

Distributed Estimation with Information-Seeking Control in Agent Networks

Florian Meyer, *Student Member, IEEE*, Henk Wymeersch, *Member, IEEE*, Markus Fröhle, and Franz Hlawatsch, *Fellow, IEEE*

Abstract—We introduce a distributed, cooperative framework and method for Bayesian estimation and control in decentralized agent networks. Our framework combines joint estimation of time-varying global and local states with information-seeking control optimizing the behavior of the agents. It is suited to nonlinear and non-Gaussian problems and, in particular, to location-aware networks. For cooperative estimation, a combination of belief propagation message passing and consensus is used. For cooperative control, the negative posterior joint entropy of all states is maximized via a gradient ascent. The estimation layer provides the control layer with probabilistic information in the form of sample representations of probability distributions. Simulation results demonstrate intelligent behavior of the agents and excellent estimation performance for a simultaneous self-localization and target tracking problem. In a cooperative localization scenario with only one anchor, mobile agents can localize themselves after a short time with an accuracy that is higher than the accuracy of the performed distance measurements.

Index Terms—Agent networks, distributed estimation, distributed control, information-seeking control, distributed target tracking, cooperative localization, belief propagation, message passing, consensus, sensor networks, sequential estimation.

I. INTRODUCTION

A. Motivation and State of the Art

Recent research on distributed estimation and control in mobile agent networks [1]–[4] has frequently been motivated by location-aware scenarios and problems including environmental and agricultural monitoring [5], healthcare monitoring [6], target tracking [7], pollution source localization [8], chemical plume tracking [3], and surveillance [9]. The agents in a mobile agent network are generally equipped with sensors, wireless communication interfaces, a processing unit, and actuators, all together forming a cyber-physical system [10], [11] with a tight coupling between sensing, computing, communication, and control. A common task in mobile agent networks is seeking information, either about external phenomena or about the network itself. This task relies on *estimation* (quantifying, fusing, and disseminating

information) and *control* (configuring the network to increase information). A common theme in previous works is the reliance on position information for estimation and/or control.

Estimation methods for mobile agent networks (our focus will be on distributed Bayesian estimation) address estimation of common global states [7], [12]–[16], estimation of local states [17]–[22], or combined estimation of local and global states [23]–[25]. In the first case, the agents obtain local measurements with respect to external objects or the surrounding environment, which are fused across the network. Global fusion methods that require only local communication include consensus [26] and gossip [27]. Example applications are distributed target tracking [7], cooperative exploration of the environment [15], and chemical plume tracking [3]. In the second case (estimation of local states), the agents cooperate such that each agent is better able to estimate its own local state. Here, the dimensionality of the total state grows with the network size, which leads to more complex factorizations of the joint posterior probability density function (pdf). When the factor graph [28] of this factorization matches the network topology, efficient message passing methods for distributed inference can be used, such as the belief propagation (BP) [28] and mean field [29] methods. Example applications are cooperative localization [17], synchronization [19], [22], and simultaneous localization and synchronization [20], [21]. In the third case (estimation of both global states and local states), a message passing algorithm can be combined with a network-wide information dissemination technique. An example application is cooperative simultaneous self-localization and target tracking [23], [25].

In many cooperative estimation scenarios, it is advantageous to control certain properties of the agent network, such as the agent positions or the measurement characteristics (“controlled sensing”) [1]–[9]. In particular, here we will address the problem of combining distributed estimation and distributed control in mobile agent networks. We will limit our discussion to *information-seeking* control, which seeks to maximize the information carried by the measurements of all agents about the global and/or local states to be estimated. The use of information measures for the control of a single agent or a network of agents was introduced in [30] and [31], respectively. Suitable measures of information include negative posterior entropy [32], mutual information [32], and scalar-valued functions of the Fisher information matrix [33]. In particular, the determinant, trace, and spectral norm of the Fisher information matrix were considered in [34], where the control objective is to maximize the information related to

Final manuscript, April 15, 2015

F. Meyer and F. Hlawatsch are with the Institute of Telecommunications, Vienna University of Technology, Vienna, Austria (e-mail: {florian.meyer, franz.hlawatsch}@nt.tuwien.ac.at). H. Wymeersch and M. Fröhle are with the Department of Signals and Systems, Chalmers University of Technology, Gothenburg, Sweden (e-mail: {henkw, frohle}@chalmers.se). This work was supported by the Austrian Science Fund (FWF) under Grants S10603 and P27370 and by the European Commission under ERC Grant No. 258418 (COOPNET), EU FP7 Marie Curie Initial Training Network MULTI-POS under Grant 316528, and the Newcom# Network of Excellence in Wireless Communications. Parts of this paper were previously presented at ICASSP 2015, Brisbane, Australia, April 2015.

the positions of the agents and of a target. The maximization of negative posterior entropy was considered in [12]–[16]. In [14], a central controller steers agents with known positions along the gradient of negative posterior entropy to optimally sense a global state. A distributed solution for global state estimation was proposed in [12], [13] based on a pairwise neighboring-agents approximation of mutual information and in [15], [16] by using a consensus algorithm. However, the methods proposed in [12]–[16] did not use BP, did not allow for multiple time-varying states, and did not include estimation of local (controlled) states.

B. Contribution and Paper Organization

Here, we present a unified Bayesian framework and method for (i) distributed, cooperative joint estimation of time-varying global and local states and (ii) distributed, cooperative information-seeking control. Our framework and method are suited to nonlinear and non-Gaussian problems, they require only communication with neighboring agents, and they are able to cope with a changing network topology. Thereby, they are particularly suited to localization and tracking tasks in location-aware scenarios involving mobile networks and nonlinear models.

For distributed estimation, following [24], [25], we combine BP message passing, consensus, and sample-based representations of the involved probability distributions. For distributed control, we define a global (holistic) objective function as the negative joint posterior entropy of all states in the network at the next time step conditioned on all measurements at the next time step. This objective function is optimized jointly by all agents via a gradient ascent. This reduces to the evaluation of local gradients at each agent, which is performed by using Monte Carlo integration based on the sample representations provided by the estimation stage and a distributed evaluation of the joint (networkwide) likelihood function. Our method advances beyond [12]–[16] in the following respects:

- It constitutes a more general information-maximizing control framework based on BP for estimation problems involving multiple time-varying states.
- It includes estimation of the local (controlled) states of the agents, thus enabling its use in a wider range of applications.

Contrary to [24], [25], which introduces the distributed joint *estimation* of time-varying local and global states in agent networks, here we focus on the information-maximizing *control* of the agents. Our main contribution is a derivation and sample-based formulation of a new information-seeking controller that maximizes the negative joint posterior entropy of time-varying local and global states. Compared to the information-seeking controller proposed in [12]–[16], where maximization of the negative posterior entropy reduces to maximization of the mutual information between observations and states, our controller includes an additional term that arises because the posterior entropy involves also the local (controlled) states of the agents. Due to this more general formulation, our controller is suited to decentralized estimation tasks where the agents cooperatively infer also their own states.




Fig. 1. Agent network with CAs and targets. The neighborhood sets $\mathcal{A}_l^{(n)}$, $\mathcal{C}_l^{(n)}$, $\mathcal{T}_l^{(n)}$, and $\mathcal{C}_m^{(n)}$ for one CA $l \in \mathcal{C}$ and one target $m \in \mathcal{T}$ are also shown.

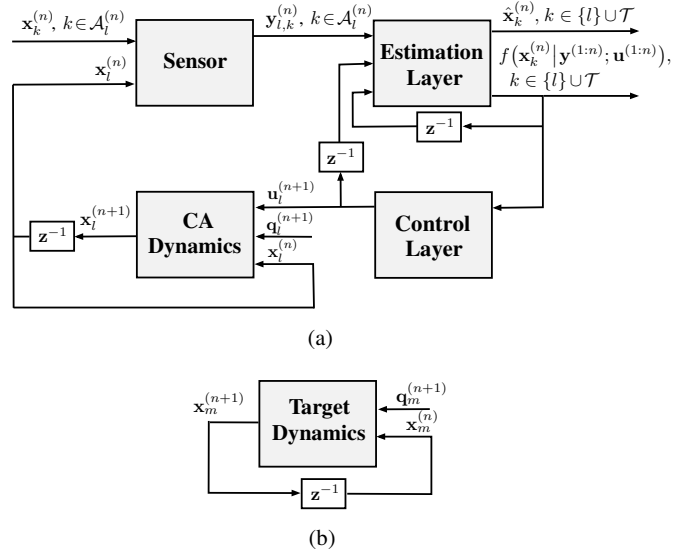


Fig. 2. Block diagram of the overall “signal processing system” for (a) a CA $l \in \mathcal{C}$ and (b) a target $m \in \mathcal{T}$.

This paper is organized as follows. In Section II, the system model is described and the joint estimation and control problem is formulated. Section III reviews joint local and global state estimation [25]. In Section IV, we introduce the proposed gradient-based controller. The distributed computation of the gradient is discussed in Sections V–VII. Section VIII considers two special cases of the joint estimation and control framework. Finally, in Section IX, we present simulation results demonstrating the performance of our method for a simultaneous self-localization and target tracking problem.

II. SYSTEM MODEL AND PROBLEM FORMULATION

We consider a network of mobile agents $k \in \mathcal{A}$ as shown in Fig. 1. The set of all agents, \mathcal{A} , consists of the set of cooperative agents (CAs), $\mathcal{C} \subseteq \mathcal{A}$, and the set of “targets,” $\mathcal{T} = \mathcal{A} \setminus \mathcal{C}$. Here, a target may be anything that does not cooperate and cannot be controlled, such as a noncooperative agent or a relevant feature of the environment. We will typically use the indices $k \in \mathcal{A}$, $l \in \mathcal{C}$, and $m \in \mathcal{T}$ to denote a generic agent, a CA, and a target, respectively. A block diagram of the overall “signal processing system” is shown in Fig. 2 for a CA (including the estimation and control layers of the proposed method) and for a target. This system is described below.

A. Agent States and Sensor Measurements

The *state* of agent $k \in \mathcal{A}$ at discrete time $n \in \{0, 1, \dots\}$ is denoted by the vector $\mathbf{x}_k^{(n)}$. For example, in a localization scenario, $\mathbf{x}_k^{(n)}$ may consist of the current position and motion-related quantities such as velocity, acceleration, and angular velocity [35]. The states evolve according to

$$\mathbf{x}_l^{(n)} = g_l(\mathbf{x}_l^{(n-1)}, \mathbf{u}_l^{(n)}, \mathbf{q}_l^{(n)}), \quad l \in \mathcal{C} \quad (1)$$

or

$$\mathbf{x}_m^{(n)} = g_m(\mathbf{x}_m^{(n-1)}, \mathbf{q}_m^{(n)}), \quad m \in \mathcal{T}, \quad (2)$$

where $g_k(\cdot)$ is a possibly nonlinear function, $\mathbf{q}_k^{(n)}$ is process (driving) noise, and $\mathbf{u}_l^{(n)} \in \mathcal{U}_l$ is a deterministic control vector that controls the l th CA. Since $\mathbf{u}_l^{(n)}$ is deterministic [36, Chap. 5], it is either completely unknown (before it is determined) or perfectly known (after being determined by the control layer). Note that also targets may have control variables. However, as these are hidden from the CAs, we will subsume any control for target m in the noise $\mathbf{q}_m^{(n)}$. For the derivation of the controller, we assume that for $l \in \mathcal{C}$, $g_l(\mathbf{x}_l^{(n-1)}, \mathbf{u}_l^{(n)}, \mathbf{q}_l^{(n)})$ is bijective with respect to $\mathbf{x}_l^{(n-1)}$ and differentiable with respect to $\mathbf{u}_l^{(n)}$. The statistical relation between $\mathbf{x}_k^{(n-1)}$ and $\mathbf{x}_k^{(n)}$ as defined by (1) or (2) can also be described by the *state-transition pdf* $f(\mathbf{x}_l^{(n)} | \mathbf{x}_l^{(n-1)}; \mathbf{u}_l^{(n)})$ for $l \in \mathcal{C}$ and $f(\mathbf{x}_m^{(n)} | \mathbf{x}_m^{(n-1)})$ for $m \in \mathcal{T}$.

The measurement and communication topology of the network is described by the neighborhood sets $\mathcal{C}_l^{(n)}$, $\mathcal{T}_l^{(n)}$, and $\mathcal{A}_l^{(n)}$ as follows. CA l acquires a measurement $\mathbf{y}_{l,l'}$ relative to CA l' if $l' \in \mathcal{C}_l^{(n)}$. This relation is symmetric, i.e., $l' \in \mathcal{C}_l^{(n)}$ implies $l \in \mathcal{C}_{l'}^{(n)}$. It is assumed that CAs that acquire measurements relative to each other are able to communicate, i.e., to transmit data via a communication link. Furthermore, CA $l \in \mathcal{C}$ acquires a measurement $\mathbf{y}_{l,m}^{(n)}$ relative to target m if $m \in \mathcal{T}_l^{(n)} \subseteq \mathcal{T}$. The targets are noncooperative in that they do not communicate and do not acquire any measurements. We also define $\mathcal{A}_l^{(n)} \triangleq \mathcal{C}_l^{(n)} \cup \mathcal{T}_l^{(n)}$. Finally, the set $\mathcal{C}_m^{(n)}$ contains all CAs measuring target m , i.e., all $l \in \mathcal{C}$ with $m \in \mathcal{T}_l^{(n)}$. The sets $\mathcal{C}_l^{(n)}$, $\mathcal{T}_l^{(n)}$, $\mathcal{A}_l^{(n)}$, and $\mathcal{C}_m^{(n)}$ are generally time-dependent. An example of a measurement and communication topology is given in Fig. 1.

We consider “pairwise” measurements¹ $\mathbf{y}_{l,k}^{(n)}$ that depend on the state $\mathbf{x}_l^{(n)}$ of a measuring CA $l \in \mathcal{C}$ and the state $\mathbf{x}_k^{(n)}$ of a measured agent (CA or target) $k \in \mathcal{A}_l^{(n)}$ according to

$$\mathbf{y}_{l,k}^{(n)} = d_l(\mathbf{x}_l^{(n)}, \mathbf{x}_k^{(n)}, \mathbf{v}_{l,k}^{(n)}), \quad l \in \mathcal{C}, k \in \mathcal{A}_l^{(n)}, \quad (3)$$

where $d_l(\cdot)$ is a possibly nonlinear function and $\mathbf{v}_{l,k}^{(n)}$ is measurement noise. An example is the scalar measurement

$$y_{l,k}^{(n)} = \|\mathbf{x}_l^{(n)} - \mathbf{x}_k^{(n)}\| + v_{l,k}^{(n)}, \quad (4)$$

where $\mathbf{x}_k^{(n)}$ equals the position of agent k and, hence, $\|\mathbf{x}_l^{(n)} - \mathbf{x}_k^{(n)}\|$ is the spatial distance between agents l and k . The

¹The proposed framework can be easily extended to self-measurements (measurements that involve only the own state) and cluster measurements (measurements that involve the states of several other agents).

statistical relation between the measurement $\mathbf{y}_{l,k}^{(n)}$ and the involved states $\mathbf{x}_l^{(n)}$ and $\mathbf{x}_k^{(n)}$ is also described by the *local likelihood function* $f(\mathbf{y}_{l,k}^{(n)} | \mathbf{x}_l^{(n)}, \mathbf{x}_k^{(n)})$. For the derivation of the controller, $f(\mathbf{y}_{l,k}^{(n)} | \mathbf{x}_l^{(n)}, \mathbf{x}_k^{(n)})$ is assumed differentiable with respect to $\mathbf{x}_l^{(n)}$.

We also make the following assumptions. The number of targets is known, and the targets can be identified by the CAs, i.e., target-to-measurement assignments are known. Furthermore, CA l knows the state evolution models $g_k(\cdot)$ and process noise pdfs $f(\mathbf{q}_k^{(n)})$ for $k \in \{l\} \cup \mathcal{C}_l^{(n)} \cup \mathcal{T}$; the initial prior pdfs of the agent states, $f(\mathbf{x}_k^{(0)})$, for $k \in \{l\} \cup \mathcal{C}_l^{(0)} \cup \mathcal{T}$; the measurement models $d_{l'}(\cdot)$ for $l' \in \{l\} \cup \mathcal{C}_l^{(n)}$; and the measurement noise pdfs $f(\mathbf{v}_{l',k}^{(n)})$, $k \in \mathcal{A}_{l'}^{(n)}$ and $f(\mathbf{v}_{l',l}^{(n)})$, $l' \in \mathcal{C}_l^{(n)}$.

B. Problem Formulation

The following tasks are to be performed at each time n :

- 1) Each CA $l \in \mathcal{C}$ estimates the states $\mathbf{x}_k^{(n)}$, $k \in \{l\} \cup \mathcal{T}$ (i.e., its own local state and the states of all targets) from prior information and all past and present measurements in the network.
- 2) The state of each CA is controlled such that the negative joint posterior entropy of all states in the network at the next time, conditioned on all measurements in the network at the next time, is maximized.

We solve these two problems in a distributed and recursive manner. Our method consists of an *estimation layer* and a *control layer*, as shown in Fig. 2(a). In the estimation layer, CA l computes an approximation of the marginal posterior pdfs of the states $\mathbf{x}_k^{(n)}$, $k \in \{l\} \cup \mathcal{T}$ given all the past and present measurements and control vectors in the entire network. In the control layer, CA l uses these marginal posteriors and the statistical model to determine a quasi-optimal control variable $\mathbf{u}_l^{(n+1)}$. In both layers, the CAs communicate with neighbor CAs.

III. ESTIMATION LAYER

The estimation layer performs distributed estimation of the local and global states by using the BP- and consensus-based method introduced in [24], [25]. We will review this method in our present context. Let us denote by $\mathbf{x}^{(n)} \triangleq [\mathbf{x}_k^{(n)}]_{k \in \mathcal{A}}$, $\mathbf{u}^{(n)} \triangleq [\mathbf{u}_l^{(n)}]_{l \in \mathcal{C}}$, and $\mathbf{y}^{(n)} \triangleq [\mathbf{y}_{l,k}^{(n)}]_{l \in \mathcal{C}, k \in \mathcal{A}_l^{(n)}}$ the stacked vectors of, respectively, all states, control variables, and measurements at time n . Furthermore, let $\mathbf{x}^{(1:n)} \triangleq [\mathbf{x}^{(1)\top}, \dots, \mathbf{x}^{(n)\top}]^\top$, $\mathbf{u}^{(1:n)} \triangleq [\mathbf{u}^{(1)\top}, \dots, \mathbf{u}^{(n)\top}]^\top$, and $\mathbf{y}^{(1:n)} \triangleq [\mathbf{y}^{(1)\top}, \dots, \mathbf{y}^{(n)\top}]^\top$. Each CA $l \in \mathcal{C}$ estimates its local state $\mathbf{x}_l^{(n)}$ and all the target states $\mathbf{x}_m^{(n)}$, $m \in \mathcal{T}$ from the measurements of all CAs up to time n , $\mathbf{y}^{(1:n)}$. This estimation is based on the posteriors $f(\mathbf{x}_k^{(n)} | \mathbf{y}^{(1:n)}; \mathbf{u}^{(1:n)})$, $k \in \{l\} \cup \mathcal{T}$, which are marginals of the joint posterior $f(\mathbf{x}^{(1:n)} | \mathbf{y}^{(1:n)}; \mathbf{u}^{(1:n)})$.

Using Bayes' rule and common assumptions [17], the joint posterior can be factorized as

$$f(\mathbf{x}^{(1:n)} | \mathbf{y}^{(1:n)}; \mathbf{u}^{(1:n)})$$

$$\begin{aligned} & \propto \left(\prod_{k \in \mathcal{A}} f(\mathbf{x}_k^{(0)}) \right) \prod_{n'=1}^n \left(\prod_{m \in \mathcal{T}} f(\mathbf{x}_m^{(n')} | \mathbf{x}_m^{(n'-1)}) \right) \\ & \quad \times \prod_{l \in \mathcal{C}} f(\mathbf{x}_l^{(n')} | \mathbf{x}_l^{(n'-1)}; \mathbf{u}_l^{(n')}) \prod_{k' \in \mathcal{A}_l^{(n')}} f(\mathbf{y}_{l,k'}^{(n')} | \mathbf{x}_l^{(n')}, \mathbf{x}_{k'}^{(n')}). \end{aligned} \quad (5)$$

The marginal posterior of state $\mathbf{x}_k^{(n)}$ is then given by

$$f(\mathbf{x}_k^{(n)} | \mathbf{y}^{(1:n)}; \mathbf{u}^{(1:n)}) = \int f(\mathbf{x}^{(1:n)} | \mathbf{y}^{(1:n)}; \mathbf{u}^{(1:n)}) d\mathbf{x}_{\sim k, n}^{(1:n)}, \quad (6)$$

where $\mathbf{x}_{\sim k, n}^{(1:n)}$ denotes $\mathbf{x}^{(1:n)}$ with $\mathbf{x}_k^{(n)}$ removed. Using $f(\mathbf{x}_k^{(n)} | \mathbf{y}^{(1:n)}; \mathbf{u}^{(1:n)})$, the minimum mean-square error (MMSE) estimator [33] of $\mathbf{x}_k^{(n)}$ is obtained as

$$\hat{\mathbf{x}}_{k, \text{MMSE}}^{(n)} \triangleq \int \mathbf{x}_k^{(n)} f(\mathbf{x}_k^{(n)} | \mathbf{y}^{(1:n)}; \mathbf{u}^{(1:n)}) d\mathbf{x}_k^{(n)}, \quad k \in \mathcal{A}. \quad (7)$$

A. Sequential Calculation

For a review of the sequential calculation of (6) proposed in [24], [25], we now switch to the following simplified notation for the sake of readability. In the conditions of the various conditional pdfs, we omit the measurements up to time $n-1$, i.e., $\mathbf{y}^{(1:n-1)}$, and the control vectors up to time n , i.e., $\mathbf{u}^{(1:n)}$, because $\mathbf{y}^{(1:n-1)}$ has already been observed and $\mathbf{u}^{(1:n)}$ has already been determined; hence both are considered fixed. Furthermore, we suppress the time index n , and we write the current and previous states of CA l as $\mathbf{x}_l^{(n)} = \mathbf{x}_l$ and $\mathbf{x}_l^{(n-1)} = \mathbf{x}_l^-$, respectively. Similarly, we write $\mathbf{u}_l^{(n+1)} = \mathbf{u}_l^+$. For sequential calculation of (6), CA $l \in \mathcal{C}$ employs a basic Bayesian recursive filtering method [17] consisting of a prediction step and a correction step. In the prediction step, CA $l \in \mathcal{C}$ computes a predictive posterior of its current state,

$$f(\mathbf{x}_l) = \int f(\mathbf{x}_l | \mathbf{x}_l^-) f(\mathbf{x}_l^-) d\mathbf{x}_l^-. \quad (8)$$

Here, $f(\mathbf{x}_l)$ and $f(\mathbf{x}_l^-)$ are short for $f(\mathbf{x}_l^{(n)} | \mathbf{y}^{(1:n-1)}; \mathbf{u}^{(1:n)})$ and $f(\mathbf{x}_l^{(n-1)} | \mathbf{y}^{(1:n-1)}; \mathbf{u}^{(1:n-1)})$, respectively. Furthermore, CA l computes predictive posteriors of the target states

$$f(\mathbf{x}_m) = \int f(\mathbf{x}_m | \mathbf{x}_m^-) f(\mathbf{x}_m^-) d\mathbf{x}_m^-, \quad m \in \mathcal{T}. \quad (9)$$

In the correction step, CA l determines the marginal posteriors $f(\mathbf{x}_l | \mathbf{y})$ and $f(\mathbf{x}_m | \mathbf{y})$, $m \in \mathcal{T}$, which are given by

$$f(\mathbf{x}_k | \mathbf{y}) \propto \int \prod_{k' \in \mathcal{A}} f(\mathbf{x}_{k'}) \prod_{l' \in \mathcal{C}} \prod_{k_1 \in \mathcal{A}_{l'}} f(\mathbf{y}_{l', k_1} | \mathbf{x}_{l'}, \mathbf{x}_{k_1}) d\mathbf{x}_{\sim k}, \quad k \in \{l\} \cup \mathcal{T}. \quad (10)$$

Here, $\mathbf{x}_{\sim k}$ denotes \mathbf{x} with \mathbf{x}_k removed and $f(\mathbf{x}_k | \mathbf{y})$ is short for $f(\mathbf{x}_k^{(n)} | \mathbf{y}^{(1:n)}; \mathbf{u}^{(1:n)})$. As shown in Fig. 2(a), these marginal posteriors are handed over to the control layer, which determines the control input for the next time step, \mathbf{u}_l^+ .

B. BP Message Passing and Consensus

Calculation of $f(\mathbf{x}_k | \mathbf{y})$ according to (10) is generally infeasible, due to the reliance on nonlocal information and the inherent complexity of the marginalization process. Based on the factorization of the joint posterior in (5), a computationally feasible approximation of (10) is provided by a distributed,

cooperative algorithm that combines BP message passing and the average consensus scheme [24], [25]. This algorithm computes an approximation (“belief”) $b(\mathbf{x}_k) \approx f(\mathbf{x}_k | \mathbf{y})$ in an iterative manner, using only communication with neighbor CAs $l' \in \mathcal{C}_l$. Its complexity scales only linearly with the number of states in the network, $|\mathcal{A}|$ (for a fixed number of iterations).

According to [24], [25], the belief of local state \mathbf{x}_l at message passing iteration $p \in \{1, \dots, P\}$ is given by

$$\begin{aligned} b^{(p)}(\mathbf{x}_l) & \propto f(\mathbf{x}_l) \prod_{l' \in \mathcal{C}_l} \int f(\mathbf{y}_{l, l'} | \mathbf{x}_l, \mathbf{x}_{l'}) b^{(p-1)}(\mathbf{x}_{l'}) d\mathbf{x}_{l'} \\ & \quad \times \prod_{m \in \mathcal{T}_l} \int f(\mathbf{y}_{l, m} | \mathbf{x}_l, \mathbf{x}_m) \psi_{m \rightarrow l}^{(p-1)}(\mathbf{x}_m) d\mathbf{x}_m, \end{aligned} \quad (11)$$

which is initialized as $b^{(0)}(\mathbf{x}_l) = f(\mathbf{x}_l)$. Similarly, the belief of target state \mathbf{x}_m at message passing iteration p is given by

$$b^{(p)}(\mathbf{x}_m) \propto f(\mathbf{x}_m) \prod_{l \in \mathcal{C}_m} \int f(\mathbf{y}_{l, m} | \mathbf{x}_l, \mathbf{x}_m) \psi_{l \rightarrow m}^{(p-1)}(\mathbf{x}_l) d\mathbf{x}_l, \quad (12)$$

with initialization $b^{(0)}(\mathbf{x}_m) = f(\mathbf{x}_m)$. Here, $\psi_{m \rightarrow l}^{(p-1)}(\mathbf{x}_m)$ and $\psi_{l \rightarrow m}^{(p-1)}(\mathbf{x}_l)$ are the “extrinsic informations” from target m to CA l and from CA l to target m , respectively, at the previous message passing iteration $p-1$. These extrinsic informations are calculated recursively as

$$\psi_{m \rightarrow l}^{(p)}(\mathbf{x}_m) = \frac{b^{(p)}(\mathbf{x}_m)}{\int f(\mathbf{y}_{l, m} | \mathbf{x}_l, \mathbf{x}_m) \psi_{l \rightarrow m}^{(p-1)}(\mathbf{x}_l) d\mathbf{x}_l} \quad (13)$$

$$\psi_{l \rightarrow m}^{(p)}(\mathbf{x}_l) = \frac{b^{(p)}(\mathbf{x}_l)}{\int f(\mathbf{y}_{l, m} | \mathbf{x}_l, \mathbf{x}_m) \psi_{m \rightarrow l}^{(p-1)}(\mathbf{x}_m) d\mathbf{x}_m}, \quad (14)$$

with initialization $\psi_{m \rightarrow l}^{(0)}(\mathbf{x}_m) = f(\mathbf{x}_m)$ and $\psi_{l \rightarrow m}^{(0)}(\mathbf{x}_l) = f(\mathbf{x}_l)$, respectively. In (11), the beliefs $b^{(p-1)}(\mathbf{x}_{l'})$ of neighbor CAs $l' \in \mathcal{C}_l$ are used instead of the extrinsic informations. This is part of the *sum-product algorithm over a wireless network* (SPAWN) scheme for cooperative localization [17], which results from a specific choice of the message schedule in loopy BP and has been observed to provide highly accurate estimates [17], [24], [25], [37].

A sample-based distributed implementation of (11)–(14) has been proposed in [24], [25]. A problem for a distributed implementation is that the products $\prod_{l \in \mathcal{C}_m} \int f(\mathbf{y}_{l, m} | \mathbf{x}_l, \mathbf{x}_m) \psi_{l \rightarrow m}^{(p-1)}(\mathbf{x}_l) d\mathbf{x}_l$, $m \in \mathcal{T}$ involved in (12) are not available at the CAs. However, an approximation of these products can be provided to each CA in a distributed manner (i.e., using only local communications) either by a consensus algorithm performed in parallel for each sample weight [25], [38], [39] or by the likelihood consensus scheme [7], [24], [40]. For the calculations in the control layer—to be described in Sections IV–VII—all CAs require a common set of samples representing the target beliefs $b^{(p)}(\mathbf{x}_m)$. This can be ensured by additionally using, e.g., a max-consensus and providing all CAs with the same seed for random number generation [39], [41].

The output of the estimation layer is the set of beliefs $b^{(P)}(\mathbf{x}_k)$, $k \in \{l\} \cup \mathcal{T}$ at the final message passing iteration $p = P$. These beliefs are handed over to the control layer, which

calculates the control variables \mathbf{u}_l^+ (see Fig. 2(a)). Each belief $b^{(P)}(\mathbf{x}_k)$ is represented by J samples $\{\mathbf{x}_k^{(j)}\}_{j=1}^J$, which is briefly denoted as $\{\mathbf{x}_k^{(j)}\}_{j=1}^J \sim b^{(P)}(\mathbf{x}_k)$. From these samples, an approximation of the MMSE estimate (7) is calculated as

$$\hat{\mathbf{x}}_k = \frac{1}{J} \sum_{j=1}^J \mathbf{x}_k^{(j)}.$$

IV. CONTROL LAYER

Next, we present the information-seeking controller. We temporarily revert to the full notation.

A. Objective Function and Controller

According to our definition at the beginning of Section III, the vector comprising all the measurements in the network at the next time $n+1$ is $\mathbf{y}^{(n+1)} = [\mathbf{y}_{l,k}^{(n+1)}]_{l \in \mathcal{C}, k \in \mathcal{A}_l^{(n+1)}}$. However, in this definition of $\mathbf{y}^{(n+1)}$, we now formally replace $\mathcal{A}_l^{(n+1)}$ by $\mathcal{A}_l^{(n)}$ since at the current time n , the sets $\mathcal{A}_l^{(n+1)}$ are not yet known. Thus, with an abuse of notation, $\mathbf{y}^{(n+1)}$ is redefined as

$$\mathbf{y}^{(n+1)} \triangleq [\mathbf{y}_{l,k}^{(n+1)}]_{l \in \mathcal{C}, k \in \mathcal{A}_l^{(n)}}. \quad (15)$$

In the proposed control approach, each CA $l \in \mathcal{C}$ determines its next control variable $\mathbf{u}_l^{(n+1)}$ such that the information about the next joint state $\mathbf{x}^{(n+1)}$ given $\mathbf{y}^{(1:n+1)}$ is maximized. We quantify this information by the negative conditional differential entropy [32, Chap. 8] of $\mathbf{x}^{(n+1)}$ given $\mathbf{y}^{(n+1)}$, with $\mathbf{y}^{(1:n)}$ being an additional condition that has been observed previously and is thus fixed:

$$\begin{aligned} & -h(\mathbf{x}^{(n+1)} | \mathbf{y}^{(n+1)}; \mathbf{y}^{(1:n)}, \mathbf{u}^{(1:n+1)}) \\ &= \iint f(\mathbf{x}^{(n+1)}, \mathbf{y}^{(n+1)} | \mathbf{y}^{(1:n)}; \mathbf{u}^{(1:n+1)}) \\ & \quad \times \log f(\mathbf{x}^{(n+1)} | \mathbf{y}^{(n+1)}, \mathbf{y}^{(1:n)}; \mathbf{u}^{(1:n+1)}) \\ & \quad \times d\mathbf{x}^{(n+1)} d\mathbf{y}^{(n+1)}, \end{aligned} \quad (16)$$

where \log denotes the natural logarithm. Note that $h(\mathbf{x}^{(n+1)} | \mathbf{y}^{(n+1)}; \mathbf{y}^{(1:n)}, \mathbf{u}^{(1:n+1)})$ depends on the *random vectors* $\mathbf{x}^{(n+1)}$ and $\mathbf{y}^{(n+1)}$, i.e., on their joint distribution but not on their values. Our notation indicates this fact by using a sans serif font for $\mathbf{x}^{(n+1)}$ and $\mathbf{y}^{(n+1)}$ in $h(\mathbf{x}^{(n+1)} | \mathbf{y}^{(n+1)}; \mathbf{y}^{(1:n)}, \mathbf{u}^{(1:n+1)})$.

According to expression (16), $-h(\mathbf{x}^{(n+1)} | \mathbf{y}^{(n+1)}; \mathbf{y}^{(1:n)}, \mathbf{u}^{(1:n+1)})$ is a function of the control vector $\mathbf{u}^{(n+1)}$. This function will be denoted as $D_h(\mathbf{u}^{(n+1)})$, i.e.,

$$D_h(\mathbf{u}^{(n+1)}) \triangleq -h(\mathbf{x}^{(n+1)} | \mathbf{y}^{(n+1)}; \mathbf{y}^{(1:n)}, \mathbf{u}^{(1:n+1)}), \quad (17)$$

and it will be used as the objective function for control at each CA. This objective function is holistic in that it involves *all* the next states (of both the CAs and the targets), $\mathbf{x}^{(n+1)}$, and *all* the next measurements, $\mathbf{y}^{(n+1)}$. Instead of a full-blown maximization of $D_h(\mathbf{u}^{(n+1)})$, we perform one step of a gradient ascent [42] at each time n . Thus, $\mathbf{u}^{(n+1)}$ is determined as

$$\hat{\mathbf{u}}^{(n+1)} = \mathbf{u}_r^{(n+1)} + c^{(n+1)} \nabla D_h(\mathbf{u}^{(n+1)}) \Big|_{\mathbf{u}^{(n+1)} = \mathbf{u}_r^{(n+1)}}, \quad (18)$$

where $\mathbf{u}_r^{(n+1)}$ is a reference vector and $c^{(n+1)} > 0$ is a step size. The choice of $\mathbf{u}_r^{(n+1)}$ depends on the manner in which the local control vectors $\mathbf{u}_l^{(n)}$ (which are subvectors of $\mathbf{u}^{(n)}$) appear in the state evolution functions $g_l(\mathbf{x}_l^{(n-1)}, \mathbf{u}_l^{(n)}, \mathbf{q}_l^{(n)})$ in (1); two common choices are $\mathbf{u}_r^{(n+1)} = \mathbf{u}^{(n)}$ and $\mathbf{u}_r^{(n+1)} = \mathbf{0}$ (cf. Section IX-A).

Since $\mathbf{u}^{(n+1)} = [\mathbf{u}_l^{(n+1)}]_{l \in \mathcal{C}}$, we have

$$\nabla D_h(\mathbf{u}^{(n+1)}) = \left[\frac{\partial D_h(\mathbf{u}^{(n+1)})}{\partial \mathbf{u}_l^{(n+1)}} \right]_{l \in \mathcal{C}},$$

and thus the gradient ascent (18) with respect to $\mathbf{u}^{(n+1)}$ is equivalent to local gradient ascents at the individual CAs l with respect to the local control vectors $\mathbf{u}_l^{(n+1)}$. At CA l , the local gradient ascent is performed as

$$\hat{\mathbf{u}}_l^{(n+1)} = \mathbf{u}_{r,l}^{(n+1)} + c_l^{(n+1)} \frac{\partial D_h(\mathbf{u}^{(n+1)})}{\partial \mathbf{u}_l^{(n+1)}} \Big|_{\mathbf{u}^{(n+1)} = \mathbf{u}_r^{(n+1)}}, \quad (19)$$

where $\mathbf{u}_{r,l}^{(n+1)}$ is the part of $\mathbf{u}_r^{(n+1)}$ that corresponds to CA l (we have $\mathbf{u}_r^{(n+1)} = [\mathbf{u}_{r,l}^{(n+1)}]_{l \in \mathcal{C}}$). The local step sizes $c_l^{(n+1)}$ are constrained by the condition $\mathbf{u}_l^{(n+1)} \in \mathcal{U}_l$ for given sets \mathcal{U}_l . In practice, this condition can be easily satisfied by an appropriate scaling of the $c_l^{(n+1)}$. Note that, as in [15], we use different local step sizes $c_l^{(n+1)}$ at the individual CAs l . This heuristic modification is made to account for the possibly different sets \mathcal{U}_l and to avoid the necessity of reaching a consensus on a common step size across all the CAs; it was observed to yield good results. Because the objective function $D_h(\cdot)$ changes over time n , the local ascent described by (19) generally is not guaranteed to converge; this is similar to existing information-seeking control algorithms [13], [15]. Indeed, the goal of the proposed control algorithm is to make available informative measurements to the estimation layer; because of the dynamic scenario, this is generally not compatible with convergence.

B. Expansion of the Objective Function

A central contribution of this paper is a distributed sample-based technique for calculating the gradients $\frac{\partial D_h(\mathbf{u}^{(n+1)})}{\partial \mathbf{u}_l^{(n+1)}} \Big|_{\mathbf{u}^{(n+1)} = \mathbf{u}_r^{(n+1)}}$ in (19). As a starting point for developing this technique, we next derive an expansion of the objective function $D_h(\mathbf{u}^{(n+1)})$. We will use the following simplified notation. We do not indicate the conditioning on $\mathbf{y}^{(1:n)}$ and the dependence on $\mathbf{u}^{(1:n)}$ because at time $n+1$, $\mathbf{y}^{(1:n)}$ has already been observed and $\mathbf{u}^{(1:n)}$ has already been determined; hence both are fixed. (Note that in Section III, we did not indicate the conditioning on $\mathbf{y}^{(1:n-1)}$, rather than $\mathbf{y}^{(1:n)}$.) Also, we suppress the time index n and designate variables at time $n+1$ by the superscript “+”. For example, we write $h(\mathbf{x}^+ | \mathbf{y}^+; \mathbf{u}^+)$ instead of $h(\mathbf{x}^{(n+1)} | \mathbf{y}^{(n+1)}; \mathbf{y}^{(1:n)}, \mathbf{u}^{(1:n+1)})$.

For calculating the gradient, following [13] and [15], we disregard the unknown driving noise \mathbf{q}_l in (1) by formally replacing it with its expectation, $\bar{\mathbf{q}}_l \triangleq \int \mathbf{q}_l f(\mathbf{q}_l) d\mathbf{q}_l$. We can then rewrite (1) (with n replaced by $n+1$) as

$$\mathbf{x}_l^+ = g_l(\mathbf{x}_l, \mathbf{u}_l^+, \bar{\mathbf{q}}_l^+) = \tilde{g}_l(\mathbf{x}_l, \mathbf{u}_l^+), \quad l \in \mathcal{C}. \quad (20)$$

As shown in Appendix A, the conditional differential entropy in (17) can be expressed as

$$h(\mathbf{x}^+ | \mathbf{y}^+; \mathbf{u}^+) = h(\mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+ | \mathbf{y}^+; \mathbf{u}^+) + \sum_{l \in \mathcal{C}} G_l(\mathbf{u}_l^+), \quad (21)$$

where $\mathbf{x}_\mathcal{C} \triangleq [\mathbf{x}_l]_{l \in \mathcal{C}}$, $\mathbf{x}_\mathcal{T}^+ \triangleq [\mathbf{x}_m^+]_{m \in \mathcal{T}}$, and

$$G_l(\mathbf{u}_l^+) \triangleq \int f(\mathbf{x}_l) \log |J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}_l^+)| d\mathbf{x}_l, \\ \text{with } J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}_l^+) \triangleq \det \frac{\partial \tilde{g}_l(\mathbf{x}_l, \mathbf{u}_l^+)}{\partial \mathbf{x}_l}. \quad (22)$$

The first term on the right-hand side of (21) can be decomposed as [32, Chap. 8]

$$h(\mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+ | \mathbf{y}^+; \mathbf{u}^+) = h(\mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+) - I(\mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+; \mathbf{y}^+; \mathbf{u}^+). \quad (23)$$

Here, $I(\mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+; \mathbf{y}^+; \mathbf{u}^+)$ denotes the two-variable mutual information between $(\mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+)$ and \mathbf{y}^+ (with \mathbf{u}^+ being a deterministic parameter, i.e., not a third random variable), which is given by [32, Chap. 8]

$$I(\mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+; \mathbf{y}^+; \mathbf{u}^+) \\ = \iiint f(\mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+, \mathbf{y}^+; \mathbf{u}^+) \log \frac{f(\mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+, \mathbf{y}^+; \mathbf{u}^+)}{f(\mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+)f(\mathbf{y}^+; \mathbf{u}^+)} \\ \times d\mathbf{x}_\mathcal{C} d\mathbf{x}_\mathcal{T}^+ d\mathbf{y}^+. \quad (24)$$

Note that $h(\mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+)$ in (23) does not depend on \mathbf{u}^+ , since neither the CA states $\mathbf{x}_\mathcal{C}$ nor the future target states $\mathbf{x}_\mathcal{T}^+$ are controlled by the future control variable \mathbf{u}^+ . We explicitly express the dependence of $I(\mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+; \mathbf{y}^+; \mathbf{u}^+)$ on \mathbf{u}^+ by defining the function

$$D_I(\mathbf{u}^+) \triangleq I(\mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+; \mathbf{y}^+; \mathbf{u}^+). \quad (25)$$

Combining (17), (21), (23), and (25) then yields the following expansion of the objective function:

$$D_h(\mathbf{u}^+) = -h(\mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+) + D_I(\mathbf{u}^+) - \sum_{l \in \mathcal{C}} G_l(\mathbf{u}_l^+). \quad (26)$$

This entails the following expansion of the gradient in (19):

$$\frac{\partial D_h(\mathbf{u}^+)}{\partial \mathbf{u}_l^+} = \frac{\partial D_I(\mathbf{u}^+)}{\partial \mathbf{u}_l^+} - \frac{\partial G_l(\mathbf{u}_l^+)}{\partial \mathbf{u}_l^+}. \quad (27)$$

In Sections V–VII, we will develop sample-based techniques for calculating $\frac{\partial D_I(\mathbf{u}^+)}{\partial \mathbf{u}_l^+} \Big|_{\mathbf{u}^+ = \mathbf{u}_l^+}$ and $\frac{\partial G_l(\mathbf{u}_l^+)}{\partial \mathbf{u}_l^+} \Big|_{\mathbf{u}_l^+ = \mathbf{u}_l^+}$. The calculation of $\frac{\partial D_I(\mathbf{u}^+)}{\partial \mathbf{u}_l^+} \Big|_{\mathbf{u}^+ = \mathbf{u}_l^+}$ is cooperative and distributed; it requires communication with neighbor CAs $l' \in \mathcal{C}_l$. The calculation of $\frac{\partial G_l(\mathbf{u}_l^+)}{\partial \mathbf{u}_l^+} \Big|_{\mathbf{u}_l^+ = \mathbf{u}_l^+}$ is performed locally at each CA l . Both calculations use the samples of relevant marginal posteriors that were computed by the estimation layer. The operations performed in the control layer as described in this section and in Sections V–VII are summarized in Figs. 3 and 4 for two alternative distributed implementations (discussed in Section VI).

V. CALCULATION OF THE GRADIENT OF $D_I(\mathbf{u}^+)$

In this section, we develop a Monte Carlo approximation of $\frac{\partial D_I(\mathbf{u}^+)}{\partial \mathbf{u}_l^+} \Big|_{\mathbf{u}^+ = \mathbf{u}_l^+}$ that uses importance sampling. The distributed computation of this approximation will be addressed in Section VI.

The mutual information in (24) can be rewritten as

$$D_I(\mathbf{u}^+) = \iiint f(\mathbf{y}^+ | \mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+; \mathbf{u}^+) f(\mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+) \\ \times \log \frac{f(\mathbf{y}^+ | \mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+; \mathbf{u}^+)}{f(\mathbf{y}^+; \mathbf{u}^+)} d\mathbf{x}_\mathcal{C} d\mathbf{x}_\mathcal{T}^+ d\mathbf{y}^+.$$

Invoking [15, Th. 1], we obtain

$$\frac{\partial D_I(\mathbf{u}^+)}{\partial \mathbf{u}_l^+} = \iiint \frac{\partial f(\mathbf{y}^+ | \mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+; \mathbf{u}^+)}{\partial \mathbf{u}_l^+} f(\mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+) \\ \times \log \frac{f(\mathbf{y}^+ | \mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+; \mathbf{u}^+)}{f(\mathbf{y}^+; \mathbf{u}^+)} d\mathbf{x}_\mathcal{C} d\mathbf{x}_\mathcal{T}^+ d\mathbf{y}^+. \quad (28)$$

The likelihood function $f(\mathbf{y}^+ | \mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+; \mathbf{u}^+)$ involved in (28) can be written as

$$f(\mathbf{y}^+ | \mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+; \mathbf{u}^+) = \prod_{l \in \mathcal{C}} \prod_{l' \in \mathcal{C}_l} f(\mathbf{y}_{l,l'}^+ | \mathbf{x}_l, \mathbf{x}_{l'}; \mathbf{u}_l^+, \mathbf{u}_{l'}^+) \\ \times \prod_{m \in \mathcal{T}_l} f(\mathbf{y}_{l,m}^+ | \mathbf{x}_l, \mathbf{x}_m^+; \mathbf{u}_l^+), \quad (29)$$

with

$$f(\mathbf{y}_{l,l'}^+ | \mathbf{x}_l, \mathbf{x}_{l'}; \mathbf{u}_l^+, \mathbf{u}_{l'}^+) \\ = f(\mathbf{y}_{l,l'}^+ | \mathbf{x}_l^+, \mathbf{x}_{l'}^+) \Big|_{\mathbf{x}_l^+ = \tilde{g}_l(\mathbf{x}_l, \mathbf{u}_l^+), \mathbf{x}_{l'}^+ = \tilde{g}_{l'}(\mathbf{x}_{l'}, \mathbf{u}_{l'}^+)}, \quad l \in \mathcal{C}, l' \in \mathcal{C}_l \\ f(\mathbf{y}_{l,m}^+ | \mathbf{x}_l, \mathbf{x}_m^+; \mathbf{u}_l^+) \\ = f(\mathbf{y}_{l,m}^+ | \mathbf{x}_l^+, \mathbf{x}_m^+) \Big|_{\mathbf{x}_l^+ = \tilde{g}_l(\mathbf{x}_l, \mathbf{u}_l^+)}, \quad l \in \mathcal{C}, m \in \mathcal{T}_l.$$

(The latter expressions are obtained using (3) and (20).)

Let $\tilde{\mathbf{y}}_l^+$ denote the subvector of $\mathbf{y}^+ = [\mathbf{y}_{l,k}^+]_{l \in \mathcal{C}, k \in \mathcal{A}_l}$ (cf. (15)) whose likelihood function includes all those factors of $f(\mathbf{y}^+ | \mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+; \mathbf{u}^+)$ in (29) that depend on the local control vector \mathbf{u}_l^+ . This subvector is given by

$$\tilde{\mathbf{y}}_l^+ \triangleq \left[[\mathbf{y}_{l,k}^+]_{k \in \mathcal{A}_l} [\mathbf{y}_{l',l}^+]_{l' \in \mathcal{C}_l} \right]^T, \quad (31)$$

and its likelihood function is obtained as

$$f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_l, \mathbf{x}_{\mathcal{C}_l}, \mathbf{x}_{\mathcal{T}_l}^+; \mathbf{u}_{\mathcal{C}_l}^+) \\ = \prod_{l' \in \mathcal{C}_l} f(\mathbf{y}_{l,l'}^+ | \mathbf{x}_l, \mathbf{x}_{l'}; \mathbf{u}_l^+, \mathbf{u}_{l'}^+) f(\mathbf{y}_{l',l}^+ | \mathbf{x}_{l'}, \mathbf{x}_l; \mathbf{u}_{l'}^+, \mathbf{u}_l^+) \\ \times \prod_{m \in \mathcal{T}_l} f(\mathbf{y}_{l,m}^+ | \mathbf{x}_l, \mathbf{x}_m^+; \mathbf{u}_l^+), \quad (32)$$

with $\mathbf{u}_{\mathcal{C}_l}^+ \triangleq [\mathbf{u}_{l'}^+]_{l' \in \{l\} \cup \mathcal{C}_l}$. By comparing (32) with (29), it is seen that $\tilde{\mathbf{y}}_l^+$ is also the subvector of \mathbf{y}^+ whose likelihood function includes all those factors of $f(\mathbf{y}^+ | \mathbf{x}_\mathcal{C}, \mathbf{x}_\mathcal{T}^+; \mathbf{u}^+)$ in (29) that involve the state \mathbf{x}_l^+ .

Using (29) and (32), it is shown in Appendix B that a Monte Carlo (i.e., sample-based) approximation of (28) evaluated at $\mathbf{u}^+ = \mathbf{u}_l^+$ is given by²

²With an abuse of notation, the superscript (j) now indicates the j th sample, whereas previously, in our full-blown notation, the superscript (n) indicated the n th time step.

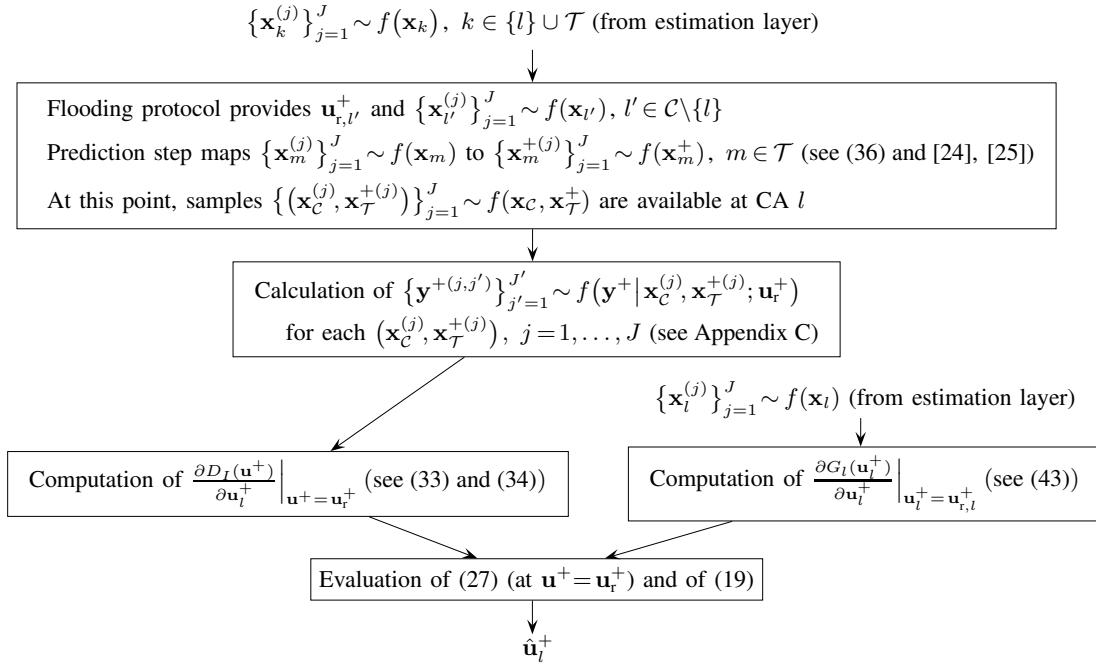


Fig. 3. Flow chart of the flooding-based implementation of the control layer at CA l (see Section VI-A).

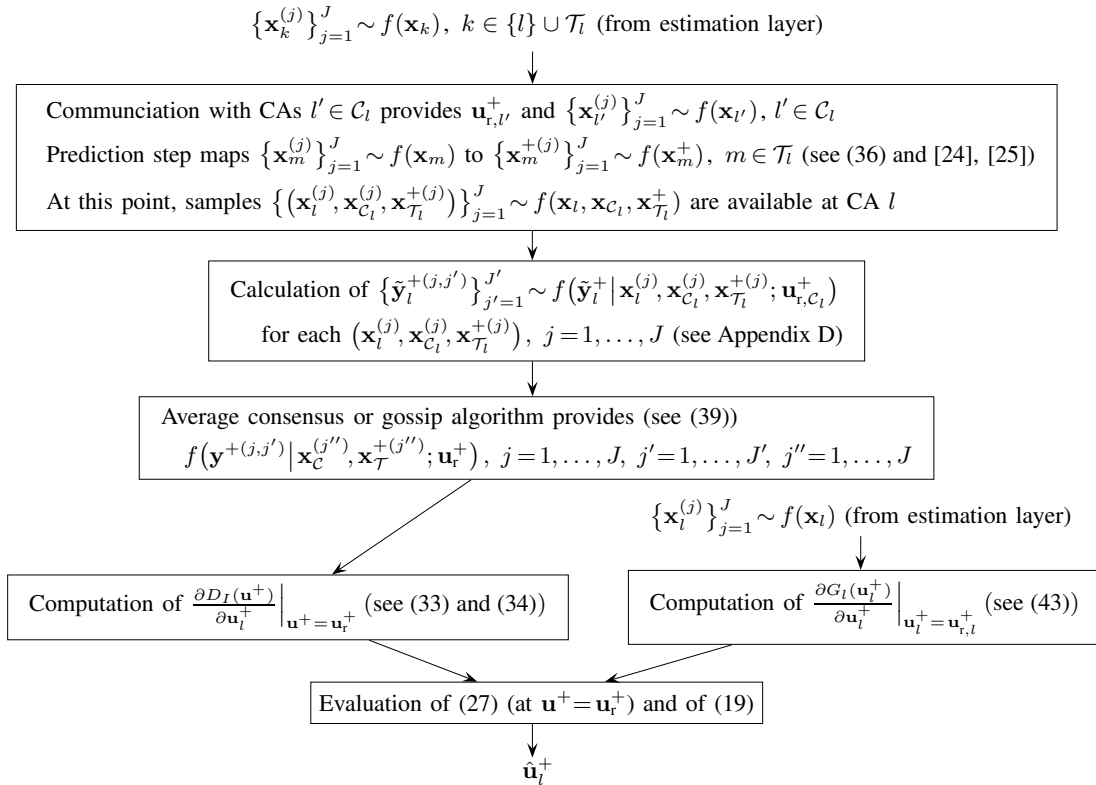


Fig. 4. Flow chart of the consensus-based implementation of the control layer at CA l (see Section VI-B).

$$\begin{aligned}
& \left. \frac{\partial D_I(\mathbf{u}^+)}{\partial \mathbf{u}_l^+} \right|_{\mathbf{u}^+ = \mathbf{u}_r^+} && \times \left. \frac{\partial f(\tilde{\mathbf{y}}_l^{+(j,j')} | \mathbf{x}_l^{(j)}, \mathbf{x}_{c_l}^{(j)}, \mathbf{x}_{T_l}^{+(j)}; \mathbf{u}_{c_l}^+)}{\partial \mathbf{u}_l^+} \right|_{\mathbf{u}_{c_l}^+ = \mathbf{u}_{r,c_l}^+} \\
& \approx \frac{1}{JJ'} \sum_{j=1}^J \sum_{j'=1}^{J'} \frac{1}{f(\tilde{\mathbf{y}}_l^{+(j,j')} | \mathbf{x}_l^{(j)}, \mathbf{x}_{c_l}^{(j)}, \mathbf{x}_{T_l}^{+(j)}; \mathbf{u}_{r,c_l}^+)} && \times \log \frac{f(\mathbf{y}^{+(j,j')} | \mathbf{x}_c^{(j)}, \mathbf{x}_T^{+(j)}; \mathbf{u}_r^+)}{f(\mathbf{y}^{+(j,j')} | \mathbf{u}_r^+)}, \quad (33)
\end{aligned}$$

with

$$f(\mathbf{y}^{+(j,j')}; \mathbf{u}_r^+) \approx \frac{1}{J} \sum_{j''=1}^J f(\mathbf{y}^{+(j,j')} | \mathbf{x}_C^{(j'')}, \mathbf{x}_T^{+(j'')}; \mathbf{u}_r^+). \quad (34)$$

Here, $\mathbf{y}^{+(j,j')}$, $\mathbf{x}_C^{(j)}$, and $\mathbf{x}_T^{+(j)}$ are samples of \mathbf{y}^+ , \mathbf{x}_C , and \mathbf{x}_T^+ , respectively that are drawn from the importance density [43] $q(\mathbf{y}^+, \mathbf{x}_C, \mathbf{x}_T^+) \triangleq f(\mathbf{x}_C, \mathbf{x}_T^+) f(\mathbf{y}^+ | \mathbf{x}_C, \mathbf{x}_T^+; \mathbf{u}_r^+)$ via the following two-stage procedure. First, samples $\{(\mathbf{x}_C^{(j)}, \mathbf{x}_T^{+(j)})\}_{j=1}^J$ are drawn from

$$f(\mathbf{x}_C, \mathbf{x}_T^+) = \prod_{l \in \mathcal{C}} f(\mathbf{x}_l) \prod_{m \in \mathcal{T}} f(\mathbf{x}_m^+). \quad (35)$$

(This factorization expresses the conditional statistical independence of the \mathbf{x}_l , $l \in \mathcal{C}$ and the \mathbf{x}_m^+ , $m \in \mathcal{T}$ given $\mathbf{y}^{(1:n)}$. This is a common approximation, which was introduced in [17] and is also used in the estimation layer [24], [25].) Then, for each sample $(\mathbf{x}_C^{(j)}, \mathbf{x}_T^{+(j)})$, samples $\{\mathbf{y}^{+(j,j')}\}_{j'=1}^{J'}$ are drawn from the conditional pdf $f(\mathbf{y}^+ | \mathbf{x}_C^{(j)}, \mathbf{x}_T^{+(j)}; \mathbf{u}_r^+)$. The distributed calculation of these samples will be discussed in Section VI and in Appendix C. Finally, we note that using (32), one easily obtains a (rather unwieldy) expression of the derivative $\frac{\partial f(\mathbf{y}_l^+ | \mathbf{x}_l, \mathbf{x}_{C_l}, \mathbf{x}_{T_l}^+; \mathbf{u}_{C_l}^+)}{\partial \mathbf{u}_l^+}$ occurring in (33). This expression involves the factors in (32) and the derivatives $\frac{\partial g_l(\mathbf{x}_l, \mathbf{u}_l^+)}{\partial \mathbf{u}_l^+}$, $\frac{\partial f(\mathbf{y}_{l,l'}^+ | \mathbf{x}_l^+, \mathbf{x}_{l'}^+)}{\partial \mathbf{x}_l^+}$ for $l' \in \mathcal{C}_l$, and $\frac{\partial f(\mathbf{y}_{l,m}^+ | \mathbf{x}_l^+, \mathbf{x}_m^+)}{\partial \mathbf{x}_l^+}$ for $m \in \mathcal{T}_l$.

VI. DISTRIBUTED PROCESSING

In this section, we present two alternative schemes for a distributed, sample-based computation of $\frac{\partial D_l(\mathbf{u}^+)}{\partial \mathbf{u}_l^+} \Big|_{\mathbf{u}^+ = \mathbf{u}_l^+}$ according to (33) and (34). Both schemes are distributed in that they require only local communication, i.e., each CA $l \in \mathcal{C}$ transmits data only to its neighbors $l' \in \mathcal{C}_l$.

A. Flooding-Based Processing

We first discuss a distributed scheme where each CA $l \in \mathcal{C}$ performs an exact (“quasi-centralized”) calculation of (33) and (34). As a result of the estimation layer, samples $\{\mathbf{x}_k^{(j)}\}_{j=1}^J \sim f(\mathbf{x}_k)$, $k \in \{l\} \cup \mathcal{T}$ are available at CA l (see (10), noting that $f(\mathbf{x}_k)$ was denoted $f(\mathbf{x}_k | \mathbf{y})$ there). A flooding protocol [44] is now used to make available to each CA l the reference vectors $\mathbf{u}_{r,l'}^+$ and the samples $\{\mathbf{x}_{l'}^{(j)}\}_{j=1}^J \sim f(\mathbf{x}_{l'})$ of all the other CAs $l' \in \mathcal{C} \setminus \{l\}$. The flooding protocol requires each CA l to communicate with its neighbor CAs $l' \in \mathcal{C}_l$. In addition, CA l locally calculates predictive marginal posteriors for all target states via the following prediction step (which is (9) with n replaced by $n+1$):

$$f(\mathbf{x}_m^+) = \int f(\mathbf{x}_m^+ | \mathbf{x}_m) f(\mathbf{x}_m) d\mathbf{x}_m, \quad m \in \mathcal{T}. \quad (36)$$

An implementation of (36) using the samples $\{\mathbf{x}_m^{(j)}\}_{j=1}^J$, $m \in \mathcal{T}$ produced by the estimation layer (which are available at CA l) and yielding samples $\{\mathbf{x}_m^{+(j)}\}_{j=1}^J \sim f(\mathbf{x}_m^+)$, $m \in \mathcal{T}$ is described in [24], [25]. At this point, samples $\{\mathbf{x}_{l'}^{(j)}\}_{j=1}^J \sim$

$f(\mathbf{x}_{l'})$ for $l' \in \mathcal{C}$ and $\{\mathbf{x}_m^{+(j)}\}_{j=1}^J \sim f(\mathbf{x}_m^+)$ for $m \in \mathcal{T}$ are available at CA l . Because all states \mathbf{x}_l , $l \in \mathcal{C}$ and \mathbf{x}_m^+ , $m \in \mathcal{T}$ are conditionally independent given $\mathbf{y}^{(1:n)}$ (see (35)), CA l can now obtain samples $\{(\mathbf{x}_C^{(j)}, \mathbf{x}_T^{+(j)})\}_{j=1}^J \sim f(\mathbf{x}_C, \mathbf{x}_T^+)$ by the simple stacking operations $\mathbf{x}_C^{(j)} = [\mathbf{x}_{l'}^{(j)}]_{l' \in \mathcal{C}}$ and $\mathbf{x}_T^{+(j)} = [\mathbf{x}_m^{+(j)}]_{m \in \mathcal{T}}$. Finally, for each $(\mathbf{x}_C^{(j)}, \mathbf{x}_T^{+(j)})$, CA l computes samples $\{\mathbf{y}^{+(j,j')}\}_{j'=1}^{J'} \sim f(\mathbf{y}^+ | \mathbf{x}_C^{(j)}, \mathbf{x}_T^{+(j)}; \mathbf{u}_r^+)$ as described in Appendix C.

Using the samples $(\mathbf{x}_C^{(j)}, \mathbf{x}_T^{+(j)})$ and $\mathbf{y}^{+(j,j')}$, $j = 1, \dots, J$, $j' = 1, \dots, J'$, as well as the reference vectors $\mathbf{u}_{r,l'}^+$, $l' \in \mathcal{C}$, CA l can compute the gradient $\frac{\partial D_l(\mathbf{u}^+)}{\partial \mathbf{u}_l^+} \Big|_{\mathbf{u}^+ = \mathbf{u}_l^+}$ locally according to (33) and (34). Note, however, that this flooding-based scheme presupposes that each CA l knows the state evolution models (1) and the measurement models (3) of all the other CAs $l' \in \mathcal{C} \setminus \{l\}$.

The communication cost of the flooding-based scheme, in terms of the number of real values transmitted by each CA, is $(JM + M_u)W \approx JMW$. Here, M and M_u are the dimensions of the vectors \mathbf{x}_l and \mathbf{u}_l , respectively, and W depends on the network size and topology and is bounded as $1 \leq W \leq |\mathcal{C}|$. Thus, the number of transmissions scales linearly with J and does not depend on J' . In large networks, flooding protocols tend to require a large memory and book-keeping overhead and introduce a significant delay [45]. If the network formed by the CAs is *fully* connected, i.e., $\mathcal{C} = \{l\} \cup \mathcal{C}_l$, then all the samples $\{(\mathbf{x}_C^{(j)}, \mathbf{x}_T^{+(j)})\}_{j=1}^J$ can be obtained without flooding: CA l simply broadcasts its reference vector $\mathbf{u}_{r,l}^+$ and its samples $\{\mathbf{x}_l^{(j)}\}_{j=1}^J \sim f(\mathbf{x}_l)$ to all the other CAs in the network and receives their reference vectors and samples. Here, the number of real values transmitted by each CA is only $JM + M_u$.

Finally, the computational complexity per CA of the flooding-based scheme—i.e., evaluation of (33) and (34), with J and J' fixed—scales linearly with the number of agents in the network. Because the computational complexity and the communication cost increase with the network size, the flooding-based distributed processing scheme is primarily suited to small networks.

B. Consensus-Based Processing

Next, we present an alternative distributed computation of (33) and (34) that avoids the use of a flooding protocol and does not require each CA to know the state evolution and measurement models of all the other CAs. As a first step, CA l broadcasts its own samples $\{\mathbf{x}_l^{(j)}\}_{j=1}^J \sim f(\mathbf{x}_l)$ calculated in the estimation layer to all neighbor CAs $l' \in \mathcal{C}_l$, and it receives from them their own samples $\{\mathbf{x}_{l'}^{(j)}\}_{j=1}^J \sim f(\mathbf{x}_{l'})$, $l' \in \mathcal{C}_l$. In addition, CA l locally calculates samples $\{\mathbf{x}_m^{+(j)}\}_{j=1}^J \sim f(\mathbf{x}_m^+)$ for all $m \in \mathcal{T}_l$ via the prediction step (36) (with \mathcal{T} replaced by \mathcal{T}_l), using the sample-based implementation described in [24], [25]. Thus, after the stacking operations $\mathbf{x}_{C_l}^{(j)} = [\mathbf{x}_{l'}^{(j)}]_{l' \in \mathcal{C}_l}$ and $\mathbf{x}_{T_l}^{+(j)} = [\mathbf{x}_m^{+(j)}]_{m \in \mathcal{T}_l}$, samples $\{(\mathbf{x}_l^{(j)}, \mathbf{x}_{C_l}^{(j)}, \mathbf{x}_{T_l}^{+(j)})\}_{j=1}^J \sim f(\mathbf{x}_l, \mathbf{x}_{C_l}, \mathbf{x}_{T_l}^+)$ are available at CA l . Then, for each sample $(\mathbf{x}_l^{(j)}, \mathbf{x}_{C_l}^{(j)}, \mathbf{x}_{T_l}^{+(j)})$, J' samples $\{\tilde{\mathbf{y}}_l^{+(j,j')}\}_{j'=1}^{J'} \sim f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_l^{(j)}, \mathbf{x}_{C_l}^{(j)}, \mathbf{x}_{T_l}^{+(j)}; \mathbf{u}_{r,C_l}^+)$ are

computed as described in Appendix D. This only involves communication between neighboring CAs.

The key question at this point is as to whether the quantities $f(\tilde{\mathbf{y}}_l^{+(j,j')} | \mathbf{x}_l^{(j)}, \mathbf{x}_{\mathcal{C}_l}^{(j)}, \mathbf{x}_{\mathcal{T}_l}^{+(j)}; \mathbf{u}_{r,\mathcal{C}_l}^+)$, $\frac{\partial f(\tilde{\mathbf{y}}_l^{+(j,j')} | \mathbf{x}_l^{(j)}, \mathbf{x}_{\mathcal{C}_l}^{(j)}, \mathbf{x}_{\mathcal{T}_l}^{+(j)}; \mathbf{u}_{r,\mathcal{C}_l}^+)}{\partial \mathbf{u}_l^+} \Big|_{\mathbf{u}_{\mathcal{C}_l}^+ = \mathbf{u}_{r,\mathcal{C}_l}^+}$, and $f(\mathbf{y}^{+(j,j')} | \mathbf{x}_{\mathcal{C}}^{(j'')}, \mathbf{x}_{\mathcal{T}}^{+(j'')}; \mathbf{u}_r^+)$ (and, in particular, $f(\mathbf{y}^{+(j,j')} | \mathbf{x}_{\mathcal{C}}^{(j)}, \mathbf{x}_{\mathcal{T}}^{+(j)}; \mathbf{u}_r^+)$) involved in (33) and (34) are locally available at CA l . The factors of $f(\tilde{\mathbf{y}}_l^{+(j,j')} | \mathbf{x}_l^{(j)}, \mathbf{x}_{\mathcal{C}_l}^{(j)}, \mathbf{x}_{\mathcal{T}_l}^{+(j)}; \mathbf{u}_{r,\mathcal{C}_l}^+)$ (see (32)) correspond to measurements to be acquired by CA l or by its neighbor CAs $l' \in \mathcal{C}_l$; they are known to CA l since its own state evolution and measurement models and those of its neighbors are known to CA l (cf. (30)). Thus, we conclude that the $f(\tilde{\mathbf{y}}_l^{+(j,j')} | \mathbf{x}_l^{(j)}, \mathbf{x}_{\mathcal{C}_l}^{(j)}, \mathbf{x}_{\mathcal{T}_l}^{+(j)}; \mathbf{u}_{r,\mathcal{C}_l}^+)$ are available at CA l . On the other hand, many of the factors of $f(\mathbf{y}^+ | \mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+; \mathbf{u}_r^+)$ (see (29)) correspond to measurements to be acquired by CAs that are not in the neighborhood of CA l ; they are not known to CA l since, typically, the respective state evolution and measurement models are unknown to CA l . Therefore, the $f(\mathbf{y}^{+(j,j')} | \mathbf{x}_{\mathcal{C}}^{(j'')}, \mathbf{x}_{\mathcal{T}}^{+(j'')}; \mathbf{u}_r^+)$ are not available at CA l .

We will now present a distributed computation of $f(\mathbf{y}^{+(j,j')} | \mathbf{x}_{\mathcal{C}}^{(j'')}, \mathbf{x}_{\mathcal{T}}^{+(j'')}; \mathbf{u}_r^+)$. Let \mathbf{y}_l^+ denote the subvector of $\mathbf{y}^+ = [\mathbf{y}_{l,k}^+]_{l \in \mathcal{C}, k \in \mathcal{A}_l}$ in (15) that comprises the measurements acquired by CA l at the next time, i.e.,

$$\mathbf{y}_l^+ \triangleq [\mathbf{y}_{l,k}^+]_{k \in \mathcal{A}_l}. \quad (37)$$

The likelihood function of \mathbf{y}_l^+ combines all the factors in (29) that involve the entries of \mathbf{y}_l^+ , i.e.,

$$f(\mathbf{y}_l^+ | \mathbf{x}_l, \mathbf{x}_{\mathcal{C}_l}, \mathbf{x}_{\mathcal{T}_l}^+; \mathbf{u}_{\mathcal{C}_l}^+) = \prod_{l' \in \mathcal{C}_l} f(\mathbf{y}_{l,l'}^+ | \mathbf{x}_l, \mathbf{x}_{l'}; \mathbf{u}_l^+, \mathbf{u}_{l'}^+) \prod_{m \in \mathcal{T}_l} f(\mathbf{y}_{l,m}^+ | \mathbf{x}_l, \mathbf{x}_m; \mathbf{u}_l^+). \quad (38)$$

Using (29) and (38), one can show that

$$f(\mathbf{y}^{+(j,j')} | \mathbf{x}_{\mathcal{C}}^{(j'')}, \mathbf{x}_{\mathcal{T}}^{+(j'')}; \mathbf{u}_r^+) = \exp(|\mathcal{C}| F_{j,j',j''}), \quad (39)$$

where

$$F_{j,j',j''} \triangleq \frac{1}{|\mathcal{C}|} \sum_{l \in \mathcal{C}} F_{j,j',j''}^{(l)} \quad (40)$$

with

$$F_{j,j',j''}^{(l)} \triangleq \log f(\mathbf{y}_l^{+(j,j')} | \mathbf{x}_l^{(j'')}, \mathbf{x}_{\mathcal{C}_l}^{(j'')}, \mathbf{x}_{\mathcal{T}_l}^{+(j'')}; \mathbf{u}_{r,\mathcal{C}_l}^+), \quad (41)$$

for $j = 1, \dots, J$, $j' = 1, \dots, J'$, and $j'' = 1, \dots, J$. To compute $F_{j,j',j''}^{(l)}$ in (41), CA l needs samples $\{\mathbf{y}_l^{+(j,j')}\}_{j'=1}^{J'}$ $\sim f(\mathbf{y}_l^+ | \mathbf{x}_l^{(j)}, \mathbf{x}_{\mathcal{C}_l}^{(j)}, \mathbf{x}_{\mathcal{T}_l}^{+(j)}; \mathbf{u}_{r,\mathcal{C}_l}^+)$ and the reference vectors $\mathbf{u}_{r,l'}^+$ for $l' \in \{l\} \cup \mathcal{C}_l$. The samples $\{\mathbf{y}_l^{+(j,j')}\}_{j'=1}^{J'}$ are already available at CA l since \mathbf{y}_l^+ is a subvector of $\tilde{\mathbf{y}}_l^+$ (see (31) and (37)) and samples $\{\tilde{\mathbf{y}}_l^{+(j,j')}\}_{j'=1}^{J'}$ $\sim f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_l^{(j)}, \mathbf{x}_{\mathcal{C}_l}^{(j)}, \mathbf{x}_{\mathcal{T}_l}^{+(j)}; \mathbf{u}_{r,\mathcal{C}_l}^+)$ have already been computed as described above. The $\mathbf{u}_{r,l'}^+$ can be obtained at CA l through communication with the neighbor CAs $l' \in \mathcal{C}_l$.

Once the $F_{j,j',j''}^{(l)}$ have been calculated at CA l , their averages $F_{j,j',j''}$ in (40) can be computed in a distributed way by using $J^2 J'$ parallel instances of an average consensus or gossip scheme [26], [27]. These schemes are iterative;

they are initialized at each CA l with $F_{j,j',j''}^{(l)}$. They are robust to communication link failures [26], [27] and use only communication between neighbor CAs (i.e., each CA $l \in \mathcal{C}$ transmits data to each neighbor $l' \in \mathcal{C}_l$). After convergence of the consensus or gossip scheme, $F_{j,j',j''}$ and, hence, $f(\mathbf{y}^{+(j,j')} | \mathbf{x}_{\mathcal{C}}^{(j'')}, \mathbf{x}_{\mathcal{T}}^{+(j'')}; \mathbf{u}_r^+)$ for all j, j', j'' is available at each CA l .

At this point, CA l has available $f(\tilde{\mathbf{y}}_l^{+(j,j')} | \mathbf{x}_l^{(j)}, \mathbf{x}_{\mathcal{C}_l}^{(j)}, \mathbf{x}_{\mathcal{T}_l}^{+(j)}; \mathbf{u}_{r,\mathcal{C}_l}^+)$ and $\frac{\partial f(\tilde{\mathbf{y}}_l^{+(j,j')} | \mathbf{x}_l^{(j)}, \mathbf{x}_{\mathcal{C}_l}^{(j)}, \mathbf{x}_{\mathcal{T}_l}^{+(j)}; \mathbf{u}_{r,\mathcal{C}_l}^+)}{\partial \mathbf{u}_l^+} \Big|_{\mathbf{u}_{\mathcal{C}_l}^+ = \mathbf{u}_{r,\mathcal{C}_l}^+}$, and an approximation of $f(\mathbf{y}^{+(j,j')} | \mathbf{x}_{\mathcal{C}}^{(j'')}, \mathbf{x}_{\mathcal{T}}^{+(j'')}; \mathbf{u}_r^+)$ has been provided by the consensus or gossip scheme, for $j = 1, \dots, J$, $j' = 1, \dots, J'$, and $j'' = 1, \dots, J$. Therefore, CA l is now able to evaluate (33) and (34).

In the course of the overall distributed computation, CA l transmits $J^2 J' |\mathcal{C}_l| R + J J' |\mathcal{C}_l| M_y + J M + M_u \approx J^2 J' |\mathcal{C}_l| R$ real values, where R is the number of iterations used for one instance of the consensus or gossip scheme and M_y is the dimension of the vectors $\mathbf{y}_{l,k}$. Asymptotically, for $R \rightarrow \infty$, this distributed computation of $\frac{\partial D_l(\mathbf{u}_l^+)}{\partial \mathbf{u}_l^+} \Big|_{\mathbf{u}_l^+ = \mathbf{u}_r^+}$ converges to the exact centralized result given by (33) and (34). The speed of convergence depends on the topology and size of the network [26], [27]. As R increases, the information available at each agent converges, which means that local data is disseminated over large distances in the network. However, because the control vector of a given CA might not be strongly affected by information from far away CAs, a small R might be sufficient for good performance. Because the communication requirements are proportional to $J^2 J'$, they are typically higher than those of the flooding-based scheme discussed in Section VI-A unless the network is large and R is small.

Finally, the computational complexity of the distributed processing—i.e., evaluation of (33) and (34), with J , J' , and R fixed—is constant in the number of agents in the network.

VII. CALCULATION OF THE GRADIENT OF $G_l(\mathbf{u}_l^+)$

Next, we consider the second gradient in the expansion (27), i.e., $\frac{\partial G_l(\mathbf{u}_l^+)}{\partial \mathbf{u}_l^+} \Big|_{\mathbf{u}_l^+ = \mathbf{u}_{r,l}^+}$. Using (22), we obtain

$$\begin{aligned} & \frac{\partial G_l(\mathbf{u}_l^+)}{\partial \mathbf{u}_l^+} \Big|_{\mathbf{u}_l^+ = \mathbf{u}_{r,l}^+} \\ &= \int f(\mathbf{x}_l) \frac{\partial \log |J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}_l^+)|}{\partial \mathbf{u}_l^+} \Big|_{\mathbf{u}_l^+ = \mathbf{u}_{r,l}^+} d\mathbf{x}_l \\ &= \int f(\mathbf{x}_l) \frac{1}{|J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}_l^+)|} \frac{\partial |J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}_l^+)|}{\partial \mathbf{u}_l^+} \Big|_{\mathbf{u}_l^+ = \mathbf{u}_{r,l}^+} d\mathbf{x}_l. \quad (42) \end{aligned}$$

Here, we assumed that $|J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}_l^+)|$ is continuous and satisfies $|f(\mathbf{x}_l) \partial \log |J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}_l^+)| / \partial \mathbf{u}_l^+| \leq \alpha(\mathbf{x}_l, \mathbf{u}_l^+)$ for all $(\mathbf{x}_l, \mathbf{u}_l^+)$, for some function $\alpha(\mathbf{x}_l, \mathbf{u}_l^+) \geq 0$ that is integrable with respect to \mathbf{x}_l for each \mathbf{u}_l^+ [46, Cor. 5.9]. Furthermore, we assumed that for each value of \mathbf{x}_l , $|J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}_l^+)|$ is differentiable with respect to \mathbf{u}_l^+ at $\mathbf{u}_{r,l}^+$. A sufficient condition is that $J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}_l^+)$ is differentiable with respect to \mathbf{u}_l^+ at $\mathbf{u}_{r,l}^+$ and nonzero for all \mathbf{u}_l^+ in some (arbitrarily small) neighborhood of $\mathbf{u}_{r,l}^+$.

Based on the samples $\{\mathbf{x}_l^{(j)}\}_{j=1}^J \sim f(\mathbf{x}_l)$ that were calculated in the estimation layer, a Monte Carlo approximation of (42) is obtained as

$$\begin{aligned} & \left. \frac{\partial G_l(\mathbf{u}_l^+)}{\partial \mathbf{u}_l^+} \right|_{\mathbf{u}_l^+ = \mathbf{u}_{r,l}^+} \\ & \approx \frac{1}{J} \sum_{j=1}^J \frac{1}{|J_{\tilde{g}_l}(\mathbf{x}_l^{(j)}; \mathbf{u}_{r,l}^+)|} \left. \frac{\partial |J_{\tilde{g}_l}(\mathbf{x}_l^{(j)}; \mathbf{u}_l^+)|}{\partial \mathbf{u}_l^+} \right|_{\mathbf{u}_l^+ = \mathbf{u}_{r,l}^+}. \end{aligned} \quad (43)$$

For many practically relevant state evolution models (20), the computation of $\left. \frac{\partial G_l(\mathbf{u}_l^+)}{\partial \mathbf{u}_l^+} \right|_{\mathbf{u}_l^+ = \mathbf{u}_{r,l}^+}$ can be avoided altogether or $\left. \frac{\partial G_l(\mathbf{u}_l^+)}{\partial \mathbf{u}_l^+} \right|_{\mathbf{u}_l^+ = \mathbf{u}_{r,l}^+}$ can be calculated in closed form, without a sample-based approximation. Some examples are considered in the following.

- 1) $J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}_l^+)$ does not depend on \mathbf{u}_l^+ : In this case, $\left. \frac{\partial G_l(\mathbf{u}_l^+)}{\partial \mathbf{u}_l^+} \right|_{\mathbf{u}_l^+ = \mathbf{u}_{r,l}^+} = \mathbf{0}$. An important example is the ‘‘linear additive’’ state evolution model $\tilde{g}_l(\mathbf{x}_l, \mathbf{u}_l^+) = \mathbf{A}\mathbf{x}_l + \zeta(\mathbf{u}_l^+)$ with some matrix \mathbf{A} and function $\zeta(\cdot)$ of suitable dimensions. Here, we obtain $J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}_l^+) = \det \mathbf{A}$ and thus $\left. \frac{\partial G_l(\mathbf{u}_l^+)}{\partial \mathbf{u}_l^+} \right|_{\mathbf{u}_l^+ = \mathbf{u}_{r,l}^+} = \mathbf{0}$. A second important example is the *odometry motion model* [47, Sec. 5.3]. Here, the local state \mathbf{x}_l is the pose of a robot, which consists of the 2D position $(x_{l,1}, x_{l,2})$ and the orientation θ_l , and the control vector \mathbf{u}_l consists of the translational velocity ν_l and the rotational velocity ω_l . The state evolution model is given by

$$\tilde{g}_l(\mathbf{x}_l, \mathbf{u}_l^+) = \begin{bmatrix} x_{l,1} + \nu_l^+ \cos(\theta_l + \omega_l^+) \\ x_{l,2} + \nu_l^+ \sin(\theta_l + \omega_l^+) \\ \theta_l + \omega_l^+ \end{bmatrix}.$$

Here, $J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}_l^+) = 1$ and thus $\left. \frac{\partial G_l(\mathbf{u}_l^+)}{\partial \mathbf{u}_l^+} \right|_{\mathbf{u}_l^+ = \mathbf{u}_{r,l}^+} = \mathbf{0}$.

- 2) $J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}_l^+)$ does not depend on \mathbf{x}_l : If $J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}_l^+) = J_{\tilde{g}_l}(\mathbf{u}_l^+)$, then (22) simplifies to $G_l(\mathbf{u}_l^+) = \log |J_{\tilde{g}_l}(\mathbf{u}_l^+)|$. Thus, we have

$$\left. \frac{\partial G_l(\mathbf{u}_l^+)}{\partial \mathbf{u}_l^+} \right|_{\mathbf{u}_l^+ = \mathbf{u}_{r,l}^+} = \frac{1}{|J_{\tilde{g}_l}(\mathbf{u}_l^+)|} \left. \frac{\partial |J_{\tilde{g}_l}(\mathbf{u}_l^+)|}{\partial \mathbf{u}_l^+} \right|_{\mathbf{u}_l^+ = \mathbf{u}_{r,l}^+},$$

which can be calculated in closed form.

VIII. TWO SPECIAL CASES

A. Cooperative Estimation of Local States

Here, we assume that there are no targets, and thus the task considered is only the distributed, cooperative estimation of the local states.

1) *Estimation Layer*: The marginal posteriors corresponding to the targets are no longer calculated. In the calculation of the marginal posterior of CA l , the correction step (10) simplifies to

$$f(\mathbf{x}_l | \mathbf{y}) \propto \prod_{l' \in \mathcal{C}_l} f(\mathbf{x}_{l'}) \prod_{l_1 \in \mathcal{C}_{l'}} f(\mathbf{y}_{l_1, l'} | \mathbf{x}_{l_1}, \mathbf{x}_{l'}) d\mathbf{x}_{\sim l}, \quad (44)$$

while the prediction step (8) remains unchanged. A feasible and, typically, accurate approximation of $f(\mathbf{x}_l | \mathbf{y})$ in (44) can be obtained by evaluating

$$b^{(p)}(\mathbf{x}_l) \propto f(\mathbf{x}_l) \prod_{l' \in \mathcal{C}_l} \int f(\mathbf{y}_{l_1, l'} | \mathbf{x}_l, \mathbf{x}_{l'}) b^{(p-1)}(\mathbf{x}_{l'}) d\mathbf{x}_{l'} \quad (45)$$

for iteration index $p = 1, \dots, P$, where $b^{(0)}(\mathbf{x}_{l'}) = f(\mathbf{x}_{l'})$, $l' \in \mathcal{C}_l$. This amounts to the BP-based SPAWN scheme presented in [17]. All quantities involved in (45) are locally available at CA l or can be made available by communicating only with the neighbor CAs $l' \in \mathcal{C}_l$. A sample-based implementation of (45) is discussed in [48] and [37].

2) *Control Layer*: Since there are no targets, the component $D_I(\mathbf{u}^+) = I(\mathbf{x}_C, \mathbf{x}_T^+; \mathbf{y}^+; \mathbf{u}^+)$ of the objective function in (26) simplifies to $D_I(\mathbf{u}^+) = I(\mathbf{x}_C; \mathbf{y}^+; \mathbf{u}^+)$. The expression of the gradient of $D_I(\mathbf{u}^+)$ in (33) and (34) simplifies as well because $f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_l, \mathbf{x}_C, \mathbf{x}_T^+; \mathbf{u}_{C_l}^+) = f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_l, \mathbf{x}_C; \mathbf{u}_{C_l}^+)$ and $f(\mathbf{y}^+ | \mathbf{x}_C, \mathbf{x}_T^+; \mathbf{u}^+) = f(\mathbf{y}^+ | \mathbf{x}_C; \mathbf{u}^+)$ (according to (32) and (29), since $\mathcal{T} = \emptyset$); furthermore, sampling from $f(\mathbf{x}_C, \mathbf{x}_T^+)$ (see Section VI) reduces to sampling from $f(\mathbf{x}_C)$.

B. Cooperative Estimation of Global States

Next, we discuss the case where the local states of the CAs are known, and thus our task is only the distributed, cooperative estimation of the target states.

1) *Estimation Layer*: The marginal posteriors corresponding to the CAs are no longer calculated, and the correction step (10) in the calculation of the marginal posterior of the m th target simplifies to

$$f(\mathbf{x}_m | \mathbf{y}) \propto f(\mathbf{x}_m) \prod_{l \in \mathcal{C}_m} f(\mathbf{y}_{l,m} | \mathbf{x}_l, \mathbf{x}_m), \quad (46)$$

where $f(\mathbf{x}_m)$ is calculated according to (9). A computationally feasible sample-based approximation of sequential state estimation as given by (46) and (9) is provided by the particle filter [7], [40], [49].

The product of local likelihood functions $\prod_{l \in \mathcal{C}_m} f(\mathbf{y}_{l,m} | \mathbf{x}_l, \mathbf{x}_m)$ is not available at the CAs. However, as in Section III-B, an approximation of these products can be provided to each CA in a distributed manner by a consensus (or gossip) algorithm performed in parallel for each sample weight [25], [38], [39] or by the likelihood consensus scheme [7], [24], [40].

For the calculations in the control layer (described presently), a common set of samples is required at each CA. This can be ensured by additionally using, e.g., a max-consensus and providing all CAs with the same seed for random number generation [39], [41].

2) *Control Layer*: Since there are no unknown CA states, the objective function in (26) simplifies in that $D_I(\mathbf{u}^+) = I(\mathbf{x}_T^+; \mathbf{y}^+; \mathbf{u}^+)$ and $G_l(\mathbf{u}_l^+) = 0$ for all $l \in \mathcal{C}$. The expression of the gradient of $D_I(\mathbf{u}^+)$ in (33) and (34) simplifies because $f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_l, \mathbf{x}_C, \mathbf{x}_T^+; \mathbf{u}_{C_l}^+) = f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_l, \mathbf{x}_T^+; \mathbf{u}_{C_l}^+)$ and $f(\mathbf{y}^+ | \mathbf{x}_C, \mathbf{x}_T^+; \mathbf{u}^+) = f(\mathbf{y}^+ | \mathbf{x}_T^+; \mathbf{u}^+)$; furthermore, sampling from $f(\mathbf{x}_C, \mathbf{x}_T^+)$ reduces to sampling from $f(\mathbf{x}_T^+)$.

This special case was previously considered in [15]. More specifically, [15] studied estimation of one static global

state and proposed a distributed, gradient-based, information-seeking controller and a sample-based implementation.

IX. SIMULATION RESULTS

We demonstrate the performance of the proposed method for three different localization scenarios. In Section IX-B, we study the behavior of the controller by considering noncooperative self-localization of four mobile CAs based on distance measurements relative to an anchor. In Section IX-C, we consider cooperative self-localization of three mobile CAs. Finally, in Section IX-D, two mobile CAs perform cooperative simultaneous self-localization and tracking of a target. Further simulation results demonstrating the performance of the estimation layer in larger networks are reported in [25]. Simulation source files and animated plots are available at <http://www.nt.tuwien.ac.at/about-us/staff/florian-meyer/>.

A. Simulation Setup

The following aspects of the simulation setup are common to all three scenarios. The states of the CAs consist of their 2D position, i.e., $\mathbf{x}_l^{(n)} \triangleq [x_{l,1}^{(n)}, x_{l,2}^{(n)}]^T$ in a global reference frame. In addition to the mobile CAs, there is one anchor CA (indexed by $l = 1$), i.e., a static CA that broadcasts its own (true) position to the mobile CAs but does not perform any measurements. The CA network is fully connected. The states of the mobile CAs evolve independently according to [35]

$$\mathbf{x}_l^{(n)} = \mathbf{x}_l^{(n-1)} + \mathbf{u}_l^{(n)} + \mathbf{q}_l^{(n)}, \quad n = 1, 2, \dots \quad (47)$$

Here, $\mathbf{q}_l^{(n)} \in \mathbb{R}^2$ is zero-mean Gaussian with independent and identically distributed entries, i.e., $\mathbf{q}_l^{(n)} \sim \mathcal{N}(\mathbf{0}, \sigma_q^2 \mathbf{I})$ with $\sigma_q^2 = 10^{-3}$, and $\mathbf{q}_l^{(n)}$ and $\mathbf{q}_{l'}^{(n')}$ are independent unless $(l, n) = (l', n')$. The admissible set \mathcal{U}_l of the control vector $\mathbf{u}_l^{(n)}$ is defined by the norm constraint $\|\mathbf{u}_l^{(n)}\| \leq u_l^{\max}$. For the interpretation of $\mathbf{u}_l^{(n)}$ within (47), it is assumed that the CAs know the orientation of the global reference frame. In the initialization of the algorithms, at time $n = 0$, we use a state prior that is uniform on $[-200, 200] \times [-200, 200]$.

The mobile CAs acquire distance measurements according to (4), i.e., $y_{l,k}^{(n)} = \|\mathbf{x}_l^{(n)} - \mathbf{x}_k^{(n)}\| + v_{l,k}^{(n)}$, where the measurement noise $v_{l,k}^{(n)}$ is independent across l, k , and n and Gaussian with variance

$$\sigma_{l,k}^{(n)2} = \begin{cases} \sigma_0^2, & \|\mathbf{x}_l^{(n)} - \mathbf{x}_k^{(n)}\| \leq d_0 \\ \sigma_0^2 \left[\left(\frac{\|\mathbf{x}_l^{(n)} - \mathbf{x}_k^{(n)}\|}{d_0} - 1 \right)^\kappa + 1 \right], & \|\mathbf{x}_l^{(n)} - \mathbf{x}_k^{(n)}\| > d_0. \end{cases} \quad (48)$$

That is, $\sigma_{l,k}^{(n)2}$ is a function of $\|\mathbf{x}_l^{(n)} - \mathbf{x}_k^{(n)}\|$ that stays constant up to some distance d_0 and then increases polynomially with some exponent κ . This is a simple model for time-of-arrival distance measurements [50]. We set $\sigma_0^2 = 50$ and $\kappa = 2$ and, if not stated otherwise, $d_0 = 50$.

In the estimation layer, we use $J = 3.600$ samples. (Choosing J below 3.000 was observed in some rare cases to lead to a convergence to the wrong estimate.) A resampling step is performed to avoid weight degeneracy [51]. Resampling transforms weighted samples $\{\{\tilde{\mathbf{x}}_k^{(n)(j)}, w_k^{(n)(j)}\}_{j=1}^J$ representing

the belief $b(\mathbf{x}_k^{(n)})$ into nonweighted samples $\{\mathbf{x}_k^{(n)(j)}\}_{j=1}^J$. (We note that weighted samples arise in the estimation layer, as discussed in [24], [25].) We use a somewhat nonorthodox type of resampling that helps move samples to positions with high probability mass, thereby reducing the number of samples needed. More specifically, at every L th time step n , we sample from a kernel approximation of the belief; at all other time steps, we perform standard systematic resampling [51]. The kernel approximation of the belief $b(\mathbf{x}_k^{(n)})$ is obtained as [52]

$$\tilde{b}(\mathbf{x}_k^{(n)}) = \sum_{j=1}^J w_k^{(n)(j)} K(\mathbf{x}_k^{(n)} - \tilde{\mathbf{x}}_k^{(n)(j)}),$$

with the Gaussian kernel $K(\mathbf{x}) = (2\pi\sigma_K^2)^{-1} \exp(-\|\mathbf{x}\|^2 / (2\sigma_K^2))$. Here, the variance σ_K^2 is chosen as $\sigma_K^2 = J^{1/3} T_k^{(n)} / 2$ if $T_k^{(n)} < 2\sigma_0^2$ and $\sigma_K^2 = \sigma_0^2$ otherwise, where $T_k^{(n)}$ denotes the trace of the weighted sample covariance matrix defined as

$$\mathbf{C}_k^{(n)} = \sum_{j=1}^J w_k^{(n)(j)} \tilde{\mathbf{x}}_k^{(n)(j)} \tilde{\mathbf{x}}_k^{(n)(j)T} - \boldsymbol{\mu}_k^{(n)} \boldsymbol{\mu}_k^{(n)T},$$

with $\boldsymbol{\mu}_k^{(n)} = \sum_{j=1}^J w_k^{(n)(j)} \tilde{\mathbf{x}}_k^{(n)(j)}$. This case distinction in choosing σ_K^2 is used since $\sigma_K^2 = J^{1/3} T_k^{(n)} / 2$ is only accurate for a unimodal distribution [52] whereas $\sigma_K^2 = \sigma_0^2$ is suitable for annularly shaped distributions (here, the width of the annulus is determined by σ_0^2 [48]). We choose $L = 40$ if $T_k^{(n)} < 80$, $L = 20$ if $80 \leq T_k^{(n)} < 1000$, and $L = 10$ if $T_k^{(n)} \geq 1000$; this led to good results in our simulation setting.

We employ a censoring scheme [37] to reduce the number of samples and avoid numerical problems during the first time steps, where the mobile CAs still have uninformative beliefs. More specifically, only CAs l with $T_k^{(n)} < 10$ are used as localization partners by neighbor CAs and (in our third scenario) are involved in localizing the target. In the control layer, this censoring scheme corresponds to the following strategy: as long as CA l is not localized (i.e., $T_k^{(n)} \geq 10$), its objective function is $\tilde{D}_h(\mathbf{u}^{(n+1)}) \triangleq -h(\mathbf{x}_l^{(n+1)} | y_{l,1}^{(n+1)}; y_{l,1}^{(1:n)}, \mathbf{u}_l^{(1:n+1)})$, i.e., the negative differential entropy of only the own state conditioned on only the own measurement relative to the anchor CA, $y_{l,1}^{(n+1)}$.

The local gradient ascents in the controller (see (19)) use the reference points $\mathbf{u}_{r,l}^{(n)} = \mathbf{0}$, which are consistent with the state evolution model (47), and step sizes $c_l^{(n)}$ chosen such that $\|\mathbf{u}_l^{(n)}\| = u_l^{\max}$. Thus, each CA $l \in \mathcal{C}$ moves with maximum nominal speed (determined by u_l^{\max}) in the direction of maximum local increase of the objective function. If not stated otherwise, the number of samples used in the control layer is $JJ' = 60.000$, with $J = 1.200$ and $J' = 50$. (The $J = 1.200$ samples are obtained by random selection from the 3.600 samples produced by the estimation layer.) We note that a reduction of J' to $J'=1$ was observed to result in more jagged CA trajectories and a slightly slower reduction of the estimation error over time.

B. Noncooperative Self-Localization

To study the behavior of the controller, we consider four mobile CAs $l = 2, 3, 4, 5$ that perform self-localization without

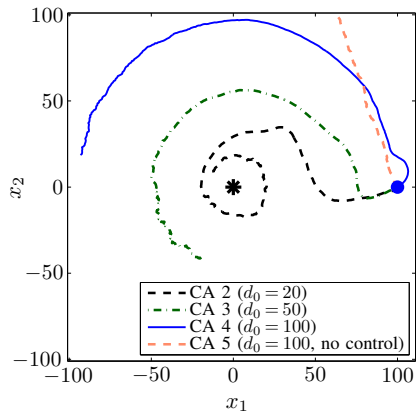


Fig. 5. Example trajectories for noncooperative self-localization with information-seeking control (except CA 5). The initial CA position and the anchor position are indicated by a bullet and a star, respectively.

any cooperation during 300 time steps n . The mobile CAs measure their distance to the static anchor CA ($l = 1$), which is located at position $[0, 0]^T$, but they do not measure any distances between themselves. Their measurement models use different values of d_0 , namely, $d_0 = 20, 50, 100$, and 100 for $l = 2, 3, 4$, and 5 , respectively. The mobile CAs start at position $[100, 0]^T$ and move with identical nominal speed determined by $u_l^{\max} = 1$. The objective function for the control of CAs 2, 3, and 4 is $\tilde{D}_h(\mathbf{u}_l^{(n+1)}) \triangleq -h(\mathbf{x}_l^{(n+1)} | y_{l,1}^{(n+1)}; y_{l,1}^{(1:n)}, \mathbf{u}_l^{(1:n+1)})$. CA 5 is not controlled; it randomly chooses a direction at time $n = 1$ and then moves in that direction with constant nominal speed determined by $u_l^{\max} = 1$. Fig. 5 shows an example of the trajectories of the four mobile CAs. These trajectories are quite different because of the different values of d_0 and the fact that CA 5 is not controlled. CA 4, after an initial turn, is roughly localized in the sense that the shape of its marginal posterior has changed from an annulus to only a segment of an annulus. Thereafter, CA 4 turns around the anchor, which is reasonable in view of the single distance measurement available at each time n and the fact that, since $d_0 = 100$, the measurement noise cannot be decreased by approaching the anchor. CA 3 (with $d_0 = 50$) initially approaches the anchor. At a distance of 50 to the anchor, the measurement noise cannot be decreased any more, and thus CA 3 turns around the anchor without approaching it further. A similar behavior is exhibited by CA 2 (with $d_0 = 20$).

Fig. 6 shows the self-localization root-mean-square errors (RMSEs) of the four mobile CAs. These RMSEs were determined at each time n by averaging over 300 simulation runs. As can be seen, the three CAs performing information-seeking control ($l = 2, 3, 4$) are fairly well localized after about 100 time steps. CA 2 (with $d_0 = 20$) takes longer to localize itself than CAs 3 and 4 since, prior to reaching a distance of 20 to the anchor, it has a larger noise variance (see (48)). The performance of CA 3 and CA 4 is almost identical; the larger noise variance of CA 3 during the initial time steps is compensated by a smaller turning radius once a distance of 50 to the anchor has been reached. CA 5 is unable to localize itself, due to the lack of intelligent control.




Fig. 6. Self-localization RMSE for noncooperative self-localization with information-seeking control (except CA 5).

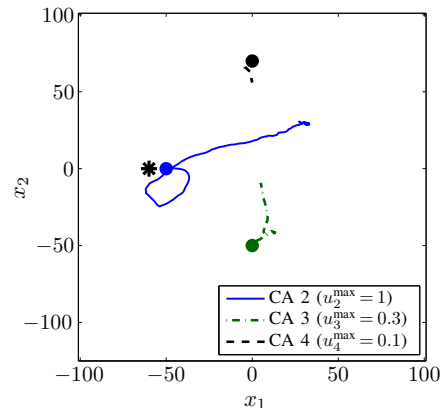


Fig. 7. Example trajectories for cooperative self-localization with information-seeking control (C-C scheme). The initial CA positions and the anchor position are indicated by bullets and a star, respectively.

C. Cooperative Self-Localization

Next, we study the proposed method for cooperative self-localization with information-seeking control (abbreviated as C-C). There are three mobile CAs $l = 2, 3, 4$ with different start points $([-50, 0]^T, [0, -50]^T, \text{ and } [0, 70]^T$ for $l = 2, 3$, and 4 , respectively) and different nominal speeds ($u_l^{\max} = 1, 0.3$, and 0.1 for $l = 2, 3$, and 4 , respectively). The mobile CAs measure their distances to a static anchor $l = 1$ located at $[-60, 0]^T$ and to each other, using $d_0 = 50$. Example trajectories are shown in Fig. 7. For comparison, we also consider noncooperative self-localization with information-seeking control as studied in Section IX-B (abbreviated as N-C). Finally, we consider another scheme (abbreviated as C-N) where the CAs cooperate in the estimation layer but no intelligent control is performed. Here, each CA randomly chooses a direction and then moves in that direction with constant nominal speed determined by u_l^{\max} .

Fig. 8 shows the self-localization RMSEs of the three schemes, which were determined by averaging over the three mobile CAs and over 300 simulation runs. It is seen that the RMSEs of the two reference schemes N-C and C-N decrease only very slowly whereas, after about 100 time steps, the RMSE of the proposed C-C scheme has decreased to a low value. This behavior can be explained as follows. Without cooperation (N-C) or without intelligent control (C-N), CAs 3




Fig. 8. Self-localization RMSE of the proposed estimation/control method and of two reference methods.

and 4 need a long time to localize themselves because they are slow and initially far away from the anchor. On the other hand, CA 2 localizes itself very quickly because it is fast and initially close to the anchor. With cooperation and control (C-C), CA 2 moves in such a way that it supports the self-localization of the two other CAs. In fact, as shown by Fig. 7, CA 2 first localizes itself by starting to turn around the anchor and then makes a sharp turn to approach CAs 3 and 4, which helps them localize themselves. This demonstrates the function and benefits of cooperative estimation and control.

D. Cooperative Self-Localization and Target Tracking

Finally, we consider cooperative simultaneous self-localization and target tracking. Two mobile CAs $l = 2, 3$ starting at position $[20, 20]^T$ and $[-10, -10]^T$, respectively and with nominal speed determined by $u_l^{\max} = 1$ cooperatively localize and track themselves and a mobile target. There is also a static anchor $l = 1$ at position $[-50, 0]^T$. The target state $\mathbf{x}_m^{(n)} = \mathbf{x}_4^{(n)}$ consists of position and velocity, i.e., $\mathbf{x}_4^{(n)} \triangleq [x_{4,1}^{(n)}, x_{4,2}^{(n)}, \dot{x}_{4,1}^{(n)}, \dot{x}_{4,2}^{(n)}]^T$. The target state evolves according to [35]

$$\mathbf{x}_4^{(n)} = \mathbf{G}\mathbf{x}_4^{(n-1)} + \mathbf{W}\mathbf{q}_4^{(n)}, \quad n=1, 2, \dots,$$

where

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

and $\mathbf{q}_4^{(n)} \in \mathbb{R}^2$ is zero-mean Gaussian with independent and identically distributed entries, i.e., $\mathbf{q}_4^{(n)} \sim \mathcal{N}(\mathbf{0}, \tilde{\sigma}_q^2 \mathbf{I})$ with $\tilde{\sigma}_q^2 = 10^{-5}$, and with $\mathbf{q}_4^{(n)}$ and $\mathbf{q}_4^{(n')}$ independent unless $n = n'$. The target trajectory is initialized with position $[x_{4,1}^{(0)}, x_{4,2}^{(0)}]^T = [50, 0]^T$ and velocity $[\dot{x}_{4,1}^{(0)}, \dot{x}_{4,2}^{(0)}]^T = [0.05, 0.05]^T$. In the initialization of the algorithms, we use a target position prior that is uniform on $[-200, 200] \times [-200, 200]$ and a target velocity prior that is Gaussian with mean $[0, 0]^T$ and covariance matrix $\text{diag}\{10^{-1}, 10^{-1}\}$. The number of samples used in the estimation layer is $J = 120,000$; the number of samples used in the control layer is $JJ' = 6,000$, with $J = 1,200$ and $J' = 5$. Fig. 9 shows an example of CA and target trajectories obtained with the proposed method for cooperative localization with




Fig. 9. Example trajectories for cooperative simultaneous self-localization and target tracking with information-seeking control (C-C scheme). The initial CA positions are indicated by bullets, the initial target position by a cross, and the anchor position by a star.

information-seeking control (C-C). One can observe that the two CAs first start turning around the anchor to localize themselves and then approach the target. Finally, at a distance of 50 to the target, where further approaching the target would no longer decrease the measurement noise, the CAs spread out to achieve a geometric formation that is favorable for cooperatively localizing and tracking the target.

As before, we compare our C-C method with two reference methods, namely, noncooperative localization with information-seeking control (N-C) and cooperative localization with fixed, randomly chosen directions of movement (C-N). Fig. 10 shows the self-localization RMSEs and target localization RMSEs of the three schemes, which were determined by averaging over the two CAs and over 100 simulation runs. The following observations can be made:

- The self-localization performance of C-N is very poor: after an initial decrease, the RMSE slowly increases. In fact, typically, no cooperation actually takes place, since the CAs are unable to localize themselves and thus each CA is censored by the respective other CA. The self-localization RMSEs of C-C and N-C decrease rather quickly to a low value. They are very similar, which can be explained as follows. Because both CAs move with the same nominal speed, they localize themselves approximately in the same manner. Therefore, as long as the CAs are not localized, no cooperation takes place due to censoring, and after they are localized, no further gain can be achieved by cooperation.
- The target localization RMSEs of the three methods are initially equal to 50 and slowly increase during the first 40 time steps. Indeed, due to the censoring scheme, the CAs start localizing the target only when they are localized themselves. Therefore, during the first 40 time steps, no measurements of the distance to the target are used by the CAs, and thus the CAs' target position estimation is solely based on the prior distribution, which is uniform. This leads to a target position estimate of $[0, 0]^T$ and in turn (since the target is initially located at $[50, 0]^T$) to an initial target localization RMSE of 50 at time $n = 1$.

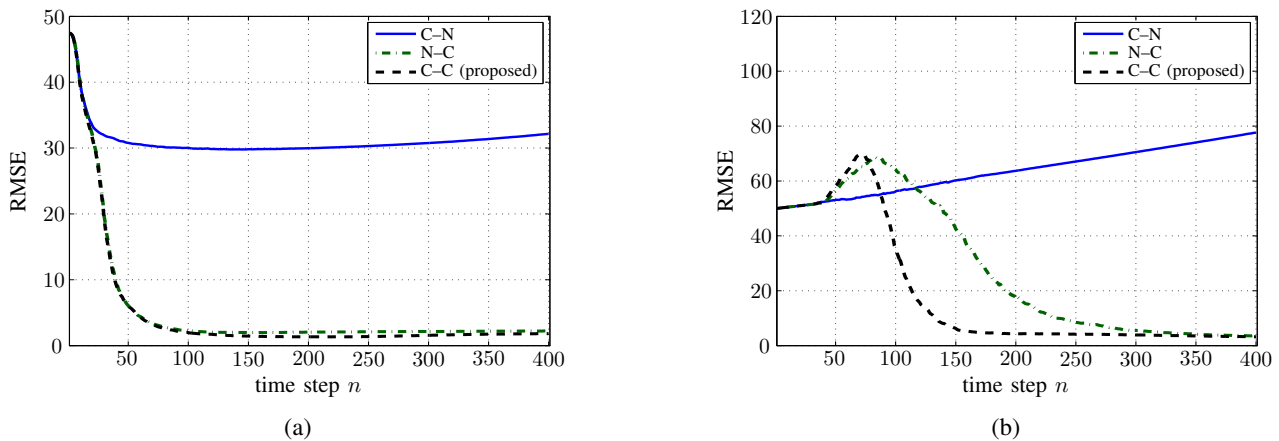


Fig. 10. Performance of three different methods for simultaneous self-localization and target localization: (a) Self-localization RMSE, (b) target-localization RMSE.

During the first 40 time steps, the RMSE slowly increases since the target slowly moves away from $[0, 0]^T$. The RMSE of C–N continues to increase in this manner even after $n = 40$ since with C–N, the CAs are never localized and therefore never start localizing the target. For C–C and N–C (both employing information-seeking control), after $n = 40$, the RMSE first increases and then decreases. The RMSE of C–C decreases sooner and more quickly than that of N–C, which again shows the benefits of cooperative estimation.

The initial increase and subsequent decrease of the target localization RMSE observed with C–C and N–C after $n = 40$ can be explained as follows. After the CAs localized themselves and start localizing the target, the target position posterior at a given CA is roughly annularly shaped, with the center of the annulus being the CA position. (This position is equal to the turning point of the respective CA trajectory in Fig. 9.) The resulting target position estimate is located at that center. Thus, it is more distant from the true target position than the estimate $[0, 0]^T$ that was obtained when the CA was not yet localized and the target position posterior was still uniform. As the CAs approach the target, the target position posterior becomes unimodal and the target can be localized, resulting in a decrease of the target localization RMSE.

X. CONCLUSION

We proposed a Bayesian framework and method for distributed estimation with information-seeking control in agent networks. Distributed, cooperative estimation is performed for time-varying global states (related to noncooperative targets or features of the environment) and/or time-varying local states (related to individual cooperative agents), using a combination of belief propagation message passing and consensus. The distributed, cooperative control seeks to optimize the behavior of the cooperative agents by maximizing the negative joint posterior entropy of the agent states via a gradient ascent. A probabilistic information transfer from the estimation layer to the control layer enables effective control strategies and thus leads to excellent estimation performance.

A major advantage of the proposed approach is its generality. Our method relies on general state evolution and measurement models, an information-theoretic objective function for control, and sample-based representations of probability distributions. These characteristics make it suitable for nonlinear and non-Gaussian systems, such as those arising in location-aware networks. Numerical simulations for a simultaneous agent self-localization and target tracking problem demonstrated intelligent behavior of the cooperative agents and a resulting improvement of estimation performance.

Possible directions for future research include an extension of the myopic controller (i.e., optimizing only one time step ahead) to a receding horizon [53]; this can be expected to improve the performance in scenarios with multiple time-varying global states. Furthermore, the complexity and communication cost of the proposed method can be reduced by introducing Gaussian or Gaussian mixture approximations [54] and using cubature points [55] instead of random samples.

APPENDIX A PROOF OF EQUATION (21)

We will use the following transformation rule for differential entropy [56, Eq. 18]: For a continuous random vector \mathbf{a} and a transformed random vector of identical dimension $\mathbf{b} = g(\mathbf{a})$, where $g(\cdot)$ is a bijective differentiable function with Jacobian determinant $J_g(\mathbf{a}) = \det \frac{\partial g(\mathbf{a})}{\partial \mathbf{a}}$,

$$h(\mathbf{b}) = h(\mathbf{a}) + e(\mathbf{a}), \quad \text{with } e(\mathbf{a}) \triangleq \int f(\mathbf{a}) \log |J_g(\mathbf{a})| d\mathbf{a}. \quad (49)$$

The conditional differential entropy $h(\mathbf{x}^+ | \mathbf{y}^+; \mathbf{u}^+)$ can be expanded as [32, Chap. 8]

$$h(\mathbf{x}^+ | \mathbf{y}^+; \mathbf{u}^+) = h(\mathbf{x}^+, \mathbf{y}^+; \mathbf{u}^+) - h(\mathbf{y}^+; \mathbf{u}^+). \quad (50)$$

The vector \mathbf{x}^+ consists of \mathbf{x}_l^+ and $\mathbf{x}_{A \setminus \{l\}}^+ \triangleq [\mathbf{x}_k^+]_{k \in A \setminus \{l\}}$, and there is $\mathbf{x}_l^+ = \tilde{g}_l(\mathbf{x}_l, \mathbf{u}_l^+)$ (see (20)). Thus, the first term on the right-hand side of (50) can be expressed as $h(\mathbf{x}^+, \mathbf{y}^+; \mathbf{u}^+) = h(\tilde{g}_l(\mathbf{x}_l, \mathbf{u}_l^+), \mathbf{x}_{A \setminus \{l\}}^+, \mathbf{y}^+; \mathbf{u}^+)$. Applying the transformation rule (49) to the “extended state evolution mapping” $\tilde{g}_l^* : [\mathbf{x}_l^T, \mathbf{x}_{A \setminus \{l\}}^T, \mathbf{y}^{+T}]^T \mapsto [(\tilde{g}_l(\mathbf{x}_l, \mathbf{u}_l^+))^T, \mathbf{x}_{A \setminus \{l\}}^T, \mathbf{y}^{+T}]^T$, we then obtain

$$h(\mathbf{x}^+, \mathbf{y}^+; \mathbf{u}^+) = h(\mathbf{x}_l, \mathbf{x}_{\mathcal{A} \setminus \{l\}}^+, \mathbf{y}^+; \mathbf{u}^+) + e(\mathbf{x}_l, \mathbf{x}_{\mathcal{A} \setminus \{l\}}^+, \mathbf{y}^+; \mathbf{u}_l^+), \quad (51)$$

where

$$e(\mathbf{x}_l, \mathbf{x}_{\mathcal{A} \setminus \{l\}}^+, \mathbf{y}^+; \mathbf{u}_l^+) \triangleq \int \int \int f(\mathbf{x}_l, \mathbf{x}_{\mathcal{A} \setminus \{l\}}^+, \mathbf{y}^+) \log |J_{\tilde{g}_l^*}(\mathbf{x}_l, \mathbf{x}_{\mathcal{A} \setminus \{l\}}^+, \mathbf{y}^+; \mathbf{u}_l^+)| \times d\mathbf{x}_l d\mathbf{x}_{\mathcal{A} \setminus \{l\}}^+ d\mathbf{y}^+.$$

Here, $J_{\tilde{g}_l^*}(\mathbf{x}_l, \mathbf{x}_{\mathcal{A} \setminus \{l\}}^+, \mathbf{y}^+; \mathbf{u}_l^+)$ is the Jacobian determinant of $\tilde{g}_l^*(\mathbf{x}_l, \mathbf{x}_{\mathcal{A} \setminus \{l\}}^+, \mathbf{y}^+; \mathbf{u}_l^+)$. It is easily seen that $J_{\tilde{g}_l^*}(\mathbf{x}_l, \mathbf{x}_{\mathcal{A} \setminus \{l\}}^+, \mathbf{y}^+; \mathbf{u}^+) = J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}^+)$, and thus we obtain further

$$\begin{aligned} e(\mathbf{x}_l, \mathbf{x}_{\mathcal{A} \setminus \{l\}}^+, \mathbf{y}^+; \mathbf{u}_l^+) &= \int \left[\int \int f(\mathbf{x}_l, \mathbf{x}_{\mathcal{A} \setminus \{l\}}^+, \mathbf{y}^+) d\mathbf{x}_{\mathcal{A} \setminus \{l\}}^+ d\mathbf{y}^+ \right] \\ &\quad \times \log |J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}_l^+)| d\mathbf{x}_l \\ &= \int f(\mathbf{x}_l) \log |J_{\tilde{g}_l}(\mathbf{x}_l; \mathbf{u}_l^+)| d\mathbf{x}_l \\ &= G_l(\mathbf{u}_l^+). \end{aligned} \quad (52)$$

Inserting (52) into (51) and the resulting expression of $h(\mathbf{x}^+, \mathbf{y}^+; \mathbf{u}^+)$ into (50) gives

$$h(\mathbf{x}^+ | \mathbf{y}^+; \mathbf{u}^+) = h(\mathbf{x}_l, \mathbf{x}_{\mathcal{A} \setminus \{l\}}^+, \mathbf{y}^+; \mathbf{u}^+) + G_l(\mathbf{u}_l^+) - h(\mathbf{y}^+; \mathbf{u}^+). \quad (53)$$

Next, we repeat this transformation procedure but apply it to the term $h(\mathbf{x}_l, \mathbf{x}_{\mathcal{A} \setminus \{l\}}^+, \mathbf{y}^+; \mathbf{u}^+)$ in (53) instead of $h(\mathbf{x}^+, \mathbf{y}^+; \mathbf{u}^+)$. Consider an arbitrary $l' \in \mathcal{C} \setminus \{l\}$, and note that $\mathbf{x}_{\mathcal{A} \setminus \{l\}}^+$ consists of $\mathbf{x}_{l'}^+$ and $\mathbf{x}_{\mathcal{A} \setminus \{l, l'\}}^+ \triangleq [\mathbf{x}_k^+]_{k \in \mathcal{A} \setminus \{l, l'\}}$, where $\mathbf{x}_{l'}^+ = \tilde{g}_{l'}(\mathbf{x}_{l'}, \mathbf{u}_{l'}^+)$ according to (20). Proceeding as above and inserting the resulting expression of $h(\mathbf{x}_l, \mathbf{x}_{\mathcal{A} \setminus \{l\}}^+, \mathbf{y}^+; \mathbf{u}^+)$ into (53) yields

$$h(\mathbf{x}^+ | \mathbf{y}^+; \mathbf{u}^+) = h(\mathbf{x}_l, \mathbf{x}_{l'}, \mathbf{x}_{\mathcal{A} \setminus \{l, l'\}}^+, \mathbf{y}^+; \mathbf{u}^+) + G_{l'}(\mathbf{u}_{l'}^+) + G_l(\mathbf{u}_l^+) - h(\mathbf{y}^+; \mathbf{u}^+).$$

We continue this procedure in a recursive fashion, splitting off CA state vectors from $\mathbf{x}_{\mathcal{A} \setminus \{l, l'\}}^+$ until only the target states (contained in $\mathbf{x}_{\mathcal{T}}^+$) are left, and applying the transformation rule at each recursion. In the end, we obtain

$$h(\mathbf{x}^+ | \mathbf{y}^+; \mathbf{u}^+) = h(\mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+, \mathbf{y}^+; \mathbf{u}^+) + \sum_{l \in \mathcal{C}} G_l(\mathbf{u}_l^+) - h(\mathbf{y}^+; \mathbf{u}^+).$$

Finally, Equation (21) is obtained by noting that $h(\mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+, \mathbf{y}^+; \mathbf{u}^+) = h(\mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+ | \mathbf{y}^+; \mathbf{u}^+) + h(\mathbf{y}^+; \mathbf{u}^+)$.

APPENDIX B

DERIVATION OF (33) AND (34)

1) *Derivation of (33)*: Let us first define the vector $\check{\mathbf{y}}_l^+ \triangleq [\mathbf{y}_{l', k}^+]_{l' \in \mathcal{C} \setminus \{l\}, k \in \mathcal{A} \setminus \{l\}}$, which contains all those measurements $\mathbf{y}_{l', k}^+$ that are not contained in $\tilde{\mathbf{y}}_l^+$ (cf. (31)). The corresponding likelihood function is given by

$$f(\check{\mathbf{y}}_l^+ | \mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+; \mathbf{u}^+) = \frac{f(\mathbf{y}^+ | \mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+; \mathbf{u}^+)}{f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_l, \mathbf{x}_{\mathcal{C}_l}, \mathbf{x}_{\mathcal{T}_l}^+; \mathbf{u}_{\mathcal{C}_l}^+)}, \quad (54)$$

which, according to (29) and (32), involves all factors of $f(\mathbf{y}^+ | \mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+; \mathbf{u}^+)$ that do not depend on the local control vector \mathbf{u}_l^+ . Using (54) in (28) yields

$$\begin{aligned} \frac{\partial D_I(\mathbf{u}^+)}{\partial \mathbf{u}_l^+} &= \int \int \int f(\mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+) f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+; \mathbf{u}^+) \\ &\quad \times \frac{\partial f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_l, \mathbf{x}_{\mathcal{C}_l}, \mathbf{x}_{\mathcal{T}_l}^+; \mathbf{u}_{\mathcal{C}_l}^+)}{\partial \mathbf{u}_l^+} \\ &\quad \times \log \frac{f(\mathbf{y}^+ | \mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+; \mathbf{u}^+)}{f(\mathbf{y}^+; \mathbf{u}^+)} d\mathbf{x}_{\mathcal{C}} d\mathbf{x}_{\mathcal{T}}^+ d\mathbf{y}^+. \end{aligned} \quad (55)$$

Setting $\mathbf{u}^+ = \mathbf{u}_r^+$, and multiplying and dividing the integrand in (55) by $f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_l, \mathbf{x}_{\mathcal{C}_l}, \mathbf{x}_{\mathcal{T}_l}^+; \mathbf{u}_{r, \mathcal{C}_l}^+)$, we obtain further

$$\begin{aligned} \frac{\partial D_I(\mathbf{u}^+)}{\partial \mathbf{u}_l^+} \Big|_{\mathbf{u}^+ = \mathbf{u}_r^+} &= \int \int \int q(\mathbf{y}^+, \mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+) \frac{1}{f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_l, \mathbf{x}_{\mathcal{C}_l}, \mathbf{x}_{\mathcal{T}_l}^+; \mathbf{u}_{r, \mathcal{C}_l}^+)} \\ &\quad \times \frac{\partial f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_l, \mathbf{x}_{\mathcal{C}_l}, \mathbf{x}_{\mathcal{T}_l}^+; \mathbf{u}_{\mathcal{C}_l}^+)}{\partial \mathbf{u}_l^+} \Big|_{\mathbf{u}_{\mathcal{C}_l}^+ = \mathbf{u}_{r, \mathcal{C}_l}^+} \\ &\quad \times \log \frac{f(\mathbf{y}^+ | \mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+; \mathbf{u}_r^+)}{f(\mathbf{y}^+; \mathbf{u}_r^+)} d\mathbf{x}_{\mathcal{C}} d\mathbf{x}_{\mathcal{T}}^+ d\mathbf{y}^+, \end{aligned} \quad (56)$$

where

$$q(\mathbf{y}^+, \mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+) \triangleq f(\mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+) f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_l, \mathbf{x}_{\mathcal{C}_l}, \mathbf{x}_{\mathcal{T}_l}^+; \mathbf{u}_{r, \mathcal{C}_l}^+) \times f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+; \mathbf{u}_r^+).$$

Then, (33) is recognized to be a Monte Carlo approximation of (56) that is obtained by performing importance sampling [43] using $q(\mathbf{y}^+, \mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+)$ as importance density, i.e., the samples $\mathbf{y}^{+(j, j')}$, $\mathbf{x}_{\mathcal{C}}^{(j)}$, and $\mathbf{x}_{\mathcal{T}}^{+(j)}$ occurring in (33) are drawn from $q(\mathbf{y}^+, \mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+)$. Using (54), this importance density can be expressed as

$$\begin{aligned} q(\mathbf{y}^+, \mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+) &= f(\mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+) f(\mathbf{y}^+ | \mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+; \mathbf{u}_r^+) \\ &= f(\mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+, \mathbf{y}^+; \mathbf{u}_r^+). \end{aligned}$$

The first expression, $f(\mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+) f(\mathbf{y}^+ | \mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+; \mathbf{u}_r^+)$, underlies the two-stage sampling procedure described in Section V.

2) *Derivation of (34)*: We have

$$f(\mathbf{y}^+; \mathbf{u}_r^+) = \int \int f(\mathbf{y}^+ | \mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+; \mathbf{u}_r^+) f(\mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+) d\mathbf{x}_{\mathcal{C}} d\mathbf{x}_{\mathcal{T}}^+. \quad (57)$$

Using samples $\{\mathbf{x}_{\mathcal{C}}^{(j)}, \mathbf{x}_{\mathcal{T}}^{+(j)}\}_{j=1}^J \sim f(\mathbf{x}_{\mathcal{C}}, \mathbf{x}_{\mathcal{T}}^+)$ (see Section V), a Monte Carlo approximation of (57) is obtained as

$$f(\mathbf{y}^+; \mathbf{u}_r^+) \approx \frac{1}{J} \sum_{j''=1}^J f(\mathbf{y}^+ | \mathbf{x}_{\mathcal{C}}^{(j'')}, \mathbf{x}_{\mathcal{T}}^{+(j'')}; \mathbf{u}_r^+).$$

Evaluating this for $\mathbf{y}^+ = \mathbf{y}^{+(j, j')}$ (again see Section V) yields (34).

APPENDIX C

DRAWING SAMPLES FROM $f(\mathbf{y}^+ | \mathbf{x}_{\mathcal{C}}^{(j)}, \mathbf{x}_{\mathcal{T}}^{+(j)}; \mathbf{u}_r^+)$

We consider the setting of Section VI-A. As discussed there, samples $\{\mathbf{x}_{l'}^{(j)}\}_{j=1}^J \sim f(\mathbf{x}_{l'})$, $l' \in \mathcal{C}$ and $\{\mathbf{x}_m^{+(j)}\}_{j=1}^J \sim$

$f(\mathbf{x}_m^+)$, $m \in \mathcal{T}$ are available at CA l , and it is assumed that the state evolution and measurement models of all CAs $l' \in \mathcal{C}$ are known to CA l . We start by noting that by combining (15) and (3), the composite measurement vector \mathbf{y}^+ can be written as

$$\mathbf{y}^+ = [d_l(\mathbf{x}_l^+, \mathbf{x}_k^+, \mathbf{v}_{l,k}^+)]_{l \in \mathcal{C}, k \in \mathcal{A}_l}. \quad (58)$$

First, CA l obtains samples $\{\mathbf{x}_{l'}^{+(j)}\}_{j=1}^J \sim \tilde{f}(\mathbf{x}_{l'}^+) \triangleq f(\mathbf{x}_{l'}^+) |_{\mathbf{x}_{l'}^+ = \tilde{g}_{l'}(\mathbf{x}_{l'}, \mathbf{u}_{r,l'}^+)}$ (see (20)) for all $l' \in \mathcal{C}$ by evaluating $\tilde{g}_{l'}(\mathbf{x}_{l'}, \mathbf{u}_{l'}^+)$ at $\mathbf{x}_{l'} = \mathbf{x}_{l'}^{(j)}$ and $\mathbf{u}_{l'}^+ = \mathbf{u}_{r,l'}^+$, i.e.,

$$\mathbf{x}_{l'}^{+(j)} = \tilde{g}_{l'}(\mathbf{x}_{l'}^{(j)}, \mathbf{u}_{r,l'}^+), \quad j = 1, \dots, J. \quad (59)$$

Thus, at this point, samples $\{\mathbf{x}_k^{+(j)}\}_{j=1}^J$ for all $k \in \mathcal{A}$ are available at CA l . Next, for each $j \in \{1, \dots, J\}$, CA l draws samples $\{\mathbf{v}_{l',k}^{+(j,j')}\}_{j'=1}^{J'}$ for $l' \in \mathcal{C}$ and $k \in \mathcal{A}_{l'}$. Finally, CA l obtains samples $\{\mathbf{y}^{+(j,j')}\}_{j'=1}^{J'}$ for $j \in \{1, \dots, J\}$ by evaluating (58) using the appropriate samples, i.e.,

$$\mathbf{y}^{+(j,j')} = [d_{l'}(\mathbf{x}_{l'}^{+(j)}, \mathbf{x}_k^{+(j)}, \mathbf{v}_{l',k}^{+(j,j')})]_{l' \in \mathcal{C}, k \in \mathcal{A}_{l'}}, \quad j' = 1, \dots, J'.$$

APPENDIX D

DRAWING SAMPLES FROM $f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_l^{(j)}, \mathbf{x}_{\mathcal{C}_l}^{(j)}, \mathbf{x}_{\mathcal{T}_l}^{+(j)}; \mathbf{u}_{r,\mathcal{C}_l}^+)$

In the setting of Section VI-B, samples $\{\mathbf{x}_{l'}^{(j)}\}_{j=1}^J \sim f(\mathbf{x}_{l'})$, $l' \in \{l\} \cup \mathcal{C}_l$ and $\{\mathbf{x}_m^+\}_{j=1}^J \sim f(\mathbf{x}_m^+)$, $m \in \mathcal{T}_l$ are available at CA l . We start by noting that combining (37) and (3) yields

$$\mathbf{y}_l^+ = [d_l(\mathbf{x}_l^+, \mathbf{x}_k^+, \mathbf{v}_{l,k}^+)]_{k \in \mathcal{A}_l}. \quad (60)$$

Based on the analogy of this expression to (58), CA l first obtains samples $\{\mathbf{y}_l^{+(j,j')}\}_{j'=1}^{J'} \sim f(\mathbf{y}_l^+ | \mathbf{x}_l^{(j)}, \mathbf{x}_{\mathcal{C}_l}^{(j)}, \mathbf{x}_{\mathcal{T}_l}^{+(j)}; \mathbf{u}_{r,\mathcal{C}_l}^+)$ by carrying out the steps of Appendix C with obvious modifications—in particular, \mathbf{y}^+ is replaced by \mathbf{y}_l^+ , \mathcal{C} by $\{l\} \cup \mathcal{C}_l$, and \mathcal{T} by \mathcal{T}_l . More specifically, CA l obtains samples $\{\mathbf{x}_{l'}^{+(j)}\}_{j=1}^J$ for $l' \in \{l\} \cup \mathcal{C}_l$ according to (59). Then, for each $j \in \{1, \dots, J\}$, CA l draws samples $\{\mathbf{v}_{l,k}^{+(j,j')}\}_{j'=1}^{J'} \sim f(\mathbf{v}_{l,k}^+)$ for $k \in \mathcal{A}_l$ and, in turn, obtains samples $\{\mathbf{y}_l^{+(j,j')}\}_{j'=1}^{J'}$ by evaluating (60) using the appropriate samples, i.e.,

$$\mathbf{y}_l^{+(j,j')} = [d_l(\mathbf{x}_l^{+(j)}, \mathbf{x}_k^{+(j)}, \mathbf{v}_{l,k}^{+(j,j')})]_{k \in \mathcal{A}_l}, \quad j' = 1, \dots, J'.$$

It remains to obtain samples of those entries of $\tilde{\mathbf{y}}_l^+$ that are not contained in \mathbf{y}_l^+ (cf. (31) and (37)). More specifically, for each sample $\mathbf{y}_l^{+(j,j')}$, CA l needs to obtain samples $\mathbf{y}_{l',l}^{+(j,j')}$, $l' \in \mathcal{C}_l$. This is done through communication with neighbor CAs: CA l transmits to each neighbor CA $l' \in \mathcal{C}_l$ the samples $\{\mathbf{y}_{l,l'}^{+(j,j')}\}_{j'=1}^{J'}$, $j = 1, \dots, J$, and it receives from CA $l' \in \mathcal{C}_l$ the samples $\{\mathbf{y}_{l',l}^{+(j,j')}\}_{j'=1}^{J'}$, $j = 1, \dots, J$. Thus, finally, samples $\{\tilde{\mathbf{y}}_l^{+(j,j')}\}_{j'=1}^{J'} \sim f(\tilde{\mathbf{y}}_l^+ | \mathbf{x}_l^{(j)}, \mathbf{x}_{\mathcal{C}_l}^{(j)}, \mathbf{x}_{\mathcal{T}_l}^{+(j)}; \mathbf{u}_{r,\mathcal{C}_l}^+)$ are locally available at CA l .

ACKNOWLEDGMENT

The authors would like to thank Dr. Günther Koliander for illuminating discussions.

REFERENCES

- [1] T. Shima and S. Rasmussen, *UAV Cooperative Decision and Control: Challenges and Practical Approaches*. Philadelphia, PA: SIAM, 2009.
- [2] F. Bullo, J. Cortés, and S. Martínez, *Distributed Control of Robotic Networks: A Mathematical Approach to Motion Coordination Algorithms*. Princeton, NJ: Princeton University Press, 2009.
- [3] A. Nayak and I. Stojmenović, *Wireless Sensor and Actuator Networks: Algorithms and Protocols for Scalable Coordination and Data Communication*. Hoboken, NJ: Wiley, 2010.
- [4] F. Zhao and L. J. Guibas, *Wireless Sensor Networks: An Information Processing Approach*. Amsterdam, The Netherlands: Morgan Kaufmann, 2004.
- [5] P. Corke, T. Wark, R. Jurdak, W. Hu, P. Valencia, and D. Moore, “Environmental wireless sensor networks,” *Proc. IEEE*, vol. 98, no. 11, pp. 1903–1917, 2010.
- [6] J. Ko, C. Lu, M. B. Srivastava, J. A. Stankovic, A. Terzis, and M. Welsh, “Wireless sensor networks for healthcare,” *Proc. IEEE*, vol. 98, no. 11, pp. 1947–1960, 2010.
- [7] 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, no. 1, pp. 61–81, 2013.
- [8] T. Zhao and A. Nehorai, “Distributed sequential Bayesian estimation of a diffusive source in wireless sensor networks,” *IEEE Trans. Signal Process.*, vol. 55, no. 4, pp. 1511–1524, 2007.
- [9] H. Aghajan and A. Cavallaro, *Multi-Camera Networks: Principles and Applications*. Burlington, MA: Academic Press, 2009.
- [10] K.-D. Kim and P. R. Kumar, “Cyber physical systems: A perspective at the centennial,” *Proc. IEEE*, vol. 100, pp. 1287–1308, 2012.
- [11] S. Haykin, *Cognitive Dynamic Systems: Perception-Action Cycle, Radar and Radio*. New York, NY: Cambridge University Press, 2012.
- [12] A. D. Ryan, H. Durrant-Whyte, and J. K. Hedrick, “Information-theoretic sensor motion control for distributed estimation,” in *Proc. IMECE '07*, Seattle, WA, Nov. 2007.
- [13] G. M. Hoffmann and C. J. Tomlin, “Mobile sensor network control using mutual information methods and particle filters,” *IEEE Trans. Autom. Control*, vol. 55, no. 1, pp. 32–47, 2010.
- [14] M. Schwager, P. Dames, D. Rus, and V. Kumar, “A multi-robot control policy for information gathering in the presence of unknown hazards,” in *Proc. ISRR-11*, Flagstaff, AZ, Aug. 2011.
- [15] B. J. Julian, M. Angermann, M. Schwager, and D. Rus, “Distributed robotic sensor networks: An information-theoretic approach,” *Int. J. Robot. Res.*, vol. 31, pp. 1134–1154, Sep. 2012.
- [16] N. A. Atanasov, J. L. Ny, and G. J. Pappas, “Distributed algorithms for stochastic source seeking with mobile robot networks,” *J. Dyn. Sys., Meas., Control*, vol. 137, Mar. 2015.
- [17] H. Wymeersch, J. Lien, and M. Z. Win, “Cooperative localization in wireless networks,” *Proc. IEEE*, vol. 97, pp. 427–450, Feb. 2009.
- [18] T. Sathyan and M. Hedley, “Fast and accurate cooperative tracking in wireless networks,” *IEEE Trans. Mobile Comput.*, vol. 12, no. 9, pp. 1801–1813, 2013.
- [19] Y.-C. Wu, Q. M. Chaudhari, and E. Serpedin, “Clock synchronization of wireless sensor networks,” *IEEE Signal Process. Mag.*, vol. 28, Jan. 2011.
- [20] 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.
- [21] B. Etzlinger, F. Meyer, A. Springer, F. Hlawatsch, and H. Wymeersch, “Cooperative simultaneous localization and synchronization: A distributed hybrid message passing algorithm,” in *Proc. Asilomar Conf. Sig., Syst., Comput.*, Pacific Grove, CA, pp. 1978–1982, Nov. 2013.
- [22] B. Etzlinger, H. Wymeersch, and A. Springer, “Cooperative synchronization in wireless networks,” *IEEE Trans. Signal Process.*, vol. 62, pp. 2837–2849, Jun. 2014.
- [23] 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, no. 5, pp. 2305–2318, 2012.

- [24] 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.
- [25] F. Meyer, O. Hlinka, H. Wymeersch, E. Riegler, and F. Hlawatsch, "Cooperative simultaneous localization and tracking in mobile agent networks," *IEEE Trans. Signal Process.*, 2014, submitted. Available online: <http://arxiv.org/abs/1403.1824>.
- [26] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proc. IEEE*, vol. 95, no. 1, pp. 215–233, 2007.
- [27] 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.
- [28] 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.
- [29] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York, NY: Springer, 2006.
- [30] W. Burgard, D. Fox, and S. Thrun, "Active mobile robot localization by entropy minimization," in *Proc. EUROBOT '97*, Brescia, Italy, pp. 155–162, Oct. 1997.
- [31] B. Grocholsky, *Information-Theoretic Control of Multiple Sensor Platforms*. PhD thesis, University of Sydney, Sydney, Australia, 2002.
- [32] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY: Wiley, 2006.
- [33] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Upper Saddle River, NJ: Prentice-Hall, 1993.
- [34] F. Morbidì and G. L. Mariottini, "Active target tracking and cooperative localization for teams of aerial vehicles," *IEEE Trans. Control Syst. Technol.*, vol. 21, no. 5, pp. 1694–1707, 2013.
- [35] 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.
- [36] Y. Bar-Shalom, X.-R. Li, and T. Kirubarajan, *Estimation with applications to tracking and navigation*. New York, NY: Wiley, 2001.
- [37] 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.
- [38] 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.
- [39] 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.
- [40] 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.
- [41] C. Lindberg, L. S. Muppirisetty, K.-M. Dahlen, V. Savic, and H. Wymeersch, "MAC delay in belief consensus for distributed tracking," in *Proc. IEEE WPNC-13*, Dresden, Germany, Mar. 2013.
- [42] R. Fletcher, *Practical Methods of Optimization*. New York, NY: Wiley, 1987.
- [43] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*. New York, NY: Springer, 2001.
- [44] H. Lim and C. Kim, "Multicast tree construction and flooding in wireless ad hoc networks," in *Proc. MSWIM '00*, New York, NY, pp. 61–68, Aug. 2000.
- [45] L. Xiao, S. Boyd, and S. Lall, "A scheme for robust distributed sensor fusion based on average consensus," in *Proc. IPSN '05*, Los Angeles, CA, pp. 63–70, Apr. 2005.
- [46] R. G. Bartle, *The Elements of Integration and Lebesgue Measure*. New York, NY: Wiley, 1995.
- [47] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge, MA: MIT Press, 2005.
- [48] 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.
- [49] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Norwood, MA: Artech House, 2004.
- [50] G. E. Garcia, L. S. Muppirisetty, E. M. Schiller, and H. Wymeersch, "On the trade-off between accuracy and delay in cooperative UWB localization: Performance bounds and scaling laws," *IEEE Trans. Wireless Commun.*, vol. 13, pp. 4574–4585, Aug. 2014.
- [51] R. Douc and O. Cappé, "Comparison of resampling schemes for particle filtering," in *Proc. ISPA-05*, Zagreb, Croatia, pp. 64–69, Sep. 2005.
- [52] B. W. Silverman and P. J. Green, *Density Estimation for Statistics and Data Analysis*. London, UK: Chapman and Hall, 1986.
- [53] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scaokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, pp. 789–814, Jun. 2000.
- [54] D. L. Alspach and H. W. Sorenson, "Nonlinear Bayesian estimation using Gaussian sum approximations," *IEEE Trans. Autom. Control*, vol. 17, pp. 439–448, Aug. 1972.
- [55] I. Arasaratnam and S. Haykin, "Cubature Kalman filters," *IEEE Trans. Autom. Control*, vol. 54, pp. 1254–1269, Jun. 2009.
- [56] Y. Deville and A. Deville, "Exact and approximate quantum independent component analysis for qubit uncoupling," in *Proc. LVA/ICA '12*, Tel Aviv, Israel, pp. 58–65, Mar. 2012.



Florian Meyer (S'12) received the Dipl.-Ing. (M.Sc.) and Ph.D. degrees in electrical engineering from Vienna University of Technology, Vienna, Austria in 2011 and 2015, respectively. Since 2011, he has been a Research and Teaching Assistant with the Institute of Telecommunications, Vienna University of Technology, Vienna, Austria. He was a visiting scholar at the Department of Signals and Systems, Chalmers University of Technology, Sweden in 2013 and at the NATO STO Centre of Maritime Research and Experimentation (CMRE), La Spezia, Italy in 2014. His research interests include signal processing for wireless sensor networks, localization and tracking, information-seeking control, message passing algorithms, and finite set statistics.



Henk Wymeersch (S'99, M'05) received the Ph.D. degree in Electrical Engineering/Applied Sciences in 2005 from Ghent University, Belgium. He is currently an Associate Professor with the Department of Signals and Systems at Chalmers University of Technology, Sweden. Prior to joining Chalmers, he was a Postdoctoral Associate with the Laboratory for Information and Decision Systems (LIDS) at the Massachusetts Institute of Technology (MIT). He served as Associate Editor for IEEE COMMUNICATION LETTERS (2009–2013), IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS (2013–present), and TRANSACTIONS ON EMERGING TELECOMMUNICATIONS TECHNOLOGIES (2011–present).



Markus Fröhle received the B.Sc. degree in 2009 and the M.Sc. degree in 2012 in Telematics, both from Graz University of Technology, Austria. He is currently pursuing the Ph.D. degree in electrical engineering at the Department of Signals and Systems, Chalmers University of Technology, Sweden. From 2012 to 2013, he was a Research Assistant with the Signal Processing and Speech Communication Laboratory, Graz University of Technology, Austria. His research interests include indoor positioning and navigation, UWB, and the interplay between estimation theory, communications, and control in multi-agent systems.



Franz Hlawatsch (S'85–M'88–SM'00–F'12) received the Diplom-Ingenieur, Dr. techn., and Univ.-Dozent (habilitation) degrees in electrical engineering/signal processing from Vienna University of Technology, Vienna, Austria in 1983, 1988, and 1996, respectively. Since 1983, he has been with the Institute of Telecommunications, Vienna University of Technology, where he is currently an Associate Professor. During 1991–1992, as a recipient of an Erwin Schrödinger Fellowship, he spent a sabbatical year with the Department of Electrical Engineering, University of Rhode Island, Kingston, RI, USA. In 1999, 2000, and 2001, he held one-month Visiting Professor positions with INP/ENSEEIH, Toulouse, France and IRCCyN, Nantes, France. He (co)authored a book, three review papers that appeared in the IEEE SIGNAL PROCESSING MAGAZINE, about 200 refereed scientific papers and book chapters, and three patents. He coedited three books. His research interests include statistical and compressive signal processing methods and their application to sensor networks and wireless communications.

Prof. Hlawatsch was Technical Program Co-Chair of EUSIPCO 2004 and served on the technical committees of numerous IEEE conferences. He was an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING from 2003 to 2007 and for the IEEE TRANSACTIONS ON INFORMATION THEORY from 2008 to 2011. From 2004 to 2009, he was a member of the IEEE SPCOM Technical Committee. He currently serves as an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL AND INFORMATION PROCESSING OVER NETWORKS. He coauthored papers that won an IEEE Signal Processing Society Young Author Best Paper Award and a Best Student Paper Award at IEEE ICASSP 2011.