

Distributed Map Merging with Consensus on Common Information

Rosario Aragues^a, Jorge Cortes^b, and Carlos Sagues^c

Abstract—Sensor fusion methods combine noisy measurements of *common* variables observed by several sensors, typically by averaging information matrices and vectors of the measurements. Some sensors may have also observed *exclusive* variables on their own. Examples include robots exploring different areas or cameras observing different parts of the scene in map merging or multi-target tracking scenarios. Iteratively averaging exclusive information is not efficient, since only one sensor provides the data, and the remaining ones echo this information. This paper proposes a method to average the information matrices and vectors associated only to the common variables. Sensors use this averaged common information to locally estimate the exclusive variables. Our estimates are equivalent to the ones obtained by averaging the complete information matrices and vectors. The proposed method preserves properties of convergence, unbiased mean, and consistency, and improves the memory, communication, and computation costs.

I. INTRODUCTION

Sensor networks are a powerful technology for monitoring, surveillance, and detection of events in the environment. Sensor fusion methods estimate variables by combining noisy measurements taken by individual units. Measurements are expressed in Information Filter form (IF), and their information matrices and vectors are averaged using distributed consensus to produce the global information matrix and vector, see e.g., [6], [24]. The global estimate in terms of mean vectors and covariance matrices can then be obtained by inverting the averaged information matrix. When the variables to be estimated change along time, an additional step is locally carried out at each sensor to predict the new state of the variables [1], [7], [8], [19]. Each particular sensor may have not observed the whole set of variables. This situation has been studied in [16] in terms of the observability of the system. Several multi-target tracking algorithms may fail when there are exclusive variables, and thus specific solutions have been provided to deal with these features [14]. However, during the sensor measurement averaging phases, no specific separation between common and exclusive variables is performed, and exclusive information is exchanged and averaged at each iteration. Average consensus is an interesting choice when sensors have different information

about the same variables. However, iteratively averaging exclusive information is not efficient. A single sensor provides information for the exclusive elements, whereas the remaining sensors just echo this information.




Fig. 1. Robots exploring a region, with common areas in blue.

In map merging, multi-target tracking, or surveillance scenarios, sensors have a description of what a map feature, interesting target, or intruder are (visual features, doors, furniture, or moving bodies). Each sensor individually detects these variables, and before fusing the sensorial data, sensors must establish relationships between their variables. This problem, known as data association, has been investigated in the context of distributed map merging [2] and distributed target tracking [14], [21]. Matches are established between the variables of *neighboring* sensors, using e.g. the Joint Compatibility Branch and Bound [18], Nearest Neighbor, Combined Constraint Data Association [5], Iterative Closest Point [9], or RANSAC [11]. After that, exclusive variables are identified without requiring any extra efforts: they are variables that have not been associated to any other one. For applications that do not require data association, where sensors a priori know which variables will be estimated, it would be necessary to run a voting method in the network to find out which variables are exclusive.

Common and exclusive features have been previously discussed in the context of robot exploration and submapping [20]. In multi-robot scenarios, it is interesting to send robots to explore different regions. Thus map merging scenarios are an example of situations with high amount of exclusive features (Fig. 1). Although some consensus-based approaches have been proposed [2], [3], most of the map merging solutions rely on centralized schemes, all-to-all communication, or broadcasting methods, e.g., [13] for particle filters, [23] for multi-robot submaps, or [22] for graph maps. Since these solutions are not robust to link failures or changes in the communication topology, some decentralized solutions have been proposed where robots propagate and keep track of the measurements [17] or latest local maps [10] of all the team members. Their memory cost does not scale properly, and it increases as new robots join

This work was supported by projects Labex IMobS3, DPI2009-08126, and DPI2012-32100, grant MEC BES-2007-14772 and NSF award CCF-0917166.

^aR. Aragues (corresponding author) is with the Institut Pascal, BP 10448, F-63000 Clermont-Ferrand France, raragues@unizar.es.

^bJ. Cortes is with the Department of Mechanical and Aerospace Engineering, University of California San Diego, CA 92093 USA, cortes@ucsd.edu.

^cC. Sagues is with Instituto de Investigación en Ingeniería de Aragón, Universidad de Zaragoza, 50018 Zaragoza Spain, csagues@unizar.es.

the team, even if the scene size does not change. Although propagation is reasonable for variables exclusively sensed by a robot, consensus methods are a better choice for estimating common features observed by several robots.

The contribution of this paper is the design and analysis of a novel sensor fusion strategy that relies on executing consensus only on the information matrices and vectors of the common variables. Our method allows the robots to better distribute the computations, in particular in what refers to matrix inversion operations, and greatly improves the communication usage. We show that our method leads to the same solution as algorithms that do not make this distinction. We study the convergence, consistency, and unbiasedness of the algorithm, and illustrate its behavior in a map merging scenario. Due to space limitations, all proofs are omitted and will appear elsewhere.

II. PRELIMINARIES

We let n be the number of sensors, and i, j , be indices of sensors. We use G for the global map, and e and c for exclusive and common variables respectively. We use $t \in \mathbb{N}$ for iteration numbers. We let \mathbf{I}_d be the $d \times d$ identity matrix, and $\mathbf{0}_{d_1 \times d_2}$ be a $d_1 \times d_2$ matrix with all its elements equal to zero. Given A, B , matrices, $A \preceq B$ ($A \prec B$) means that matrix $B - A$ is positive semidefinite (positive definite).

Sensor fusion methods [6], [24] estimate some variables $\mathbf{x} \in \mathbb{R}^{\mathcal{M}_G}$ using measurements acquired by n sensors $i \in \{1, \dots, n\}$ during successive measurement steps $k \in \mathbb{N}$. In this paper we focus on the static case, which takes place when all sensors take a measurement at a single step, e.g., at $k = 0$. Some of the variables \mathbf{x} are common, $\mathbf{x}^c \in \mathbb{R}^{\mathcal{M}_G^c}$, whereas others have been exclusively observed by a single sensor, $\mathbf{x}_i^e \in \mathbb{R}^{\mathcal{M}_i^e}$, with $\mathcal{M}_G = \mathcal{M}_1^e + \dots + \mathcal{M}_n^e + \mathcal{M}_G^c$,

$$\mathbf{x} = ((\mathbf{x}_1^e)^T, \dots, (\mathbf{x}_n^e)^T, (\mathbf{x}^c)^T)^T.$$

Let matrix $H_i \in \mathbb{R}^{\mathcal{M}_i \times \mathcal{M}_G}$ relate the locally observed variables and the global ones. The observation $\mathbf{y}_i \in \mathbb{R}^{\mathcal{M}_i}$ of sensor i is a noisy measurement of the true variables \mathbf{x} ,

$$\mathbf{y}_i = H_i \mathbf{x} + \mathbf{v}_i, \quad \mathbf{v}_i \sim N(\mathbf{0}, \Sigma_i), \quad (1)$$

where $\Sigma_i \in \mathbb{R}^{\mathcal{M}_i \times \mathcal{M}_i}$ is the covariance of observation noise \mathbf{v}_i , and the noises independent, $E[\mathbf{v}_i \mathbf{v}_j^T] = \mathbf{0}$, for all $i \neq j$. Measurement \mathbf{y}_i can be expressed in terms of its exclusive $\mathbf{y}_i^e \in \mathbb{R}^{\mathcal{M}_i^e}$ and common $\mathbf{y}_i^c \in \mathbb{R}^{\mathcal{M}_i^c}$ parts, with $\mathcal{M}_i = \mathcal{M}_i^e + \mathcal{M}_i^c$, and eq. (1) becomes

$$\begin{aligned} \mathbf{y}_i &= ((\mathbf{y}_i^e)^T, (\mathbf{y}_i^c)^T)^T, \\ \mathbf{y}_i^e &= \mathbf{x}_i^e + \mathbf{v}_i^e, \quad \mathbf{y}_i^c = H_i^c \mathbf{x}^c + \mathbf{v}_i^c, \end{aligned} \quad (2)$$

since only sensor i provides information of \mathbf{x}_i^e and thus

$$H_i = \begin{bmatrix} \mathbf{0} \dots \mathbf{I} & \dots \mathbf{0} & \mathbf{0} \\ \mathbf{0} \dots \mathbf{0} & \dots \mathbf{0} & H_i^c \end{bmatrix}. \quad (3)$$

Here, we observe the decomposition of the covariance matrix Σ_i associated to the noise $\mathbf{v}_i = ((\mathbf{v}_i^e)^T, (\mathbf{v}_i^c)^T)^T$,

$$\Sigma_i = \begin{bmatrix} E[\mathbf{v}_i^e (\mathbf{v}_i^e)^T] & E[\mathbf{v}_i^e (\mathbf{v}_i^c)^T] \\ E[\mathbf{v}_i^c (\mathbf{v}_i^e)^T] & E[\mathbf{v}_i^c (\mathbf{v}_i^c)^T] \end{bmatrix} = \begin{bmatrix} M_i & N_i \\ N_i^T & O_i \end{bmatrix}. \quad (4)$$

The goal is to build the global estimate $\mathbf{x}_G \in \mathbb{R}^{\mathcal{M}_G}$, which is composed of exclusive $\mathbf{d}_i \in \mathbb{R}^{\mathcal{M}_i^e}$ and common features $\mathbf{f} \in \mathbb{R}^{\mathcal{M}_G^c}$, with covariance matrix $\Sigma_G \in \mathbb{R}^{\mathcal{M}_G \times \mathcal{M}_G}$,

$$\mathbf{x}_G = \begin{bmatrix} \mathbf{d}_1 \\ \vdots \\ \mathbf{d}_n \\ \mathbf{f} \end{bmatrix}, \quad \Sigma_G = \begin{bmatrix} D_{11} & \dots & D_{1n} & E_1 \\ \vdots & \ddots & \vdots & \vdots \\ D_{n1} & \dots & D_{nn} & E_n \\ E_1^T & \dots & E_n^T & F \end{bmatrix}, \quad (5)$$

where D_{ij} are the covariances of the exclusive variables of sensors i and j , F is the covariance of the common parts, and matrices E_i relate the exclusive and common variables. In particular, it is interesting that the nodes compute \mathbf{x}_G , and the elements in the main diagonal of the covariance matrix Σ_G . Usually, these elements provide enough information for any higher level algorithm to make decisions. The unbiased, maximum likelihood estimate \mathbf{x}_G of \mathbf{x} given the measurements \mathbf{y}_i acquired by all the sensors $i \in \{1, \dots, n\}$, with covariance matrix Σ_G , is

$$\begin{aligned} \mathbf{x}_G &= I_G^{-1} \mathbf{i}_G, & \Sigma_G &= I_G^{-1}, \\ I_G &= \sum_{i=1}^n I_i, & \mathbf{i}_G &= \sum_{i=1}^n \mathbf{i}_i, \end{aligned} \quad (6)$$

where I_i , \mathbf{i}_i , are the information matrix and vector of the measurement taken by robot i ,

$$I_i = H_i^T \Sigma_i^{-1} H_i, \quad \mathbf{i}_i = H_i^T \Sigma_i^{-1} \mathbf{y}_i. \quad (7)$$

For independent, zero-mean, *non-Gaussian* measurement noises, eq. (6) gives the linear minimum-variance unbiased estimate of \mathbf{x} given the measurements.

A. Distributed sensor fusion

The global estimate in eq. (6) is computed with distributed average consensus on the information matrices $I_i \in \mathbb{R}^{\mathcal{M}_G \times \mathcal{M}_G}$, and vectors $\mathbf{i}_i \in \mathbb{R}^{\mathcal{M}_G}$ of the measurements [6], [24]. Each sensor i maintains and exchanges variables $\hat{I}_i(t) \in \mathbb{R}^{\mathcal{M}_G \times \mathcal{M}_G}$, $\hat{\mathbf{i}}_i(t) \in \mathbb{R}^{\mathcal{M}_G}$ with an estimate of the averages of the local information matrices and vectors in eq. (7) for all the robots $j \in \{1, \dots, n\}$, and updates these variables at each iteration t through the exchange of data with its neighbors. Three properties which are of interest in sensor fusion algorithms are convergence to the global estimate in eq. (6), and unbiasedness and consistency of the temporal estimates. Convergence refers to the fact that the averaged estimates $\hat{I}_i(t)$, $\hat{\mathbf{i}}_i(t)$ at each robot $i \in \{1, \dots, n\}$, and averaging iteration t , converge to the global estimate I_G , \mathbf{i}_G , as $t \rightarrow \infty$. In particular, for fixed graphs or switching graphs jointly connected [24],

$$\begin{aligned} \lim_{t \rightarrow \infty} \hat{I}_i(t) &= I_G/n, & \lim_{t \rightarrow \infty} \hat{\mathbf{i}}_i(t) &= \mathbf{i}_G/n, \\ \lim_{t \rightarrow \infty} \hat{\mathbf{x}}_i(t) &= \lim_{t \rightarrow \infty} (\hat{I}_i(t))^{-1} \hat{\mathbf{i}}_i(t) = \mathbf{x}_G, \\ \lim_{t \rightarrow \infty} \hat{\Sigma}_i(t) &= \lim_{t \rightarrow \infty} (\hat{I}_i(t))^{-1} = n \Sigma_G, \end{aligned}$$

with \mathbf{x}_G , Σ_G , I_G , \mathbf{i}_G as in eq. (6). The temporal estimates $\hat{I}_i(t)$, $\hat{\mathbf{i}}_i(t)$ are unbiased when

$$E[\hat{\mathbf{x}}_i(t)] = E[(\hat{I}_i(t))^{-1} \hat{\mathbf{i}}_i(t)] = \mathbf{x},$$

for all i and t , and they are consistent when

$$E [(\hat{\mathbf{x}}_i(t) - \mathbf{x})(\hat{\mathbf{x}}_i(t) - \mathbf{x})^T] \preceq \hat{\Sigma}_i(t).$$

Unbiasedness has been proved for the sensor fusion algorithms [6], [24], and consistency has been proved for [6] and for the consensus-based Kalman filter method in [8].

III. DISTRIBUTED CONSENSUS ON COMMON INFORMATION

The strategies in the previous section involve the exchange of information matrices and vectors, with sizes respectively quadratic and linear in $\mathcal{M}_G = \mathcal{M}_1^c + \dots + \mathcal{M}_n^c + \mathcal{M}_G^c$, for several averaging iterations. Instead, in this paper we propose to compute average consensus of the information matrices and vectors associated to the measurements of the *common* variables, $R_i \in \mathbb{R}^{\mathcal{M}_G^c \times \mathcal{M}_G^c}$, $\mathbf{r}_i \in \mathbb{R}^{\mathcal{M}_G^c}$, given by

$$R_i = (H_i^c)^T O_i^{-1} H_i^c, \quad \mathbf{r}_i = (H_i^c)^T O_i^{-1} \mathbf{y}_i^c. \quad (8)$$

We propose to make each sensor i maintain variables $\hat{\mathbf{r}}_i(t) \in \mathbb{R}^{\mathcal{M}_G^c}$, $\hat{R}_i(t) \in \mathbb{R}^{\mathcal{M}_G^c \times \mathcal{M}_G^c}$, with an estimate of the average of R_i , \mathbf{r}_i , and update them with [6], [24], so that

$$\lim_{t \rightarrow \infty} \hat{R}_i(t) = \frac{1}{n} \sum_{j=1}^n R_j, \quad \lim_{t \rightarrow \infty} \hat{\mathbf{r}}_i(t) = \frac{1}{n} \sum_{j=1}^n \mathbf{r}_j. \quad (9)$$

As in [6], [24], we assume that $\hat{R}_i(t)$, $\hat{\mathbf{r}}_i(t)$ can be expressed in terms of R_j , \mathbf{r}_j ,

$$\hat{R}_i(t) = \sum_{j=1}^n \phi_{ij}(t) R_j, \quad \hat{\mathbf{r}}_i(t) = \sum_{j=1}^n \phi_{ij}(t) \mathbf{r}_j, \quad (10)$$

with $\phi_{ij}(t) \in \mathbb{R}$; as $t \rightarrow \infty$, these weights $\phi_{ij}(t) \rightarrow 1/n$. Thus, in our approach sensors exchange matrices and vectors quadratic and linear in the size of the common elements, \mathcal{M}_G^c , instead of \mathcal{M}_G .

We first analyze how the different elements of the global estimate in eq. (5) are obtained from the measurements taken by the robots $i \in \{1, \dots, n\}$ in eqs. (2)-(4). We define the following matrices $\bar{M}_i \in \mathbb{R}^{\mathcal{M}_i^c \times \mathcal{M}_i^c}$, $\bar{N}_i \in \mathbb{R}^{\mathcal{M}_i^c \times \mathcal{M}_G^c}$, and vector $\bar{\mathbf{y}}_i \in \mathbb{R}^{\mathcal{M}_i^c}$,

$$\begin{aligned} \bar{M}_i &= M_i - N_i O_i^{-1} N_i^T, & \bar{N}_i &= N_i O_i^{-1} H_i^c, \\ \bar{\mathbf{y}}_i &= \mathbf{y}_i^c - N_i O_i^{-1} \mathbf{y}_i^c, \end{aligned} \quad (11)$$

which can be locally computed by each node $i \in \{1, \dots, n\}$.

Proposition 3.1 (Exclusive/Common Decomposition):

The exclusive and common parts of the global estimate \mathbf{x}_G , Σ_G in eq. (5) can be expressed in terms of the exclusive and common measurements \mathbf{y}_i , Σ_i in eqs. (2)-(4) as follows:

$$F = \left(\sum_{i=1}^n R_i \right)^{-1}, \quad D_{ij} = \bar{N}_i F \bar{N}_j^T, \quad E_i = \bar{N}_i F, \quad (12)$$

$$D_{ii} = \bar{M}_i + \bar{N}_i F \bar{N}_i, \quad \mathbf{f} = F \sum_{i=1}^n \mathbf{r}_i, \quad \mathbf{d}_i = \bar{\mathbf{y}}_i + \bar{N}_i \mathbf{f}.$$

with R_i , \mathbf{r}_i being the common information matrices and vectors in eq. (8), and \bar{M}_i , \bar{N}_i , $\bar{\mathbf{y}}_i$ being the locally computed matrices and vectors in eq. (11). •

Note that the terms in eq. (12) are not an approximation; instead, they are the equivalent expression to the classical estimation operations in eq. (5). The interest is that the sizes of R_i , \mathbf{r}_i only depend on the common features observed by the robots, and not on the total amount of global features $\mathcal{M}_G = \mathcal{M}_1^c + \dots + \mathcal{M}_n^c + \mathcal{M}_G^c$. We propose to perform the averaging consensus iterations on these common information matrices R_i and vectors \mathbf{r}_i , instead of on the information matrices I_i and vectors \mathbf{i}_i of the whole set of global variables. Then, the robots can compute the remaining elements in eq. (12) from F and \mathbf{f} . Besides, the only matrix inversion operations (eqs. (8) and (11)-(12)) involve O_i and R_i with sizes quadratic on \mathcal{M}_i^c and \mathcal{M}_G^c respectively. Recall that in the original sensor fusion algorithms, matrices Σ_i and I_G are inverted instead (eqs. (6)-(7)), with sizes quadratic on \mathcal{M}_i and \mathcal{M}_G .

Algorithm 3.1 (Consensus on common information): We propose to make each robot i execute an average consensus with variables $\hat{\mathbf{r}}_i(t)$, $\hat{R}_i(t)$, initialized with R_i , \mathbf{r}_i in eq. (8), using the consensus algorithms in Section II. Given $\hat{\mathbf{r}}_i(t)$, $\hat{R}_i(t)$, at robot i and iteration t , we define the following estimates of the variables associated to the global covariance matrix and vector in eqs. (5), and (12),

$$\begin{aligned} \hat{F}_i(t) &= (\hat{R}_i(t))^{-1}, & \hat{\mathbf{f}}_i(t) &= (\hat{R}_i(t))^{-1} \hat{\mathbf{r}}_i(t), \\ \hat{\mathbf{d}}_j^i(t) &= \bar{\mathbf{y}}_j + \bar{N}_j \hat{\mathbf{f}}_i(t), & \hat{D}_{jj}^i(t) &= n \bar{M}_j + \bar{N}_j \hat{F}_i(t) \bar{N}_j^T, \end{aligned} \quad (13)$$

for all $j \in \{1, \dots, n\}$, and given the estimates of two robots i, i' , we define for all $j, j' \in \{1, \dots, n\}$

$$\begin{aligned} \hat{D}_{jj'}^{ii'}(t) &= (\hat{D}_{jj'}^i(t) + \hat{D}_{jj'}^{i'}(t))/2, & \tilde{D}_{jj'}^i(t) &= \bar{N}_j \hat{F}_i(t) \bar{N}_{j'}^T, \\ \hat{E}_{jj'}^{ii'}(t) &= (\hat{E}_{jj'}^i(t) + \hat{E}_{jj'}^{i'}(t))/2, & \tilde{E}_j^i(t) &= \bar{N}_j \hat{F}_i(t). \end{aligned} \quad (14)$$

We are interested in each robot i computing the variables in eq. (13) since they usually provide enough information about the environment for any higher level algorithm to make decisions. If necessary, however, also variables in eq. (14) can be obtained by the robots. Both strategies are discussed in the next section. •

In the remaining of this paper, we study the properties of the global map estimated by each robot i and iteration t . The exclusive variables $\hat{\mathbf{d}}_j^{ii'}(t)$, $\hat{D}_{jj'}^{ii'}(t)$ in eq. (13) may have been computed by a robot i' not necessarily equal to robot i (using the averaged estimates $\hat{\mathbf{r}}_{i'}(t)$, $\hat{R}_{i'}(t)$ of robot i'), and propagated through the network. Variables in eq. (14) may have involved two different robots. In order to cover all the possible computation/propagation strategies, we give the following general definition of the global estimate of a robot i . The global mean $\hat{\mathbf{x}}_i(t)$ estimated according to our algorithm by each robot i at iteration t is,

$$\hat{\mathbf{x}}_i(t) = \left((\hat{\mathbf{d}}_1^{i1}(t))^T, \dots, (\hat{\mathbf{d}}_n^{in}(t))^T, (\hat{\mathbf{f}}_i(t))^T \right)^T, \quad (15)$$

and its associated global covariance matrix $\hat{\Sigma}_i(t)$ is

$$\hat{\Sigma}_i(t) = \begin{bmatrix} \hat{D}_{11}^{i_1}(t) & \hat{D}_{12}^{i_1 i_2}(t) & \dots & \hat{D}_{1n}^{i_1 i_n}(t) & \hat{E}_1^{i_1 i}(t) \\ \hat{D}_{21}^{i_2 i_1}(t) & & & & \\ \vdots & & \ddots & \vdots & \vdots \\ \hat{D}_{n1}^{i_n i_1}(t) & \dots & \hat{D}_{nn}^{i_n}(t) & & \hat{E}_n^{i_n i}(t) \\ (\hat{E}_1^{i_1 i}(t))^T & \dots & (\hat{E}_n^{i_n i}(t))^T & & \hat{F}_i(t) \end{bmatrix}, \quad (16)$$

which is symmetric since $\hat{D}_{jj'}^{i_i'}(t) = (\hat{D}_{j'j}^{i_i'}(t))^T$. Note that depending on the scenario considered, robot i may not estimate the full covariance matrix, but only the common part $\hat{F}_i(t)$ and the main diagonal elements $\hat{D}_{jj}^{i_i}$. Now we discuss the different scenarios and the specific global mean $\hat{\mathbf{x}}_i(t)$ and covariance $\hat{\Sigma}_i(t)$ obtained in each case.

A. Strategies for estimating the global map

Each robot i executing Algorithm 3.1 has an estimate of the common features $\hat{\mathbf{f}}_i(t)$ and covariance $\hat{F}_i(t)$. In addition, it can compute an estimate of its exclusive features $\hat{\mathbf{d}}_i^i(t)$, and covariance matrix $\hat{D}_{ii}^i(t)$, using only its own local data $\bar{\mathbf{y}}_i, \bar{M}_i, \bar{N}_i$ in eq. (11), and $\hat{\mathbf{f}}_i(t)$ and $\hat{F}_i(t)$. In some scenarios, it may be interesting to let each robot i have only this information, i.e., the improved estimates of its local features due to the sensor fusion operation. In this case, the global estimate would be spread along the robot network. Instead, if we are interested in each robot having a copy of the global estimate, we can use the following strategies.

1) *Strategy A*: Robots exchange periodically messages (flooding or propagation) of linear size $2M_i^e$ with their estimates of their exclusive features $\hat{\mathbf{d}}_i^i(t)$, and the main diagonal of their covariances matrices $\hat{D}_{ii}^i(t)$. Usually, these elements provide enough information for any higher level algorithm to make decisions. The global mean $\hat{\mathbf{x}}_i(t)$ and covariance $\hat{\Sigma}_i(t)$ estimated by robot i at iteration t as in eqs. (15), (16), with $i_1 = 1, \dots, i_n = n$. This is the most important strategy and it is the one we implement in our simulations.

2) *Strategy B*: Alternatively, if robots were required to estimate *all the entries* of the covariance and mean of the global map, they could use the strategy described here. Each robot j propagates (once, e.g., at the beginning) vector $\bar{\mathbf{y}}_j \in \mathbb{R}^{M_j^e}$; matrix $\bar{N}_j \in \mathbb{R}^{M_j^e \times M_j^e}$; and matrix $\bar{M}_j \in \mathbb{R}^{M_j^e \times M_j^e}$ in eq. (11). Once robot i receives these matrices and vectors, it computes variables in eqs. (13)-(16) using its own $\hat{F}_i(t), \hat{\mathbf{f}}_i(t)$. The global mean $\hat{\mathbf{x}}_i(t)$ and covariance $\hat{\Sigma}_i(t)$ estimated by robot i at iteration t would be in this case in as eqs. (15), (16) with $i_1 = \dots = i_n = i$.

Remark 3.1 (Global map mean and covariance): In this paper, we consider the case where the final aim is that the robots compute the mean and covariance of the global map. We plan to investigate in future work the alternative case, where agents seek to obtain the global information matrices and vectors instead.

Remark 3.2 (Performance comparison): Regardless of the strategy selected, the computation load of our algorithm is lower than in classical sensor fusion algorithms, where

each robot must invert the whole information matrix of the global estimate to obtain the global map. Here, the sizes of the matrices $\hat{R}_i(t)$ which are inverted depend on the common features only. In the worst case, when all the robots observe exactly the same features, the performance is the same. The improvement of our method over classical sensor fusion solutions becomes more important in scenarios with a high amount of exclusive variables.

B. Algorithm properties

Here, we analyze the convergence, unbiasedness and consistency of Algorithm 3.1. The proof of the next result follows from (9) and Proposition 3.1.

Proposition 3.2 (Convergence): Assume eq. (9) holds. Then, for all i, i', j, j' , the following variables in eqs. (13)-(14) converge to

$$\begin{aligned} \lim_{t \rightarrow \infty} \hat{F}_i(t) &= nF, & \lim_{t \rightarrow \infty} \hat{D}_{jj}^j(t) &= nD_{jj}, \\ \lim_{t \rightarrow \infty} \hat{D}_{jj'}^{i_i'}(t) &= nD_{jj'}, & \lim_{t \rightarrow \infty} \hat{E}_j^{i_i'}(t) &= nE_j, \\ \lim_{t \rightarrow \infty} \hat{\mathbf{f}}_i(t) &= \mathbf{f}, & \lim_{t \rightarrow \infty} \hat{\mathbf{d}}_j^j(t) &= \mathbf{d}_j, \end{aligned} \quad (17)$$

with $F, D_{jj}, D_{jj'}, E_j, \mathbf{f}, \mathbf{d}_j$ being the blocks of the global covariance matrix and the global mean vector as in eq. (12).

Proposition 3.3 (Unbiasedness): Assume that $\hat{R}_i(t), \hat{\mathbf{r}}_i(t)$, are as in eq. (10). Then, the global map mean $\hat{\mathbf{x}}_i(t)$ (eq. (15)) estimated by each robot i at each iteration t is unbiased,

$$E[\hat{\mathbf{x}}_i(t)] = \mathbf{x}. \quad (18)$$

Theorem 3.1 (Consistency): Assume that $\hat{R}_i(t), \hat{\mathbf{r}}_i(t)$ for all i and t can be expressed as in eq. (10), and that in addition the weights $0 \leq \phi_{ij}(t) \leq 1$. Then, the numerical covariance $Q_i(t)$ of the global mean $\hat{\mathbf{x}}_i(t)$ (eq. (15)) estimated by each robot i at each iteration t , defined as

$$Q_i(t) = E[(\hat{\mathbf{x}}_i(t) - \mathbf{x})(\hat{\mathbf{x}}_i(t) - \mathbf{x})^T], \quad (19)$$

is related to the covariance matrix $\hat{\Sigma}_i(t)$ in eq. (16) as follows,

$$Q_i(t) \preceq \hat{\Sigma}_i(t). \quad (20)$$

IV. SIMULATIONS

We apply the method proposed in this paper to a map merging scenario, where 7 robots explore an unknown environment (Fig. 2 (a)). Features inside rooms are only observed by robots that enter this room (Fig. 2 (b)). When robots finish their exploration, they merge their maps. Fig. 2 (c) shows the global map (6) that could be computed if all the local maps were available to a fusion center.

Robots execute the data association algorithm in [4] to establish a relationship between the features locally observed by themselves and the ones observed by the robot team. They identify exclusive features, which have not been associated

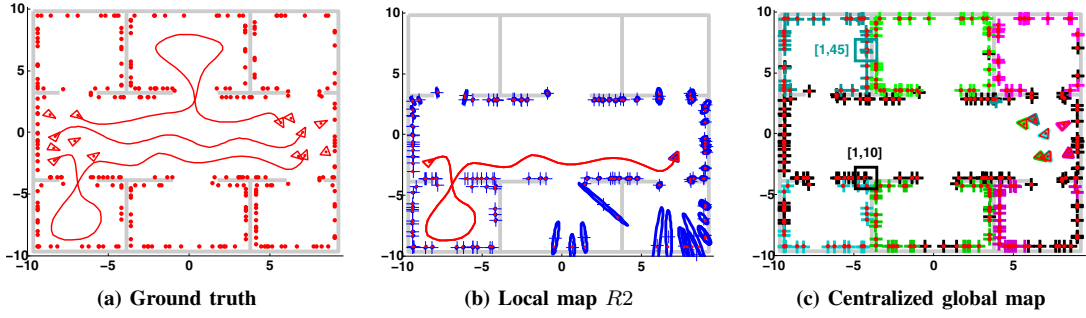


Fig. 2. 7 robots (triangles) explore an environment where there are some landmarks (dots). Robots perform different trajectories (lines) and traverse the main corridor and some rooms. They finish in the right part of the scene. Features inside rooms are observed by the robots that enter this room. (a) Ground truth robot poses (red triangles) and landmark positions (red dots). (b) Local map computed by robot R_2 . It has estimated (blue crosses and ellipses) the position of the features in the main corridor and in the rooms it has entered, as well as its own pose (blue triangle). (c) Centralized global map (eqs. (6),(12)). The common parts are shown in black, whereas the exclusive features observed by each robot are displayed in different colors. Two features ($F[1, 45]$ and $F[1, 10]$) that will be studied later are shown inside boxes.

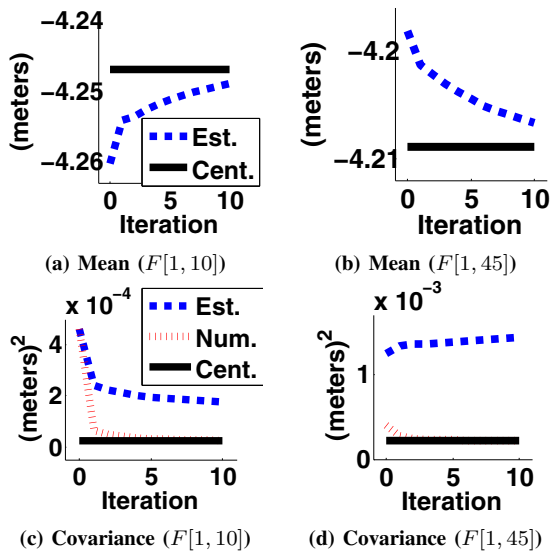


Fig. 3. Position of features $F[1, 10]$ and $F[1, 45]$ (x -coordinate) estimated by robot R_1 along $t = 10$ iterations of the algorithm. Feature $F[1, 10]$ has been observed by several robots (common), whereas $F[1, 45]$ has been only observed by R_1 . (a), (b) The estimated mean (blue dashed) approaches the centralized estimate (black solid). (c), (d) The numerical covariance (red dotted), which cannot be computed by the robots, approaches the centralized covariance (black solid); this numerical covariance (red dotted) is always bounded by the locally estimated covariance (blue dashed), which converges to n times the centralized covariance.

to any other feature. They execute the proposed sensor fusion method on their common features (Algorithm 3.1) for $t = 10$ iterations, using the strategy A (Section III-A). At each iteration t , each robot i uses its most recent estimates of the common parts $\hat{R}_i(t)$, $\hat{r}_i(t)$, to compute its exclusive features estimates $\hat{d}_i^e(t)$ and exclusive covariance matrix $\hat{D}_{ii}^e(t)$. Then, it propagates $\hat{d}_i^e(t)$ and the elements in the main diagonal of $\hat{D}_{ii}^e(t)$. In Fig. 3 we can see the evolution of the estimated mean (Fig. 3 (a)-(b)) and covariance (Fig. 3 (c)-(d)) of features $F[1, 10]$ and $F[1, 45]$ at robot R_1 . Feature $F[1, 10]$ was observed by several of the robots, whereas $F[1, 45]$ was exclusively observed by R_1 . For both features, the estimated means (Fig. 3 (a)-(b)) converge to the centralized estimate; the estimated covariances (Fig. 3 (c)-

(d)) always bound the numerical covariances (which cannot be computed by the robots) and converge to n times the centralized covariances.

We finally discuss the communication cost of our approach. We represent numbers with single precision (4 Bytes per element) and show the results in KBytes. In the first phase, robots execute a local data association in their neighborhood followed by a propagation of matches through the network. This operation has associated a communication cost of 123.34 KBytes per robot for sending their local maps (means and covariances) in the neighborhood; then the propagation of matches process takes 5 iterations with a total communication cost of 8.69 KBytes. After that, robots know which features are exclusive and execute the proposed algorithm, with the strategy A (Section III-A). They propagate their exclusive parts at each iteration t , and execute consensus on the common features. The communication costs per robot at each iteration t are shown in Fig. 4. Our method (Fig. 4 (a)) outperforms classical consensus methods (Fig. 4 (b)) due to the management of exclusive and common data. We have made a comparison against a strategy based on pure propagation of the local maps (Fig. 4 (c)). In our algorithm, the amount of data sent by every robot is almost the same, and thus the green dashed and gray solid lines are almost the same (Fig. 4 (a)), whereas for the pure propagation method (Fig. 4 (c)) some of the robots send large amounts of data (green dashed). Pure propagation methods are prone to create bottlenecks, which are even more critical for larger networks with more robots. The maximum memory used by the robots during the execution of our algorithm was 169 KBytes (105 KBytes for storing the common parts and 65 KBytes for storing the exclusive parts), whereas the classical consensus required each robot to store 454 KBytes. In order to obtain the estimate of the global map, one of the most computationally demanding operation is the inversion of the information matrix, whose size depends on the amount of common features for our method and on the total amount of landmarks for the classical consensus. The robots computed the global map estimate in 0.192 seconds when they executed the classical consensus, and only 0.017 seconds when they

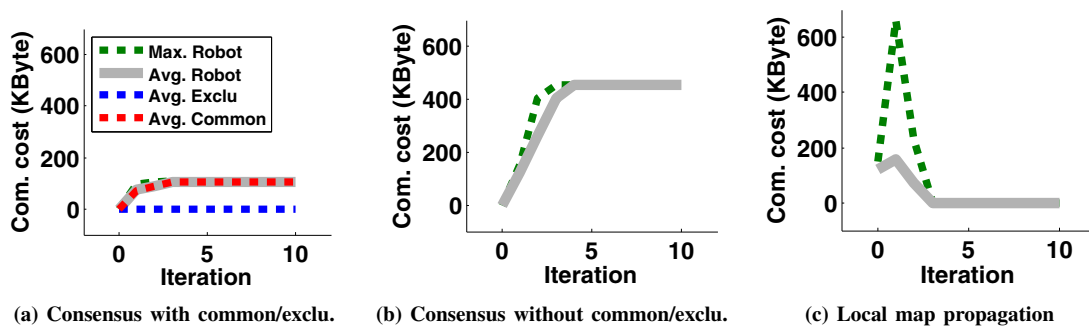


Fig. 4. Communication costs of our method (Algorithm 3.1 & Strategy A) (a) compared to consensus methods that do not make any distinction between common or exclusive features (b), and to a pure propagation of local maps (c), with numbers represented with single precision (4Bytes). The average amount of information sent by each robot is displayed in solid gray, and the maximum cost among all the robots is shown in green dashed. For our method (a), we display as well the amount of data used for sending the exclusive information (blue dashed) and for agreeing in the common data (red dashed).

used our method.

V. CONCLUSIONS

We have presented a method for sensor fusion that improves the network usage by taking advantage of the presence of exclusive features, which have been observed by a single agent. This algorithm converges to the same solution as methods that consider exclusive and common features together. We have shown that our method is convergent, unbiased, and consistent. We have applied our method to a map merging scenario, where it is usual to have a high number of exclusive features, and we have demonstrated the performance of our method compared to classical ones. Future work will explore more general distributed estimation scenarios, with time-varying variables that begin as exclusive and become common over time and also investigate the design of algorithms where robots compute the information matrices and vectors of the global estimate, instead of its covariance and mean.

REFERENCES

- [1] P. Alriksson and A. Rantzer. Distributed Kalman filtering using weighted averaging. In *Int. Symposium on Mathematical Theory of Networks and Systems*, Kyoto, Japan, July 2006.
- [2] R. Aragues, J. Cortes, and C. Sagües. Distributed consensus algorithms for merging feature-based maps with limited communication. *Robotics and Autonomous Systems*, 59(3-4):163–180, 2011.
- [3] R. Aragues, J. Cortes, and C. Sagües. Distributed consensus on robot networks for dynamically merging feature-based maps. *IEEE Transactions on Robotics*, 28(4):840–854, 2012.
- [4] R. Aragues, E. Montijano, and C. Sagües. Consistent data association in multi-robot systems with limited communications. In *Robotics: Science and Systems*, pages 97–104, Zaragoza, Spain, June 2010.
- [5] T. Bailey, E. M. Nebot, J. K. Rosenblatt, and H. F. Durrant-Whyte. Data association for mobile robot navigation: a graph theoretic approach. In *IEEE Int. Conf. on Robotics and Automation*, pages 2512–2517, San Francisco, USA, April 2000.
- [6] G. Calafiore and F. Abrate. Distributed linear estimation over sensor networks. *International Journal of Control*, 82(5):868–882, 2009.
- [7] R. Carli, A. Chiuso, L. Schenato, and S. Zampieri. Distributed Kalman filtering based on consensus strategies. *IEEE Journal on Selected Areas in Communications*, 26(4):622–633, 2008.
- [8] D. W. Casbeer and R. Beard. Distributed information filtering using consensus filters. In *American Control Conference*, pages 1882–1887, St. Louis, USA, June 2009.
- [9] A. Censi. An accurate closed-form estimate of ICP’s covariance. In *IEEE Int. Conf. on Robotics and Automation*, pages 3167–3172, Roma, Italy, April 2007.
- [10] A. Cunningham, K. M. Wurm, W. Burgard, and F. Dellaert. Fully distributed scalable smoothing and mapping with robust multi-robot data association. In *IEEE Int. Conf. on Robotics and Automation*, pages 1093–1100, St. Paul, USA, May 2012.
- [11] M. A. Fischler and R. C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, pages 381–395, 1981.
- [12] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, UK, 1985.
- [13] A. Howard. Multi-robot simultaneous localization and mapping using particle filters. *International Journal of Robotics Research*, 25(12):1243–1256, 2006.
- [14] A. T. Kamal, C. Ding, B. Song, J. A. Farrell, and A. K. Roy-Chowdhury. A generalized Kalman consensus filter for wide-area video networks. In *IEEE Conf. on Decision and Control*, pages 7863–7869, Orlando, FL, USA, December 2011.
- [15] A. T. Kamal, J. A. Farrell, and A. K. Roy-Chowdhury. Information weighted consensus. In *IEEE Conf. on Decision and Control*, Maui, Hawaii, USA, December 2012. To appear.
- [16] U. A. Khan and A. Jadbabaie. Coordinated network estimation strategies using structured systems theory. In *IEEE Conf. on Decision and Control*, pages 2112–2117, Orlando, FL, USA, December 2011.
- [17] K. Y. K. Leung, T. D. Barfoot, and H. H. T. Liu. Decentralized cooperative simultaneous localization and mapping for dynamic and sparse robot networks. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pages 3554–3561, Taipei, Taiwan, October 2010.
- [18] J. Neira and J. D. Tardós. Data association in stochastic mapping using the joint compatibility test. *IEEE Transactions on Robotics and Automation*, 17(6):890–897, 2001.
- [19] R. Olfati-Saber. Distributed Kalman filtering for sensor networks. In *IEEE Conf. on Decision and Control*, pages 5492–5498, New Orleans, LA, December 2007.
- [20] P. Pinies and J. D. Tardos. Large-scale SLAM building conditionally independent local maps: Application to monocular vision. *IEEE Transactions on Robotics*, 24(5):1094–1106, 2008.
- [21] N. F. Sandell and R. Olfati-Saber. Distributed data association for multi-target tracking in sensor networks. In *IEEE Conf. on Decision and Control*, pages 1085–1090, Cancun, Mexico, December 2008.
- [22] R. Vincent, D. Fox, J. Ko, K. Konolige, B. Limketkai, B. Morisset, C. Ortiz, D. Schulz, and B. Stewart. Distributed multirobot exploration, mapping, and task allocation. *Annals of Mathematics and Artificial Intelligence*, 52(1):229–255, 2008.
- [23] S. B. Williams and H. Durrant-Whyte. Towards multi-vehicle simultaneous localisation and mapping. In *IEEE Int. Conf. on Robotics and Automation*, pages 2743–2748, Washington, DC, USA, May 2002.
- [24] L. Xiao, S. Boyd, and S. Lall. A scheme for robust distributed sensor fusion based on average consensus. In *Symposium on Information Processing of Sensor Networks (IPSN)*, pages 63–70, Los Angeles, CA, April 2005.