that there occurs any communication involving objects.

B. Discussion

According to (6) and (7), the agent beliefs $b^{(p)}(\mathbf{x}_{l,n})$ and object beliefs $b^{(p)}(\mathbf{x}_{m,n})$ involve the product of all the messages passed to the corresponding variable node l and m , respectively. Similarly, according to (11) and (12), the extrinsic informations $\psi_{m \rightarrow l}^{(p-1)}(\mathbf{x}_{m,n})$ and $\psi_{l \rightarrow m}^{(p-1)}(\mathbf{x}_{l,n})$ involve the product of all the messages passed to the corresponding variable node m and l , respectively, except the message of the receiving factor node $f(\mathbf{y}_{l,m;n} | \mathbf{x}_{l,n}, \mathbf{x}_{m,n})$. Furthermore, according to (9), the extrinsic information $\psi_{k \rightarrow l}^{(p-1)}(\mathbf{x}_{k,n})$ passed from variable node $\mathbf{x}_{k,n}$ to factor node $f(\mathbf{y}_{k,l;n} | \mathbf{x}_{k,n}, \mathbf{x}_{l,n})$ is used for calculating the message $\phi_{k \rightarrow l}^{(p)}(\mathbf{x}_{l,n})$ passed from that factor node to the other adjacent variable node, $\mathbf{x}_{l,n}$. A similar discussion applies to (10) and $\psi_{l \rightarrow m}^{(p-1)}(\mathbf{x}_{l,n})$.

Two remarks are in order. First, for low complexity, communication requirements, and latency, messages are sent only forward in time and iterative message passing is performed for each time individually. As a consequence, the message (extrinsic information) from variable node $\mathbf{x}_{k,n-1}$ to factor node $f(\mathbf{x}_{k,n} | \mathbf{x}_{k,n-1})$ equals the belief $b^{(P)}(\mathbf{x}_{k,n-1})$ (see (8)), and $\phi_{\rightarrow n}(\mathbf{x}_{k,n})$ in (8) (for n fixed) remains unchanged during all message passing iterations. Second, for any $k \in \mathcal{A}$, as no information from the factor node $f(\mathbf{y}_{l,k;n} | \mathbf{x}_{l,n}, \mathbf{x}_{k,n})$ is used in the calculation of $b^{(p-1)}(\mathbf{x}_{k,n})$, according to (6) and (9), $b^{(p-1)}(\mathbf{x}_{k,n})$ in (9) is used as the extrinsic information passed to the factor node $f(\mathbf{y}_{l,k;n} | \mathbf{x}_{l,n}, \mathbf{x}_{k,n})$. A similar message computation order is used in the SPAWN algorithm for CS [2], [10]. This order significantly reduces the computational




Fig. 3. Block diagram of (a) separate CS and DT and (b) the proposed scheme, with $b_{A,n}^{(P)} \triangleq \{b^{(P)}(\mathbf{x}_{l,n})\}_{l \in \mathcal{A}}$, $b_{O,n}^{(P)} \triangleq \{b^{(P)}(\mathbf{x}_{m,n})\}_{m \in \mathcal{O}}$, $\psi_{A \rightarrow O,n}^{(p)} \triangleq \{\psi_{l \rightarrow m}^{(p)}(\mathbf{x}_{l,n})\}_{l \in \mathcal{A}, m \in \mathcal{O}}$, $\psi_{O \rightarrow A,n}^{(p)} \triangleq \{\psi_{m \rightarrow l}^{(p)}(\mathbf{x}_{m,n})\}_{m \in \mathcal{M}_{l,n}^{\mathcal{O}}, l \in \mathcal{A}}$ and $\hat{\mathbf{x}}_{A,n}^{(P)} \triangleq (\hat{\mathbf{x}}_{l,n}^{(P)})_{l \in \mathcal{A}}$. In separate CS and DT, the final agent state estimates $\hat{\mathbf{x}}_{A,n}^{(P)}$ are transferred from CS to DT. In the proposed scheme, probabilistic information (the extrinsic informations $\psi_{A \rightarrow O,n}^{(p)}$ and $\psi_{O \rightarrow A,n}^{(p)}$) is transferred between CS and DT at each message passing iteration p .

complexity since it avoids the computation of extrinsic informations among agent variable nodes. It also reduces the amount of communication between agents because beliefs passed between agents can be broadcast, whereas the exchange of extrinsic information would require separate point-to-point communications between agents [2], [10].

Contrary to classical sequential Bayesian filtering [44], which only exploits the temporal conditional independence structure of the estimation problem, the proposed BP scheme (6)–(11) also exploits the spatial conditional independency structure. In fact, increasing the number of agents or objects leads to additional variable nodes in the factor graph *but not to a higher dimension of the messages passed between the nodes*. As a consequence, the computational complexity scales very well in the numbers of agents and objects. As verified in Section VII-C, a comparable scaling behavior cannot be achieved with classical Bayesian filtering techniques.

The factor graph in Fig. 2 and the corresponding BP scheme (6)–(11) combine CS and DT into a unified, coherent estimation technique. Indeed, in contrast to the conventional approach of separate CS and DT—i.e., first performing CS to localize the agents and then, based on the estimated agent locations, performing DT to localize the objects—our BP scheme exchanges probabilistic information between the CS and DT parts of the factor graph. Thereby, uncertainties in one stage are taken into account by the other stage, and the performance of both stages can be improved. This information transfer, which will be further discussed in Sections V-A and V-B, is the main reason for the superior performance of the proposed joint CS-DT algorithm; it is visualized and contrasted with the conventional approach in Fig. 3.

IV. PARTICLE-BASED PROCESSING

Because of the nonlinear and non-Gaussian state transition model (1) and measurement model (2), a closed-form evaluation of the integrals and message products in (8)–(11) is typically impossible. Therefore, we next present a new low-complexity particle-based implementation of the BP scheme (8)–(11). This implementation uses particle representations (PRs) of beliefs and messages, which approximate distributions in terms of randomly drawn particles (samples) $\mathbf{x}^{(j)}$

and weights $w^{(j)}$, for $j \in \{1, \dots, J\}$. As in conventional nonparametric BP [8], [10], the operations of *message filtering* and *message multiplication* are performed. Message filtering calculates the prediction message (8) and equals that of nonparametric BP. However, for message multiplication in (6), (7), (11), and (12), the “stacking technique” introduced in [14] is used. This technique avoids an explicit calculation of the measurement messages (9) and (10) and does not use computationally intensive kernel density estimates, which are required by nonparametric BP. Thereby, its complexity is only linear in the number of particles. We note that an alternative message multiplication scheme that also avoids the use of kernel density estimates and whose complexity is linear in the number of particles was proposed in [45]. This scheme constructs an approximate importance function in order to calculate weighted particles for beliefs and messages. Our approach is different in that the importance function is formed simply by “stacking” incoming beliefs, and the calculation of particles and weights for incoming messages is avoided.

A. Message Filtering

The message filtering operation reviewed in the following is analogous to the prediction step of the sampling importance resampling particle filter [46]. Particle-based calculation of (8) means that we obtain a PR $\{\{(\mathbf{x}_{k,n-1}^{(j)}, w_{k,n-1}^{(j)})\}_{j=1}^J\}$ of $\phi_{\rightarrow n}(\mathbf{x}_{k,n}) = \int f(\mathbf{x}_{k,n} | \mathbf{x}_{k,n-1}) b^{(P)}(\mathbf{x}_{k,n-1}) d\mathbf{x}_{k,n-1}$ from a PR $\{\{(\mathbf{x}_{k,n-1}^{(j)}, w_{k,n-1}^{(j)})\}_{j=1}^J\}$ of $b^{(P)}(\mathbf{x}_{k,n-1})$. This can be easily done by recognizing that the above integral is a marginalization of $f(\mathbf{x}_{k,n} | \mathbf{x}_{k,n-1}) b^{(P)}(\mathbf{x}_{k,n-1})$. Motivated by this interpretation, we first establish a PR $\{\{(\mathbf{x}_{k,n}^{(j)}, \mathbf{x}_{k,n-1}^{(j)}, w_{k,n-1}^{(j)})\}_{j=1}^J\}$ of $f(\mathbf{x}_{k,n} | \mathbf{x}_{k,n-1}) b^{(P)}(\mathbf{x}_{k,n-1})$ by drawing for each particle $(\mathbf{x}_{k,n-1}^{(j)}, w_{k,n-1}^{(j)})$ representing $b^{(P)}(\mathbf{x}_{k,n-1})$ one particle $\mathbf{x}_{k,n}^{(j)}$ from $f(\mathbf{x}_{k,n} | \mathbf{x}_{k,n-1}^{(j)})$. Then, removing $\{\mathbf{x}_{k,n-1}^{(j)}\}_{j=1}^J$ from $\{\{(\mathbf{x}_{k,n}^{(j)}, \mathbf{x}_{k,n-1}^{(j)}, w_{k,n-1}^{(j)})\}_{j=1}^J\}$ is the Monte Carlo implementation of the above marginalization [47]. This means that $\{\{(\mathbf{x}_{k,n}^{(j)}, w_{k,n-1}^{(j)})\}_{j=1}^J\}$ with $w_{k,n}^{(j)} = w_{k,n-1}^{(j)}$ for all $j \in \{1, \dots, J\}$ constitutes the desired PR of $\phi_{\rightarrow n}(\mathbf{x}_{k,n})$.

B. Message Multiplication

Next, we propose a message multiplication scheme for calculating the beliefs in (6) and (7) and the extrinsic informations in (12) and (11). For concreteness, we present the calculation of the agent beliefs (6); the object beliefs (7) and extrinsic informations (12) and (11) are calculated in a similar manner.

Following [14], we consider the “stacked state” $\bar{\mathbf{x}}_{l,n} \triangleq (\mathbf{x}_{k,n})_{k \in \{l\} \cup \mathcal{M}_{l,n}}$, which consists of the agent state $\mathbf{x}_{l,n}$ and the states $\mathbf{x}_{k,n}$ of all measurement partners $k \in \mathcal{M}_{l,n}$ of agent l . Using (9) in (6), one readily obtains

$$b^{(p)}(\mathbf{x}_{l,n}) = \int b^{(p)}(\bar{\mathbf{x}}_{l,n}) d\bar{\mathbf{x}}_{l,n}^{\sim l}, \quad (13)$$

where

$$b^{(p)}(\bar{\mathbf{x}}_{l,n}) \propto \phi_{\rightarrow n}(\mathbf{x}_{l,n}) \times \prod_{l' \in \mathcal{M}_{l,n}^A} f(\mathbf{y}_{l',n} | \mathbf{x}_{l,n}, \mathbf{x}_{l',n}) b^{(p-1)}(\mathbf{x}_{l',n})$$

$$\times \prod_{m \in \mathcal{M}_{l,n}^{\mathcal{O}}} f(\mathbf{y}_{l,m;n} | \mathbf{x}_{l,n}, \mathbf{x}_{m,n}) \psi_{m \rightarrow l}^{(p-1)}(\mathbf{x}_{m,n}) \quad (14)$$

and $d\bar{\mathbf{x}}_{l,n}^{(j)} \triangleq \prod_{k \in \mathcal{M}_{l,n}} d\mathbf{x}_{k,n}$. To obtain a PR of $b^{(p)}(\mathbf{x}_{l,n})$, first a PR $\{(\bar{\mathbf{x}}_{l,n}^{(j)}, w_{l,n}^{(j)})\}_{j=1}^J$ of $b^{(p)}(\bar{\mathbf{x}}_{l,n})$ is calculated as explained presently. Then, $\{(\mathbf{x}_{l,n}^{(j)}, w_{l,n}^{(j)})\}_{j=1}^J$ is a PR of $b^{(p)}(\mathbf{x}_{l,n})$. This is because $\mathbf{x}_{l,n}^{(j)}$, as a subvector of $\bar{\mathbf{x}}_{l,n}^{(j)}$, can be obtained by removing from $\bar{\mathbf{x}}_{l,n}^{(j)}$ the other subvectors $\mathbf{x}_{k,n}^{(j)}$, $k \in \mathcal{M}_{l,n}$, which is the Monte Carlo implementation of the marginalization (13) (cf. Section IV-A). Finally, a resampling [47] produces equally weighted particles representing $b^{(p)}(\mathbf{x}_{l,n})$.

A PR $\{(\bar{\mathbf{x}}_{l,n}^{(j)}, w_{l,n}^{(j)})\}_{j=1}^J$ of $b^{(p)}(\bar{\mathbf{x}}_{l,n})$ can be calculated via importance sampling using the proposal distribution¹

$$q(\bar{\mathbf{x}}_{l,n}) \triangleq \phi_{\rightarrow n}(\mathbf{x}_{l,n}) \prod_{l' \in \mathcal{M}_{l,n}^{\mathcal{A}}} b^{(p-1)}(\mathbf{x}_{l',n}) \prod_{m \in \mathcal{M}_{l,n}^{\mathcal{O}}} \psi_{m \rightarrow l}^{(p-1)}(\mathbf{x}_{m,n}).$$

There is no need to draw particles $\{\bar{\mathbf{x}}_{l,n}^{(j)}\}_{j=1}^J$ from $q(\bar{\mathbf{x}}_{l,n})$ because such particles can be obtained simply by stacking particles $\{\mathbf{x}_{l,n}^{(j)}\}_{j=1}^J$ representing $\phi_{\rightarrow n}(\mathbf{x}_{l,n})$, particles $\{\mathbf{x}_{l',n}^{(j)}\}_{j=1}^J$ representing $b^{(p-1)}(\mathbf{x}_{l',n})$, $l' \in \mathcal{M}_{l,n}^{\mathcal{A}}$, and particles $\{\mathbf{x}_{m,n}^{(j)}\}_{j=1}^J$ representing $\psi_{m \rightarrow l}^{(p-1)}(\mathbf{x}_{m,n})$, $m \in \mathcal{M}_{l,n}^{\mathcal{O}}$. (Particles representing $\phi_{\rightarrow n}(\mathbf{x}_{l,n})$ were obtained by message filtering. The other particles, for $p \geq 2$, were calculated at iteration $p-1$. For $p=1$, since $b^{(0)}(\mathbf{x}_{l,n}) = \phi_{\rightarrow n}(\mathbf{x}_{l,n})$ and $\psi_{m \rightarrow l}^{(0)}(\mathbf{x}_{m,n}) = \phi_{\rightarrow n}(\mathbf{x}_{m,n})$, these particles are identical to those obtained by message filtering.) Weights $w_{l,n}^{(j)}$ corresponding to the stacked particles $\{\bar{\mathbf{x}}_{l,n}^{(j)}\}_{j=1}^J$ are then obtained by calculating $\tilde{w}_{l,n}^{(j)} \propto b^{(p)}(\bar{\mathbf{x}}_{l,n}^{(j)})/q(\bar{\mathbf{x}}_{l,n}^{(j)})$ followed by a normalization. Using (14), the nonnormalized weights are obtained as

$$\tilde{w}_{l,n}^{(j)} = \prod_{l' \in \mathcal{M}_{l,n}^{\mathcal{A}}} f(\mathbf{y}_{l,l';n} | \mathbf{x}_{l,n}^{(j)}, \mathbf{x}_{l',n}^{(j)}) \prod_{m \in \mathcal{M}_{l,n}^{\mathcal{O}}} f(\mathbf{y}_{l,m;n} | \mathbf{x}_{l,n}^{(j)}, \mathbf{x}_{m,n}^{(j)}).$$

This algorithm avoids kernel density estimation, which is required by conventional nonparametric BP [8], [10]. Its complexity scales as $\mathcal{O}(|\mathcal{M}_{l,n}|J)$, i.e., only linearly in the number of particles J . The dimension of the distribution $b^{(p)}(\bar{\mathbf{x}}_{l,n})$ involved in the importance sampling scheme is $|\mathcal{M}_{l,n}| + 1$, and thus typically considerably higher than that of the beliefs $b^{(p)}(\mathbf{x}_{l,n})$ involved in the importance sampling scheme of nonparametric BP [8], [10]. Nevertheless, we will see in Section VII that, if the number of neighbors $|\mathcal{M}_{l,n}|$ is not too large, the number of particles J required for high accuracy is not larger than for nonparametric BP.

C. Estimation

The particle-based BP algorithm described above produces PRs $\{(\mathbf{x}_{k,n}^{(j)}, w_{k,n}^{(j)})\}_{j=1}^J$ of the state beliefs $b^{(p)}(\mathbf{x}_{k,n})$, $k \in \mathcal{E}$. An approximation of the estimate $\hat{\mathbf{x}}_{k,n}^{\text{MMSE}}$ in (4) is then obtained from the respective PR as

$$\hat{\mathbf{x}}_{k,n} = \sum_{j=1}^J w_{k,n}^{(j)} \mathbf{x}_{k,n}^{(j)}.$$

¹An alternative proposal distribution that is more appropriate if agent l is static or if the prediction message is very noninformative is presented in [48].

V. DISTRIBUTED ALGORITHM

We next develop a distributed algorithm that combines the particle-based BP algorithm discussed in Section IV with a consensus scheme [19]. The overall organization of this algorithm is as follows. Each agent $l \in \mathcal{A}$ performs particle-based estimation of its own state $\mathbf{x}_{l,n}$ and of the states $\mathbf{x}_{m,n}$, $m \in \mathcal{O}$ of all the objects. Thus, the calculations required to estimate an agent state $\mathbf{x}_{l,n}$, $l \in \mathcal{A}$ are performed only once in the network (at agent l), whereas the calculations required to estimate an object state $\mathbf{x}_{m,n}$, $m \in \mathcal{O}$ are performed $|\mathcal{A}|$ times (at each agent $l \in \mathcal{A}$). Accordingly, each agent belief $b^{(p)}(\mathbf{x}_{l,n})$ is stored (temporarily, i.e., during one message passing iteration) only at the respective agent l whereas copies of all object beliefs $b^{(p)}(\mathbf{x}_{m,n})$ are stored (temporarily) at all agents $l \in \mathcal{A}$. However, all the calculations performed at any given agent $l \in \mathcal{A}$ are collaborative in that they use probabilistic information related to all the other agents and objects. The proposed distributed algorithm requires only communication between neighboring agents to disseminate this probabilistic information. The distributed calculation of the object beliefs, agent beliefs, and extrinsic informations will be discussed in the next three subsections.

A. Distributed Calculation of the Object Beliefs

Estimation of the object states $\mathbf{x}_{m,n}$, $m \in \mathcal{O}$ from $\mathbf{y}_{1:n}$ according to (4) essentially amounts to a computation of $f(\mathbf{x}_{m,n} | \mathbf{y}_{1:n})$. The following discussion describes the calculations associated with the red dotted box in Fig. 2. According to (7) and (8), the object belief $b^{(p)}(\mathbf{x}_{m,n})$, $p \in \{1, \dots, P\}$ approximating $f(\mathbf{x}_{m,n} | \mathbf{y}_{1:n})$ is given by

$$b^{(p)}(\mathbf{x}_{m,n}) \propto \phi_{\rightarrow n}(\mathbf{x}_{m,n}) \Phi_{m,n}^{(p)}(\mathbf{x}_{m,n}), \quad (15)$$

with

$$\phi_{\rightarrow n}(\mathbf{x}_{m,n}) = \int f(\mathbf{x}_{m,n} | \mathbf{x}_{m,n-1}) b^{(P)}(\mathbf{x}_{m,n-1}) d\mathbf{x}_{m,n-1} \quad (16)$$

and

$$\Phi_{m,n}^{(p)}(\mathbf{x}_{m,n}) \triangleq \prod_{l \in \mathcal{A}_{m,n}} \phi_{l \rightarrow m}^{(p)}(\mathbf{x}_{m,n}). \quad (17)$$

According to (15), each agent needs to calculate the prediction message $\phi_{\rightarrow n}(\mathbf{x}_{m,n})$ in (16) and the measurement message product $\Phi_{m,n}^{(p)}(\mathbf{x}_{m,n})$ in (17). The messages $\phi_{l \rightarrow m}^{(p)}(\mathbf{x}_{m,n})$ contained in (17) involve the extrinsic informations $\psi_{l \rightarrow m}^{(p-1)}(\mathbf{x}_{l,n})$ (see (10)); particle-based calculation of the latter will be discussed in Section V-C. However, at each agent at most one message $\phi_{l \rightarrow m}^{(p)}(\mathbf{x}_{m,n})$ is available (for a given m). We will solve this problem by means of a consensus scheme.

1) *Particle-based Calculation of $b^{(p)}(\mathbf{x}_{m,n})$* : An approximate particle-based calculation of $b^{(p)}(\mathbf{x}_{m,n}) \propto \phi_{\rightarrow n}(\mathbf{x}_{m,n}) \Phi_{m,n}^{(p)}(\mathbf{x}_{m,n})$ in (15) can be obtained via importance sampling with proposal distribution $\phi_{\rightarrow n}(\mathbf{x}_{m,n})$. First, based on (16), particles $\{\mathbf{x}_{m,n}^{(j)}\}_{j=1}^J$ representing $\phi_{\rightarrow n}(\mathbf{x}_{m,n})$ are calculated from particles representing $b^{(P)}(\mathbf{x}_{m,n-1})$ by means of message filtering (cf. Section IV-A; note that particles representing $b^{(P)}(\mathbf{x}_{m,n-1})$ were calculated by each agent at time $n-1$). Next, weights $\{w_{m,n}^{(j)}\}_{j=1}^J$ are calculated as

$$\tilde{w}_{m,n}^{(j)} = \Phi_{m,n}^{(p)}(\mathbf{x}_{m,n}^{(j)}) \quad (18)$$

followed by a normalization. Finally, resampling is performed to obtain equally weighted particles representing $b^{(p)}(\mathbf{x}_{m,n})$. However, this particle-based implementation presupposes that the message product $\Phi_{m,n}^{(p)}(\mathbf{x}_{m,n})$ evaluated at the particles $\{\mathbf{x}_{m,n}^{(j)}\}_{j=1}^J$ is available at the agents.

2) *Distributed Evaluation of $\Phi_{m,n}^{(p)}(\cdot)$* : For a distributed computation of $\Phi_{m,n}^{(p)}(\mathbf{x}_{m,n}^{(j)})$, $j \in \{1, \dots, J\}$, we first note that (17) can be written as

$$\Phi_{m,n}^{(p)}(\mathbf{x}_{m,n}^{(j)}) = \exp(|\mathcal{A}|\chi_{m,n}^{(p,j)}), \quad (19)$$

with

$$\chi_{m,n}^{(p,j)} \triangleq \frac{1}{|\mathcal{A}|} \sum_{l \in \mathcal{A}_{m,n}} \log \phi_{l \rightarrow m}^{(p)}(\mathbf{x}_{m,n}^{(j)}), \quad j \in \{1, \dots, J\}. \quad (20)$$

Thus, $\Phi_{m,n}^{(p)}(\mathbf{x}_{m,n}^{(j)})$ is expressed in terms of the arithmetic average $\chi_{m,n}^{(p,j)}$. For each j , following the ‘‘consensus-over-weights’’ approach of [19], this average can be computed in a distributed manner by a consensus or gossip scheme [49], [50]. These schemes are iterative; in each iteration i , they compute an internal state $\zeta_{l,m;n}^{(j,i)}$ at each agent l . This internal state is initialized as

$$\zeta_{l,m;n}^{(j,0)} = \begin{cases} \log \phi_{l \rightarrow m}^{(p)}(\mathbf{x}_{m,n}^{(j)}), & l \in \mathcal{A}_{m,n} \\ 0 & l \notin \mathcal{A}_{m,n}. \end{cases}$$

Here, $\phi_{l \rightarrow m}^{(p)}(\mathbf{x}_{m,n}^{(j)})$ is computed by means of a Monte Carlo approximation [44] of the integral in (10), i.e.,

$$\phi_{l \rightarrow m}^{(p)}(\mathbf{x}_{m,n}^{(j)}) = \frac{1}{J} \sum_{j'=1}^J f(\mathbf{y}_{l,m;n} | \mathbf{x}_{l,n}^{(j')}, \mathbf{x}_{m,n}^{(j)}).$$

This uses the particles $\{\mathbf{x}_{l,n}^{(j)}\}_{j=1}^J$ representing $\psi_{l \rightarrow m}^{(p-1)}(\mathbf{x}_{l,n})$, whose calculation will be discussed in Section V-C. If—as assumed in Section II-A—the communication graph is connected, then for $j \rightarrow \infty$ the internal state $\zeta_{l,m;n}^{(j,i)}$ converges to the average $\chi_{m,n}^{(p,j)}$ in (20) [49], [50]. Thus, for a sufficiently large number C of iterations, because of (19), a good approximation of $\Phi_{m,n}^{(p)}(\mathbf{x}_{m,n}^{(j)})$ is obtained at each agent by

$$\Phi_{m,n}^{(p)}(\mathbf{x}_{m,n}^{(j)}) \approx \exp(|\mathcal{A}|\zeta_{l,m;n}^{(j,C)}). \quad (21)$$

Here, the number of agents $|\mathcal{A}|$ can be determined in a distributed way by using another consensus or gossip algorithm at time $n = 0$ [51]. Furthermore, an additional max-consensus scheme [19] has to be used to obtain perfect consensus on the weights $\tilde{w}_{m,n}^{(j)}$ in (18) and, in turn, identical particles at all agents [19]. The max-consensus converges in I iterations, where I is the diameter of the communication graph [52]. Finally, the pseudo-random number generators of all agents (which are used for drawing particles) have to be synchronized, i.e., initialized with the same seed at time $n = 0$. This distributed evaluation of $\Phi_{m,n}^{(p)}(\mathbf{x}_{m,n}^{(j)})$ requires only local communication: in each iteration, for each of the J instances of the consensus or gossip scheme, J real values are broadcast by each agent to neighboring agents [49], [50]. This holds for both averaging and maximization separately.

As an alternative to this scheme, the likelihood consensus scheme [18], [20] can be employed to provide an approximation of the functional form of $\Phi_{m,n}^{(p)}(\mathbf{x}_{m,n})$ to each agent, again using only local communication [32]. The likelihood consensus scheme does not require additional max-consensus algorithms and synchronized pseudo-random number generators, but tends to require a more informative proposal distribution for message multiplication (cf. Section IV-B).

3) *Probabilistic Information Transfer*: According to (10), the messages $\phi_{l \rightarrow m}^{(p)}(\mathbf{x}_{m,n})$ occurring in $\Phi_{m,n}^{(p)}(\mathbf{x}_{m,n}) = \prod_{l \in \mathcal{A}_{m,n}} \phi_{l \rightarrow m}^{(p)}(\mathbf{x}_{m,n})$ (see (17)) involve the extrinsic informations $\psi_{l \rightarrow m}^{(p-1)}(\mathbf{x}_{l,n})$ of all agents l observing object m , i.e., $l \in \mathcal{A}_{m,n}$. Therefore, they constitute an information transfer from the CS part of the algorithm to the DT part (cf. Fig. 3(b) and, in more detail, the directed edges entering the red dotted box in Fig. 2). The estimation of object state $\mathbf{x}_{m,n}$ is based on the belief $b^{(p)}(\mathbf{x}_{m,n})$ as given by (15), and thus on $\Phi_{m,n}^{(p)}(\mathbf{x}_{m,n})$. This improves on pure DT because probabilistic information about the states of the agents $l \in \mathcal{A}_{m,n}$ —provided by $\psi_{l \rightarrow m}^{(p-1)}(\mathbf{x}_{l,n})$ —is taken into account. By contrast, pure DT [20] uses the global likelihood function—involving the measurements of all agents—instead of $\Phi_{m,n}^{(p)}(\mathbf{x}_{m,n})$. This presupposes that the agent states are known. In separate CS and DT, estimates of the agent states provided by CS are used for DT, rather than probabilistic information about the agent states as is done in the proposed combined CS–DT algorithm. The improved accuracy of object state estimation achieved by our algorithm compared to separate CS and DT will be demonstrated in Section VII-B.

B. Distributed Calculation of the Agent Beliefs

For a distributed calculation of the agent belief $b^{(p)}(\mathbf{x}_{l,n})$, $l \in \mathcal{A}$, the following information is available at agent l : (i) equally weighted particles representing $\psi_{m \rightarrow l}^{(p-1)}(\mathbf{x}_{m,n})$ for all objects $m \in \mathcal{O}$ (which were calculated as described in Section V-C); (ii) equally weighted particles representing $b^{(p-1)}(\mathbf{x}_{l',n})$ for all neighboring agents $l' \in \mathcal{M}_{l,n}^A$ (which were received from these agents); and (iii) a PR of $b^{(P)}(\mathbf{x}_{l,n-1})$ (which was calculated at time $n-1$). Using this information and the locally available measurements $\mathbf{y}_{l,k;n}$, $k \in \mathcal{M}_{l,n}$, a PR $\{(\mathbf{x}_{l,n}^{(j)}, w_{l,n}^{(j)})\}_{j=1}^J$ of $b^{(p)}(\mathbf{x}_{l,n})$ can be calculated in a distributed manner by implementing (6), using the particle-based message multiplication scheme presented in Section IV-B. Finally, resampling is performed to obtain equally weighted particles representing $b^{(p)}(\mathbf{x}_{l,n})$. This calculation of the agent beliefs improves on pure CS [2] in that it uses the probabilistic information about the states of the objects $m \in \mathcal{M}_{l,n}^{\mathcal{O}}$ provided by the messages $\psi_{m \rightarrow l}^{(p-1)}(\mathbf{x}_{m,n})$ (cf. (11)). This probabilistic information transfer from DT to CS is depicted in Fig. 3(b) and, in more detail, by the directed edges leaving the red dotted box in Fig. 2. The resulting improved accuracy of agent state estimation will be demonstrated in Section VII.

C. Distributed Calculation of the Extrinsic Informations

Since (12) is analogous to (6) and (11) is analogous to (7), particles for $\psi_{l \rightarrow m}^{(p)}(\mathbf{x}_{l,n})$ or $\psi_{m \rightarrow l}^{(p)}(\mathbf{x}_{m,n})$ can be cal-

culated similarly as for the corresponding belief. However, the following shortcut reuses previous results. To obtain particles for $\psi_{m \rightarrow l}^{(p)}(\mathbf{x}_{m,n})$, we proceed as for $b^{(p)}(\mathbf{x}_{m,n})$ (see Sections V-A1 and V-A2) but replace $\exp(|\mathcal{A}|\zeta_{l,m;n}^{(j,C)})$ in (21) with $\exp(|\mathcal{A}|\zeta_{l,m;n}^{(j,C)} - \zeta_{l,m;n}^{(j,0)})$. (This is based on the relation $\psi_{m \rightarrow l}^{(p)}(\mathbf{x}_{m,n}) \propto b^{(p)}(\mathbf{x}_{m,n})/\phi_{l \rightarrow m}^{(p)}(\mathbf{x}_{m,n})$, cf. (7) and (11).) Here, $\zeta_{l,m;n}^{(j,C)}$ and $\zeta_{l,m;n}^{(j,0)}$ are already available locally from the calculation of $b^{(p)}(\mathbf{x}_{m,n})$.

D. Statement of the Distributed Algorithm

The proposed distributed CS–DT algorithm is obtained by combining the operations discussed in Sections V-A through V-C, as summarized in the following.

ALGORITHM 1: DISTRIBUTED CS–DT ALGORITHM

Initialization: The recursive algorithm described below is initialized at time $n=0$ and agent l with particles $\{\mathbf{x}_{k,0}^{(j)}\}_{j=1}^J$ drawn from a prior pdf $f(\mathbf{x}_{k,0})$, for $k \in \{l\} \cup \mathcal{O}$.

Recursion at time n : At agent l , equally weighted particles $\{\mathbf{x}_{k,n-1}^{(j)}\}_{j=1}^J$ representing the beliefs $b^{(P)}(\mathbf{x}_{k,n-1})$ with $k \in \{l\} \cup \mathcal{O}$ are available (these were calculated at time $n-1$). At time n , agent l performs the following operations.

Step 1—Prediction: From $\{\mathbf{x}_{k,n-1}^{(j)}\}_{j=1}^J$, PRs $\{\mathbf{x}_{k,n}^{(j)}\}_{j=1}^J$ of the prediction messages $\phi_{\rightarrow n}(\mathbf{x}_{k,n})$, $k \in \{l\} \cup \mathcal{O}$ are calculated via message filtering (see Section IV-A) based on the state-transition pdf $f(\mathbf{x}_{k,n}|\mathbf{x}_{k,n-1})$, i.e., for each $\mathbf{x}_{k,n-1}^{(j)}$ one particle $\mathbf{x}_{k,n}^{(j)}$ is drawn from $f(\mathbf{x}_{k,n}|\mathbf{x}_{k,n-1}^{(j)})$.

Step 2—BP message passing: For each $k \in \{l\} \cup \mathcal{O}$, the belief is initialized as $b^{(0)}(\mathbf{x}_{k,n}) = \phi_{\rightarrow n}(\mathbf{x}_{k,n})$, in the sense that the PR of $\phi_{\rightarrow n}(\mathbf{x}_{k,n})$ is used as PR of $b^{(0)}(\mathbf{x}_{k,n})$. Then, for $p=1, \dots, P$:

- a) For each $m \in \mathcal{O}$, a PR $\{\mathbf{x}_{m,n}^{(j)}, w_{m,n}^{(j)}\}_{j=1}^J$ of $b^{(p)}(\mathbf{x}_{m,n})$ in (7) is obtained via importance sampling with proposal distribution $\phi_{\rightarrow n}(\mathbf{x}_{m,n})$ (see Section V-A1). That is, using the particles $\{\mathbf{x}_{m,n}^{(j)}\}_{j=1}^J$ representing $\phi_{\rightarrow n}(\mathbf{x}_{m,n})$ (calculated in Step 1), nonnormalized weights are calculated as $\tilde{w}_{m,n}^{(j)} = \Phi(\mathbf{x}_{m,n}^{(j)})$ (cf. (18)) for all $j \in \{1, \dots, J\}$ in a distributed manner as described in Section V-A2. The final weights $w_{m,n}^{(j)}$ are obtained by a normalization.
- b) (Not done for $p=P$) For each $m \in \mathcal{M}_{l,n}^{\mathcal{O}}$, a PR of $\psi_{m \rightarrow l}^{(p)}(\mathbf{x}_{m,n})$ is calculated in a similar manner (see Section V-C).
- c) A PR $\{\mathbf{x}_{l,n}, w_{l,n}^{(j)}\}_{j=1}^J$ of $b^{(p)}(\mathbf{x}_{l,n})$ is calculated by implementing (6) as described in Section V-B. This involves equally weighted particles of all $b^{(p-1)}(\mathbf{x}_{l',n})$, $l' \in \mathcal{M}_{l,n}^{\mathcal{A}}$ (which were received from agents $l' \in \mathcal{M}_{l,n}^{\mathcal{A}}$ at message passing iteration $p-1$) and of all $\psi_{m \rightarrow l}^{(p-1)}(\mathbf{x}_{m,n})$, $m \in \mathcal{M}_{l,n}^{\mathcal{O}}$ (which were calculated in Step 2b at message passing iteration $p-1$).
- d) (Not done for $p=P$) For each $m \in \mathcal{M}_{l,n}^{\mathcal{O}}$, a PR of $\psi_{l \rightarrow m}^{(p)}(\mathbf{x}_{l,n})$ is calculated in a similar manner.
- e) For all PRs calculated in Steps 2a–2d, resampling is performed to obtain equally weighted particles.
- f) (Not done for $p=P$) The equally weighted particles of $b^{(p)}(\mathbf{x}_{l,n})$ calculated in Step 2c are broadcast to all agents l' for which $l \in \mathcal{M}_{l',n}^{\mathcal{A}}$, and equally weighted particles of $b^{(p)}(\mathbf{x}_{l',n})$

are received from each neighboring agent $l_1 \in \mathcal{M}_{l',n}^{\mathcal{A}}$. Thus, at this point, agent l has available equally weighted particles $\{\mathbf{x}_{k,n}^{(j)}\}_{j=1}^J$ of $b^{(p)}(\mathbf{x}_{k,n})$, $k \in \{l\} \cup \mathcal{O} \cup \mathcal{M}_{l,n}^{\mathcal{A}}$ and equally weighted particles $\{\mathbf{x}_{m,n}^{(j)}\}_{j=1}^J$ of $\psi_{m \rightarrow l}^{(p)}(\mathbf{x}_{m,n})$ and of $\psi_{l \rightarrow m}^{(p)}(\mathbf{x}_{l,n})$, $m \in \mathcal{M}_{l,n}^{\mathcal{O}}$.

Step 3—Estimation: For $k \in \{l\} \cup \mathcal{O}$, an approximation of the global MMSE state estimate $\hat{\mathbf{x}}_{k,n}^{\text{MMSE}}$ in (4) is computed from the PR $\{\mathbf{x}_{k,n}^{(j)}, w_{k,n}^{(j)}\}_{j=1}^J$ of $b^{(P)}(\mathbf{x}_{k,n})$ according to

$$\hat{\mathbf{x}}_{k,n} = \sum_{j=1}^J w_{k,n}^{(j)} \mathbf{x}_{k,n}^{(j)}, \quad k \in \{l\} \cup \mathcal{O}.$$

VI. COMMUNICATION REQUIREMENTS AND DELAY

In the following discussion of the communication requirements of the proposed distributed CS–DT algorithm, we assume for simplicity that all $\tilde{\mathbf{x}}_{k,n}$, $k \in \mathcal{E}$ (i.e., the substates actually involved in the measurements, cf. (3)) have identical dimension L . Furthermore, we denote by C the number of consensus or gossip iterations used for averaging, by P the number of message passing iterations, by J the number of particles, and by I the diameter of the communication graph. For an analysis of the delay caused by communication, we assume that all agents can transmit in parallel. More specifically, broadcasting the beliefs of all the agents will be counted as one delay time slot, and broadcasting all quantities related to one consensus iteration for averaging or maximization will also be counted as one delay time slot.

- For calculation of the object beliefs $b^{(p)}(\mathbf{x}_{m,n})$, $m \in \mathcal{O}$ using the “consensus-over-weights” scheme (see Section V-A and Step 2a in Algorithm 1), at each time n , agent $l \in \mathcal{A}$ broadcasts $N^C \triangleq P(C+I)J|\mathcal{O}|$ real values to agents $l' \in \mathcal{C}_{l,n}$. The corresponding contribution to the overall delay is $P(C+I)$ time slots, because the consensus coefficients for all objects are broadcast in parallel.
- To support neighboring agents l' with $l \in \mathcal{M}_{l',n}^{\mathcal{A}}$ in calculating their own beliefs $b^{(p)}(\mathbf{x}_{l',n})$ (see Section V-B and Step 2c in Algorithm 1), at each time n , agent $l \in \mathcal{A}$ broadcasts $N^{\text{NBP}} \triangleq PJL$ real values to those neighboring agents. The delay contribution is P time slots, because each agent broadcasts a belief in each message passing iteration.

Therefore, at each time n , the total number of real values broadcast by each agent during P message passing iterations is

$$N^{\text{TOT}} = N^C + N^{\text{NBP}} = PJ((C+I)|\mathcal{O}| + L).$$

The corresponding delay is $P(C+I)$ time slots. In the extended version of this paper [48], we present methods to reduce the communication requirements and delay.

VII. SIMULATION RESULTS

We will study the performance and communication requirements of the proposed method (PM) in two dynamic scenarios and a static scenario. In addition, we will investigate the scalability of the PM by comparing it with two different particle filters and with conventional nonparametric BP [8], [10]. Simulation source files and animated plots are available at <http://www.nt.tuwien.ac.at/about-us/staff/florian-meyer/>.




Fig. 4. Network topology in the dynamic scenarios, with example trajectories of mobile agents and objects. Initial mobile agent locations are indicated by crosses, initial object locations by stars, and anchor locations by circles.

A. Dynamic Scenarios

In the dynamic scenarios, we consider $|\mathcal{A}|=12$ agents and $|\mathcal{O}|=2$ objects as shown in Fig. 4. Eight agents are mobile and four are static anchors (i.e., agents with perfect location information). Each agent has a communication range of 50 and attempts to localize itself (except for the anchors) and the two objects. The states of the mobile agents (MAs) and objects consist of location and velocity, i.e., $\mathbf{x}_{k,n} \triangleq (x_{1,k,n} \ x_{2,k,n} \ \dot{x}_{1,k,n} \ \dot{x}_{2,k,n})^T$. All agents $l \in \mathcal{A}$ acquire distance measurements according to (3), i.e., $y_{l,k;n} = \|\tilde{\mathbf{x}}_{l,n} - \tilde{\mathbf{x}}_{k,n}\| + v_{l,k;n}$, where $\tilde{\mathbf{x}}_{k,n} \triangleq (x_{1,k,n} \ x_{2,k,n})^T$ is the location of agent or object $k \in \mathcal{M}_{l,n}$ and the measurement noise $v_{l,k;n}$ is independent across l, k , and n and Gaussian with variance $\sigma_v^2 = 2$. The states of the MAs and objects evolve independently according to $\mathbf{x}_{k,n} = \mathbf{G}\mathbf{x}_{k,n-1} + \mathbf{W}\mathbf{u}_{k,n}$, $n=1, 2, \dots$ [34], where

$$\mathbf{G} = \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad \mathbf{W} = \begin{pmatrix} 0.5 & 0 \\ 0 & 0.5 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

The driving noise vectors $\mathbf{u}_{k,n} \in \mathbb{R}^2$ are Gaussian, i.e., $\mathbf{u}_{k,n} \sim \mathcal{N}(\mathbf{0}, \sigma_u^2 \mathbf{I})$, with component variance $\sigma_u^2 = 5 \cdot 10^{-5}$ for the MAs and $\sigma_u^2 = 5 \cdot 10^{-4}$ for the objects; furthermore, $\mathbf{u}_{k,n}$ and $\mathbf{u}_{k',n'}$ are independent unless $(k, n) = (k', n')$. Each MA starts moving only when it is sufficiently localized in the sense that the empirical component variance of the estimated location vector is below $5\sigma_v^2 = 10$; it then attempts to reach the center of the scene, $\tilde{\mathbf{x}}_c = (37.5 \ 37.5)^T$, in 75 time steps. For generating the measurements, the MA trajectories are initialized using a Dirac-shaped prior $f(\mathbf{x}_{l,0})$ located at $\boldsymbol{\mu}_{l,0} = (x_{1,l,0} \ x_{2,l,0} \ (\tilde{x}_{1,c} - x_{1,l,0})/75 \ (\tilde{x}_{2,c} - x_{2,l,0})/75)^T$, where $x_{1,l,0}$ and $x_{2,l,0}$ are chosen as shown in Fig. 4, and the two object trajectories are initialized using a Dirac-shaped prior $f(\mathbf{x}_{m,0})$ located at $(15 \ 0 \ 0.8 \ 0.6)^T$ and $(75 \ 20 \ -0.8 \ 0.6)^T$ (see Fig. 4). Note that knowledge of these locations is not used by the simulated algorithms.

We compare the PM with a reference method (RM) that separately performs CS by means of a particle-based implementation of [2] using nonparametric BP [8], [10] and DT by means of a consensus-based distributed particle filter [19], [53]; the latter uses the MA location estimates provided by CS. Both PM and RM use $P = 1$ message passing




Fig. 5. MA self-localization RMSE and object localization RMSE versus time n (dynamic scenario 1).

iteration and $J = 1000$ particles. (We chose $P = 1$ since we observed almost no performance gains for $P > 1$; moreover, $P > 1$ may lead to overconfident beliefs [54].) For a distributed implementation, PM and RM employ $C = 6$ iterations of an average consensus with Metropolis weights [55]. They are initialized with a location prior for objects and MAs that is uniform on $[-200, 200] \times [-200, 200]$ and a Gaussian velocity prior for the MAs (after the respective MA is sufficiently localized as described above) with mean $((\tilde{x}_{1,c} - \hat{x}_{1,l,n'})/75 \ (\tilde{x}_{2,c} - \hat{x}_{2,l,n'})/75)^T$ and covariance matrix $\text{diag}\{10^{-3}, 10^{-3}\}$. Here, $\hat{\mathbf{x}}_{l,n'}$ is the location estimate at the time n' at which MA l is sufficiently localized for the first time. We furthermore used a velocity prior for the objects that is Gaussian with mean $\boldsymbol{\mu}_{m,0}^{(v)}$ and covariance $\mathbf{C}_{m,0}^{(v)}$. Here, $\mathbf{C}_{m,0}^{(v)} = \text{diag}\{0.001, 0.001\}$ represents the uncertainty in knowing the velocity $\dot{\mathbf{x}}_{m,0}$ of object m at time $n = 0$, and $\boldsymbol{\mu}_{m,0}^{(v)}$ is a random hyperparameter that was sampled for each simulation run from $\mathcal{N}(\dot{\mathbf{x}}_{m,0}, \mathbf{C}_{m,0}^{(v)})$.

We simulated two different dynamic scenarios. In dynamic scenario 1, the measurement range of those four MAs that are initially located near the corners as shown in Fig. 4 (these agents will be termed ‘‘corner agents’’) is limited as specified later whereas all the other agents cover the entire field of size 75×75 . In dynamic scenario 2, the measurement range of all agents is limited to 20. In both dynamic scenarios, for particle-based message multiplication (cf. Section IV-B) at time $n = 1$, we used the alternative proposal distribution described in [48]. We note that these dynamic scenarios cannot be tackled by SLAT algorithms [24]–[28] since they involve MAs whereas SLAT assumes static agents.

Fig. 5 shows the root-mean-square errors (RMSEs) of MA self-localization and object localization for $n = 1, \dots, 75$ in dynamic scenario 1, with the measurement range of the corner agents chosen as 20. The MA self-localization RMSE and the object localization RMSE were determined by averaging over all MAs and all objects, respectively, and over 100 simulation runs. It is seen that the MA self-localization RMSE of PM is significantly smaller than that of RM. This is because with pure CS, the corner agents do not have enough partners for accurate self-localization, whereas with PM, they can use their measured distances to the objects to calculate the messages from the object nodes, $\phi_{m \rightarrow l}^{(p)}(\mathbf{x}_{l,n})$, which support




Fig. 6. MA self-localization RMSE and object localization RMSE versus measurement range ρ of the corner agents (dynamic scenario 1).




Fig. 7. MA self-localization RMSE and object localization RMSE versus time n (dynamic scenario 2).

self-localization. The object localization RMSEs of PM and RM are very similar at all times. This is because the objects are always measured by several well-localized agents.

Fig. 6 shows the MA self-localization and object localization RMSEs averaged over time n versus the measurement range ρ of the corner agents. For small and large ρ , PM performs similarly to RM but for different reasons: When ρ is smaller than 12.5, the objects appear in the measurement regions of the corner agents only with a very small probability. Thus, at most times, the messages $\phi_{m \rightarrow l}^{(p)}(\mathbf{x}_{l,n})$ from the object nodes cannot be calculated. For ρ larger than 25, the corner agents measure three well-localized agents at time $n=1$, and thus they are also able to localize themselves using pure CS. For ρ between 15 and 25, PM significantly outperforms RM (cf. our discussion of Fig. 5). The object localization RMSEs of PM and RM are very similar and almost independent of ρ . This is because for all ρ , the objects are again measured by several well-localized agents.

Finally, Fig. 7 shows the MA self-localization and object localization RMSEs for $n=1, \dots, 75$ in dynamic scenario 2 (i.e., the measurement range of all agents is 20). It can be seen that with both methods, the objects are roughly localized after a few initial time steps. However, with RM, due to the limited measurement range, not even a single MA can be localized. With PM, once meaningful probabilistic information about the object locations is available, also the self-localization RMSE decreases and most of the MAs can be localized after some time. This is possible since the MAs obtain additional




Fig. 8. Network topology in the static scenario, with anchor locations (indicated by circles) and example realizations of non-anchor agent locations (indicated by crosses) and of object locations (indicated by stars).

information related to the measured objects. Which MAs are localized at what times depends on the object trajectories and varies between the simulation runs.

Summarizing the results displayed in Figs. 5–7, one can conclude that the MA self-localization performance of PM is generally much better than that of RM whereas the object localization performance is not improved. In Section VII-B, we will present a scenario in which also object localization is improved. The quantities determining the communication requirements of PM according to Section VI are $N^C = 18000$ and $N^{\text{NBP}} = 2000$. The resulting total communication requirement per MA and time step is $N^{\text{TOT}} = 20000$ for both dynamic scenarios. At time $n=1$, each MA additionally broadcasts $N^{\text{AP}} = 12000$ real values for disseminating the alternative proposal distribution [48]. According to Section VI, the delay is 9 time slots in both dynamic scenarios. RM has the same communication requirements and causes the same delay as PM.

B. Static Scenario

Next, we consider a completely static scenario. This can be more challenging than the dynamic scenario considered in the last section, since at the first message passing iterations the beliefs of all entities can be highly multimodal, and thus BP algorithms using Gaussian approximations [13], [14] are typically not suitable for reliable localization. In the simulated scenario, there are $|\mathcal{A}| = 63$ static agents and $|\mathcal{O}| = 50$ static objects. 13 agents are anchors located as depicted in Fig. 8. The 50 remaining agents and the objects are randomly (uniformly) placed in a field of size 100×100 ; a realization of the locations of the non-anchor agents and objects is shown in Fig. 8. The states of the non-anchor agents and of the objects are the locations, i.e., $\mathbf{x}_{k,n} = \tilde{\mathbf{x}}_{k,n} = (x_{1,k,n} \ x_{2,k,n})^T$. Each agent performs distance measurements according to (3) with a measurement range of 22.5 and a noise variance of $\sigma_v^2 = 2$. The communication range of each agent is 50. The prior for the non-anchor agents and for the objects is uniform on $[-200, 200] \times [-200, 200]$. Both PM and RM use $J=1000$ particles and $C=15$ average consensus iterations. Since all agents and objects are static, we simulated only a single time step. This scenario is similar to that considered in [2] for pure CS, except that 50 of the agents used in [2] are replaced




Fig. 9. Non-anchor agent self-localization RMSE and object localization RMSE versus message passing iteration index p (static scenario).

by objects and also anchor nodes perform measurements. For message multiplication, we used the alternative proposal distribution described in [48].

Fig. 9 shows the localization RMSEs versus the message passing iteration index p . It is seen that the agent self-localization performance of PM is significantly better than that of RM. Again, this is because agents can use messages $\phi_{m \rightarrow l}^{(p)}(\mathbf{x}_{l,n})$ from well-localized objects to better localize themselves. Furthermore, also the object localization performance of PM is significantly better. This is because with separate CS and DT, poor self-localization of certain agents degrades the object localization performance. It is finally seen that increasing p beyond 5 does not lead to a significant reduction of the RMSEs.

In this scenario, assuming $P = 3$, we have $N^C = 2.70 \cdot 10^6$, $N^{\text{NBP}} = 6000$, and $N^{\text{TOT}} = 2.71 \cdot 10^6$. For proposal adaptation, each non-anchor agent additionally broadcasts $N^{\text{AP}} = 9.00 \cdot 10^5$ real values [48]. The delay is 54 time slots.

C. Scalability

Finally, we consider again a dynamic scenario and investigate the scalability of PM for growing network size in comparison to conventional particle filtering approximating the classical sequential Bayesian filter. Because this aspect is not fundamentally related to a distributed implementation, we consider a centralized scenario where all measurements are processed at a fusion center. We compare a centralized version of PM with the sampling importance resampling particle filter [46] (abbreviated as SPF), the unscented particle filter (UPF) [56], and a particle implementation of the proposed BP scheme (6)–(11) using conventional nonparametric belief propagation (NBP) [8], [10]. Because PM and NBP are centralized, they do not need a consensus for DT. Both SPF and UPF estimate the total “stacked” state of all MAs and objects, whose dimension grows with the network size. The state of an MA or object consists of location and velocity.

We consider mobile networks of increasing size $(|\tilde{\mathcal{A}}|, |\mathcal{O}|) = (8, 2), (16, 4), (32, 8), (64, 16),$ and $(128, 32)$, where $\tilde{\mathcal{A}} \subseteq \mathcal{A}$ is the set of MAs. In addition to the MAs and objects, four anchors are placed at locations $(-100 \ -100)^T, (-100 \ 100)^T, (100 \ -100)^T,$ and $(100 \ 100)^T$. For the MAs and objects, we use the motion model of Section VII-A




Fig. 10. Average runtime versus network size.

with driving noise variance $\sigma_u^2 = 10^{-2}$, and for the agents (MAs and anchors), we use the measurement model of Section VII-A with measurement noise variance $\sigma_v^2 = 1$. For generating the measurements, the MA and object trajectories are initialized as $\boldsymbol{\mu}_{k,0} = (x_{1,k,0} \ x_{2,k,0} \ 0 \ 0)^T$, $k \in \mathcal{E}$ where $x_{1,k,0}$ and $x_{2,k,0}$ are randomly (uniformly) chosen in a field of size 100×100 . The algorithms are initialized with the initial prior pdf $f(\mathbf{x}_{k,0}) = \mathcal{N}(\boldsymbol{\mu}_{k,0}, \mathbf{C}_{k,0})$, $k \in \mathcal{E}$. Here, $\mathbf{C}_{k,0} = \text{diag}\{10^{-2}, 10^{-2}, 10^{-2}, 10^{-2}\}$, and $\boldsymbol{\mu}_{k,0}$ is sampled for each simulation run from $\mathcal{N}(\mathbf{x}_{k,0}, \mathbf{C}_{k,0})$, where $\mathbf{x}_{k,0}$ is the true initial state used for generating the object trajectories. Since an informative initial prior is available for all entities, we do not use the alternative proposal distribution.

The measurement topology of the network is randomly determined at each time step as follows: Each MA measures, with equal probability, one or two randomly chosen anchors and two randomly chosen MAs or objects. This is done such that each object is measured by two randomly chosen MAs. In addition, each object is also measured by one or two randomly chosen anchors. The sets $\mathcal{M}_{l,n}^{\tilde{\mathcal{A}}}$, $l \in \tilde{\mathcal{A}}$ are symmetric in that $l' \in \mathcal{M}_{l,n}^{\tilde{\mathcal{A}}}$ implies $l \in \mathcal{M}_{l',n}^{\tilde{\mathcal{A}}}$. Furthermore, the sets $\mathcal{M}_{l,n}$, $l \in \tilde{\mathcal{A}}$ are chosen such that the topology graph that is formed by the object measurements performed by the MAs corresponds to a Hamiltonian cycle, i.e., each MA and object in the graph is visited exactly once [57]. This ensures that the expected number of neighbors of each entity is independent of the network size and the network is connected at all times. PM, SPF, and UPF use $J = 1000$ and 5000 particles, whereas NBP uses $J = 500$ and 1000 particles ($J = 5000$ would have resulted in excessive runtimes, due to the quadratic growth of NBP’s complexity with J). PM and NBP use $P = 2$ message passing iterations. We performed 100 simulation runs, each consisting of 100 time steps.

Fig. 10 shows the average runtime in seconds of all operations performed by the fusion center during one time step versus the network size $(|\tilde{\mathcal{A}}|, |\mathcal{O}|)$. The runtime was measured using MATLAB implementations of the algorithms on a single core of an Intel Xeon X5650 CPU. It is seen that the runtime of the BP-based methods PM and NBP scales linearly in the network size, whereas that of the particle filters SPF and UPF scales polynomially. This polynomial scaling of SPF and UPF is due to the fact that these filters perform operations involving matrices whose dimension increases with the network size. For




Fig. 11. Average localization RMSE versus network size.

the same number of particles J , PM always runs faster than UPF. Furthermore, for the considered parameters and network sizes, NBP has the highest runtime and SPF the lowest; however, for larger network sizes, the runtime of SPF and UPF will exceed that of PM and NBP due to the polynomial scaling characteristic of SPF and UPF. It is also seen that the runtime of PM is significantly lower than that of NBP.

Fig. 11 shows the average RMSE, i.e., the average of all MA self-localization and object localization errors, averaged over all time steps and simulation runs, versus the network size. SPF performs poorly since the numbers of particles it uses are not sufficient to properly represent the high-dimensional distributions. UPF performs much better; in fact, for $J=5000$, it outperforms all the other methods. This is due to a smart selection of the proposal distribution using the unscented transform. However, the RMSE of both SPF and UPF grows with the network size (although for UPF with $J=5000$, this is hardly visible in Fig. 11). In contrast, the RMSE of PM and NBP, which are both based on BP, does not depend on the network size. The RMSE of PM is only slightly higher than that of UPF with $J=5000$. The RMSE of NBP is higher than that of UPF and PM but considerably lower than that of SPF; it is reduced when J is increased from 500 to 1000. In contrast, the RMSE of PM is effectively equal for $J=1000$ and 5000. Thus, one can conclude that the performance of PM with $J=1000$ cannot be improved upon by increasing J , and the (very small) performance gap between UPF and PM is caused by the approximate nature of loopy BP. Finally, our simulations also showed that the performance gap between the BP-based methods, PM and NBP, and UPF is increased when the driving noise is increased and/or the measurement noise is decreased. This is again due to the smart selection of the proposal distribution in UPF.

These results demonstrate specific advantages of PM over particle filtering methods. In particular, PM has a very good performance-complexity tradeoff, and its scaling characteristic with respect to the network size is only linear.² For small networks, this may come at the cost of a very small performance loss relative to UPF (not, however, relative to SPF, which

²In a distributed implementation, the computational complexity of the consensus scheme depends on the number of consensus iterations, C , and, in the “consensus-over-weights” case, on the diameter of the communication graph, I . Therefore, the scaling might be slightly higher than linear.

performs much worse). For large networks, the performance of PM can be better than that of UPF, since for a fixed number of particles, the performance of UPF decreases with increasing network size (in Fig. 11, this is visible for $J=1000$ but only barely for $J=5000$).

An important further advantage of PM applies to distributed scenarios. Contrary to particle filters, PM facilitates a distributed implementation since it naturally distributes the computation effort among the agents. This distribution requires only communication with neighbor agents, and the communication cost is typically much smaller than for distributed particle filters. For example, in the case of the largest simulated network size of $|\tilde{\mathcal{A}}| + |\mathcal{O}| = 160$, for $J=1000$, each agent broadcasts 32000 real values per consensus iteration to its neighbors. In a distributed implementation of UPF, for proposal adaptation alone, each agent has to broadcast to its neighbors a covariance matrix of size 640×640 or, equivalently, 205120 real values per consensus iteration. Additional communication is required for other tasks, depending on the specific distributed particle filtering algorithm used [18]. Furthermore, for the considered joint CS-DT problem, UPF is ill-suited to large networks also because it involves the inversion and Cholesky decomposition of matrices whose dimension grows with the network size. In large networks, this may lead to numerical problems on processing units with limited dynamic range.

VIII. CONCLUSION

We proposed a Bayesian framework and methodology for distributed localization of cooperative agents and noncooperative objects in mobile networks, based on measurements between agents and objects and between different agents. Our work provides a consistent combination of cooperative self-localization (CS) and distributed object tracking (DT) for multiple mobile or static agents and objects. Starting from a factor graph formulation of the joint CS-DT problem, we developed a particle-based, distributed belief propagation (BP) message passing algorithm. This algorithm employs a consensus scheme for a distributed calculation of the product of the object messages. The proposed integration of consensus in particle-based BP solves the problem of accommodating noncooperative network nodes in distributed BP implementations. Thus, it may also be useful for other distributed inference problems.

A fundamental advantage of the proposed joint CS-DT method over both separate CS and DT and simultaneous localization and tracking (SLAT) is a probabilistic information transfer between CS and DT. This information transfer allows CS to support DT and vice versa. Our simulations demonstrated that this principle can result in significant improvements in both agent self-localization and object localization performance compared to state-of-the-art methods. Further advantages of our method are its low complexity and its very good scalability with respect to the network size. The computation effort is naturally distributed among the agents, using only a moderate amount of communication between neighboring agents. We note that the complexity can be reduced further

through an improved proposal distribution calculation that uses the *sigma point BP* technique introduced in [14].

The proposed framework and methodology can be extended to accommodate additional tasks (i.e., in addition to CS and DT) involving cooperative agents and/or noncooperative objects, such as distributed synchronization [38], [39] and cooperative mapping [58]. Another interesting direction for future work is an extension to scenarios involving an unknown number of objects [31], [59] and object-to-measurement association uncertainty [30], [31], [35], [36], [59].

REFERENCES

- [1] N. Patwari, J. N. Ash, S. Kyperountas, A. O. Hero III, R. L. Moses, and N. S. Correal, "Locating the nodes: Cooperative localization in wireless sensor networks," *IEEE Signal Process. Mag.*, vol. 22, pp. 54–69, Jul. 2005.
- [2] H. Wymeersch, J. Lien, and M. Z. Win, "Cooperative localization in wireless networks," *Proc. IEEE*, vol. 97, pp. 427–450, Feb. 2009.
- [3] J. Liu, M. Chu, and J. Reich, "Multitarget tracking in distributed sensor networks," *IEEE Signal Process. Mag.*, vol. 24, pp. 36–46, May 2007.
- [4] H. Aghajan and A. Cavallaro, *Multi-Camera Networks: Principles and Applications*. Burlington, MA: Academic Press, 2009.
- [5] P. Corke, T. Wark, R. Jurdak, W. Hu, P. Valencia, and D. Moore, "Environmental wireless sensor networks," *Proc. IEEE*, vol. 98, pp. 1903–1917, Nov. 2010.
- [6] F. Bullo, J. Cortes, and S. Martinez, *Distributed Control of Robotic Networks: A Mathematical Approach to Motion Coordination Algorithms*. Princeton, NJ: Princeton University Press, 2009.
- [7] A. Nayak and I. Stojmenović, *Wireless Sensor and Actuator Networks: Algorithms and Protocols for Scalable Coordination and Data Communication*. Hoboken, NJ: Wiley, 2010.
- [8] A. T. Ihler, J. W. Fisher, R. L. Moses, and A. S. Willsky, "Nonparametric belief propagation for self-localization of sensor networks," *IEEE J. Sel. Areas Commun.*, vol. 23, pp. 809–819, Apr. 2005.
- [9] C. Pedersen, T. Pedersen, and B. H. Fleury, "A variational message passing algorithm for sensor self-localization in wireless networks," in *Proc. IEEE ISIT-11*, Saint Petersburg, Russia, pp. 2158–2162, Aug. 2011.
- [10] J. Lien, J. Ferner, W. Srichavengsup, H. Wymeersch, and M. Z. Win, "A comparison of parametric and sample-based message representation in cooperative localization," *Int. J. Navig. Observ.*, 2012.
- [11] S. Li, M. Hedley, and I. B. Collings, "New efficient indoor cooperative localization algorithm with empirical ranging error model," *IEEE J. Sel. Areas Commun.*, vol. 33, pp. 1407–1417, Jul. 2015.
- [12] V. Savic and S. Zazo, "Reducing communication overhead for cooperative localization using nonparametric belief propagation," *IEEE Commun. Letters*, vol. 1, pp. 308–311, Aug. 2012.
- [13] T. Sathyan and M. Hedley, "Fast and accurate cooperative tracking in wireless networks," *IEEE Trans. Mobile Comput.*, vol. 12, pp. 1801–1813, Sep. 2013.
- [14] F. Meyer, O. Hlinka, and F. Hlawatsch, "Sigma point belief propagation," *IEEE Signal Process. Lett.*, vol. 21, pp. 145–149, Feb. 2014.
- [15] S. Van de Velde, G. Abreu, and H. Steendam, "Improved censoring and NLOS avoidance for wireless localization in dense networks," *IEEE J. Sel. Areas Commun.*, Nov. 2015.
- [16] A. G. O. Mutambara, *Decentralized Estimation and Control for Multi-sensor Systems*. Boca Raton, FL, USA: CRC Press, 1998.
- [17] T. Vercauteren and X. Wang, "Decentralized sigma-point information filters for target tracking in collaborative sensor networks," *IEEE Trans. Signal Process.*, vol. 53, pp. 2997–3009, Aug. 2005.
- [18] O. Hlinka, F. Hlawatsch, and P. M. Djuric, "Distributed particle filtering in agent networks: A survey, classification, and comparison," *IEEE Signal Process. Mag.*, vol. 30, pp. 61–81, Jan. 2013.
- [19] S. Farahmand, S. I. Roumeliotis, and G. B. Giannakis, "Set-membership constrained particle filter: Distributed adaptation for sensor networks," *IEEE Trans. Signal Process.*, vol. 59, pp. 4122–4138, Sep. 2011.
- [20] O. Hlinka, F. Hlawatsch, and P. M. Djuric, "Consensus-based distributed particle filtering with distributed proposal adaptation," *IEEE Trans. Signal Process.*, vol. 62, pp. 3029–3041, Jun. 2014.
- [21] M. Uney, D. Clark, and S. Julier, "Distributed fusion of PHD filters via exponential mixture densities," *IEEE J. Sel. Topics Signal Process.*, vol. 7, pp. 521–531, Jun. 2013.
- [22] G. Battistelli, L. Chisci, C. Fantacci, A. Farina, and A. Graziano, "Consensus CPHD filter for distributed multitarget tracking," *IEEE J. Sel. Topics Signal Process.*, vol. 7, pp. 508–520, Jun. 2013.
- [23] C. Fantacci, B.-N. Vo, B.-T. Vo, G. Battistelli, and L. Chisci, "Consensus labeled random finite set filtering for distributed multi-object tracking," 2015. Available online: <http://arxiv.org/abs/1501.01579>.
- [24] C. Taylor, A. Rahimi, J. Bachrach, H. Shrobe, and A. Grue, "Simultaneous localization, calibration, and tracking in an ad hoc sensor network," in *Proc. IPSN-06*, Nashville, TN, pp. 27–33, Apr. 2006.
- [25] V. Savic, H. Wymeersch, and E. Larsson, "Target tracking in confined environments with uncertain sensor positions," 2015. *IEEE Trans. Veh. Technol.*, to appear.
- [26] S. Funiak, C. Guestrin, M. Paskin, and R. Sukthankar, "Distributed localization of networked cameras," in *Proc. IPSN-06*, Nashville, TN, pp. 34–42, Apr. 2006.
- [27] N. Kantas, S. Singh, and A. Doucet, "Distributed maximum likelihood for simultaneous self-localization and tracking in sensor networks," *IEEE Trans. Signal Process.*, vol. 60, pp. 5038–5047, Oct. 2012.
- [28] J. Teng, H. Snoussi, C. Richard, and R. Zhou, "Distributed variational filtering for simultaneous sensor localization and target tracking in wireless sensor networks," *IEEE Trans. Veh. Technol.*, vol. 61, pp. 2305–2318, Jun. 2012.
- [29] M. Uney, B. Mulgrew, and D. Clark, "Cooperative sensor localisation in distributed fusion networks by exploiting non-cooperative targets," in *Proc. IEEE SSP-14*, Gold Coast, Australia, pp. 516–519, Jun. 2014.
- [30] Y. Bar-Shalom, P. Willett, and X. Tian, *Tracking and Data Fusion: A Handbook of Algorithms*. Storrs, CT: Yaakov Bar-Shalom, 2011.
- [31] R. P. S. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Boston, MA: Artech House, 2007.
- [32] F. Meyer, E. Riegler, O. Hlinka, and F. Hlawatsch, "Simultaneous distributed sensor self-localization and target tracking using belief propagation and likelihood consensus," in *Proc. 46th Asilomar Conf. Sig., Syst., Comp.*, Pacific Grove, CA, pp. 1212–1216, Nov. 2012.
- [33] F. Meyer, F. Hlawatsch, and H. Wymeersch, "Cooperative simultaneous localization and tracking (CoSLAT) with reduced complexity and communication," in *Proc. IEEE ICASSP-13*, Vancouver, Canada, pp. 4484–4488, May 2013.
- [34] X. R. Li and V. P. Jilkov, "Survey of maneuvering target tracking. Part I: Dynamic models," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, pp. 1333–1364, Oct. 2003.
- [35] J. L. Williams and R. Lau, "Approximate evaluation of marginal association probabilities with belief propagation," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 50, pp. 2942–2959, Oct. 2014.
- [36] F. Meyer, P. Braca, P. Willett, and F. Hlawatsch, "Scalable multitarget tracking using multiple sensors: A belief propagation approach," in *Proc. FUSION-15*, Washington D.C., USA, pp. 1778–1785, Jul. 2015.
- [37] Y.-C. Wu, Q. M. Chaudhari, and E. Serpedin, "Clock synchronization of wireless sensor networks," *IEEE Signal Process. Mag.*, vol. 28, pp. 124–138, Jan. 2011.
- [38] B. Etzlinger, H. Wymeersch, and A. Springer, "Cooperative synchronization in wireless networks," *IEEE Trans. Signal Process.*, vol. 62, pp. 2837–2849, Jun. 2014.
- [39] F. Meyer, B. Etzlinger, F. Hlawatsch, and A. Springer, "A distributed particle-based belief propagation algorithm for cooperative simultaneous localization and synchronization," in *Proc. Asilomar Conf. Sig., Syst., Comput.*, Pacific Grove, CA, pp. 527–531, Nov. 2013.
- [40] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Upper Saddle River, NJ: Prentice-Hall, 1993.
- [41] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. Inf. Theory*, vol. 47, pp. 498–519, Feb. 2001.
- [42] H.-A. Loeliger, "An introduction to factor graphs," *IEEE Trans. Signal Process.*, vol. 21, pp. 28–41, Jan. 2004.
- [43] M. J. Wainwright and M. I. Jordan, "Graphical models, exponential families, and variational inference," *Found. Trends Mach. Learn.*, vol. 1, Jan. 2008.
- [44] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*. New York, NY: Springer, 2001.
- [45] M. Briars, A. Doucet, and S. S. Singh, "Sequential auxiliary particle belief propagation," in *Proc. FUSION-05*, Philadelphia, PA, USA, Jul. 2005.
- [46] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Trans. Signal Process.*, vol. 50, pp. 174–188, Feb. 2002.
- [47] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Norwood, MA: Artech House, 2004.

- [48] F. Meyer, O. Hlinka, H. Wymeersch, E. Riegler, and F. Hlawatsch, "Distributed localization and tracking of mobile networks including noncooperative objects — Extended version." Available online: <http://arxiv.org/abs/1403.1824>.
- [49] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proc. IEEE*, vol. 95, pp. 215–233, Jan 2007.
- [50] A. G. Dimakis, S. Kar, J. M. F. Moura, M. G. Rabbat, and A. Scaglione, "Gossip algorithms for distributed signal processing," *Proc. IEEE*, vol. 98, pp. 1847–1864, Nov. 2010.
- [51] T.-D. Pham, H. Q. Ngo, V.-D. Le, S. Lee, and Y.-K. Lee, "Broadcast gossip based distributed hypothesis testing in wireless sensor networks," in *Proc. ATC-09*, Haiphong, Vietnam, pp. 84–87, 2009.
- [52] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Trans. Autom. Control*, vol. 49, pp. 1520–1533, Sep. 2004.
- [53] V. Savic, H. Wymeersch, and S. Zazo, "Belief consensus algorithms for fast distributed target tracking in wireless sensor networks," *Signal Processing*, vol. 95, pp. 149–160, 2014.
- [54] M. Fröhle and H. Wymeersch, "On the separation of timescales in radio-based positioning," in *Proc. ICL-GNSS-15*, Gothenburg, Sweden, Jun. 2015.
- [55] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," in *Proc. IEEE CDC-03*, vol. 5, Maui, HI, USA, pp. 4997–5002, Dec. 2003.
- [56] R. van der Merwe, A. Doucet, N. de Freitas, and E. Wan, "The unscented particle filter." Technical Report, Cambridge University, Cambridge, UK, 2000. Available online: <http://www.markirwin.net/stat220/Refs/upf2000.pdf>.
- [57] R. Balakrishnan and K. Ranganathan, *A Textbook of Graph Theory*. New York, NY: Springer, 2012.
- [58] G. Dedeoglu and G. S. Sukhatme, "Landmark-based matching algorithm for cooperative mapping by autonomous robots," in *Distributed Autonomous Robotic Systems 4* (L. E. Parker, G. Bekey, and J. Barhen, eds.), ch. 7, pp. 251–260, Tokyo, Japan: Springer, 2000.
- [59] P. Braca, S. Marano, V. Matta, and P. Willett, "Asymptotic efficiency of the PHD in multitarget/multisensor estimation," *IEEE J. Sel. Topics Signal Process.*, vol. 7, pp. 553–564, Jun. 2013.