2009 American Control Conference
Hyatt Regency Riverfront, St. Louis, MO, USA
June 10-12, 2009

Distributed Information Filtering using Consensus Filters

David W. Casbeer and Randy Beard

Abstract— **In this paper we present a new information consensus filter for distributed dynamic-state estimation. Estimation is handled by the traditional information filter, while communication of measurements is handled by a consensus filter. First and second-order statistics of local estimates are discussed. It is shown that local information consensus filter estimates are unbiased, and the actual variance of the local estimation errors is comparable to a centralized estimate. However, local agents believe their local estimates are less accurate.**

I. Introduction

There has been a recent surge of interest in distributed sensor networks. This is due, in part, to cheaper processing power, memory, communication, and sensing capability. This paper will look at the problem of tracking a dynamic-state with a distributed sensor network.

Recently, Xiao et al. introduced a distributed method, for use in sensor networks, to estimate a static state using a weighted least-squares approximation [1]. It is a novel approach in that consensus filters [2], [3] are used in the information space to communicate the information matrix and information state in a distributed fashion. This solution, however, is not easily ported over to the dynamic case because of scaling issues from the consensus filters. One proposed solution to the dynamic problem is presented in Refs. [4]–[6], where a consensus filter is run directly on the estimator state space variables. Because of this fact, it is called the *Kalman Consensus Filter*. Another similar approach decomposes the system into smaller *overlapping subsystems* that are then combined through a consensus communication scheme [7].

One question with these dynamic distributed estimators concerns the statistical properties of the local estimates. The optimality of the Kalman filter is well known [8]. However, in distributed implementations, there is a correlation between local estimates (or tracks) [9]. In general distributed networks, it is not possible to exactly determine this correlation [10], [11]. Because of this unknown correlation, optimal fusion of neighboring information into a local estimate is not possible, which results in non-optimal local tracks. This is not to say that these methods [6], [7] are futile. Just as the extended Kalman filter has proved itself to be a reliable approximation, these methods have been shown, in limited settings, to work very well.

D. Casbeer is a PhD candidate with the Dept. of Electrical and Computer Engineering, Brigham Young University, Provo, UT (corresponding author email: casbeer@byu.net)

R. Beard is with the faculty of the Dept. of Electrical and Computer Engineering, Brigham Young University, Provo, UT (email: beard@byu.edu)

This research was supported by the National Science Foundation under Information Technology Research Grant CCF-0428004.

Other techniques to accomplish distributed data fusion for dynamic systems that rely on the inverse covariance filter or information filter have been around for many years [12], [13]. These methods do not use consensus filters, and they run into the same issue of correlation between local agent estimates. There have been attempts to resolve the track-to-track correlation problem in this setting [10], [11], [14].

In this paper, a novel *information consensus filter* (ICF) filter is presented that applies consensus filters to an information filter. Information filters and consensus filters are discussed in Section II. The new ICF is presented in Section III. A statistical analysis of the local tracks is given in Section IV, which is followed by some explanatory simulations in Section V. Unfortunately, we do not solve the problem of data fusion with unknown track-to-track correlation. We do, however, give insight into the statistical effects of consensus filters in dynamic-state estimation, including the effects of track-to-track correlation. We show the ICF estimate is unbiased and the actual variance of the estimate is comparable to a centralized estimate. However, agents believe their respective local estimates are less accurate.

II. Two Decentralized Filters

Two decentralized filters are now discussed: the probabilistic information filter in Section II-A and the deterministic consensus filter in Section II-B.

A. Information Filter

Assume each agent in the distributed sensor network monitors the dynamic-state $\mathbf{x}_k$, whose evolution in time is given by the discrete time model $\mathbf{x}_k = \mathbf{F}_k \mathbf{x}_{k-1} + \mathbf{w}_k$, where $\mathbf{F}_k$ is the state transition matrix and $\mathbf{w}_k$ is zero mean white Gaussian noise with covariance matrix $\mathbf{Q}_k$. Agent i's measurement at time k follows the model $\mathbf{z}_{i,k} = \mathbf{H}_{i,k}\mathbf{x}_k + \mathbf{v}_{i,k}$, where $\mathbf{H}_{i,k}$ is agent i's measurement model and $\mathbf{v}_{i,k}$ is zero mean white Gaussian noise with covariance $\mathbf{R}_{i,k}$.

In a distributed information filter, each node monitors the state $\mathbf{x}_k$ by maintaining a local information filter. Agent i's local estimate is parameterized in information space using the *information vector* $\hat{\mathbf{y}}_{i,k} \triangleq \mathbf{P}_{i,k}^{-1}\hat{\mathbf{x}}_{i,k}$ and *information matrix* $\mathbf{Y}_{i,k} \triangleq \mathbf{P}_{i,k}^{-1}$, where $\hat{\mathbf{x}}_{i,k}$ and $\mathbf{P}_{i,k}$ are, respectively, node i's estimate and covariance of the estimate error [15]. The prediction step on the i^{th} agent is given by

$$\mathbf{Y}_{i,k|k-1} = \mathbf{M}_{i,k} - \mathbf{M}_{i,k}\boldsymbol{\Sigma}_{i,k}^{-1}\mathbf{M}_{i,k} \quad (1)$$

$$\hat{\mathbf{y}}_{i,k|k-1} = \mathbf{Y}_{i,k|k-1}\mathbf{F}_k \mathbf{Y}_{i,k-1|k-1}^{-1}\hat{\mathbf{y}}_{i,k-1|k-1}, \quad (2)$$

where

$$\mathbf{\Sigma}_{i,k} = \mathbf{M}_{i,k} + \mathbf{Q}_k^{-1},$$
$$\mathbf{M}_{i,k} = \mathbf{F}_k^{-T}\mathbf{Y}_{i,k-1|k-1}\mathbf{F}_k^{-1}.$$

New measurements are fused into the local information state estimate additively as

$$\hat{\mathbf{y}}_{i,k|k} = \hat{\mathbf{y}}_{i,k|k-1} + \mathbf{i}_k \quad (3)$$
$$\mathbf{Y}_{i,k|k} = \mathbf{Y}_{i,k|k-1} + \mathbf{I}_k, \quad (4)$$

where the new information from the observations is

$$\mathbf{i}_k \triangleq \sum_{i=1}^{n} \mathbf{H}_{i,k}^T \mathbf{R}_{i,k}^{-1} \mathbf{z}_{i,k} \quad \text{and} \quad \mathbf{I}_k \triangleq \sum_{i=1}^{n} \mathbf{H}_{i,k}^T \mathbf{R}_{i,k}^{-1} \mathbf{H}_{i,k}. \quad (5)$$

Because the observation update is additive, the information filter lends itself to distributed filtering. Any node running a local information filter needs to collect the newly observed information $\mathbf{i}_{i,k} \triangleq \mathbf{H}_{i,k}^T \mathbf{R}_{i,k}^{-1} \mathbf{z}_{i,k}$ and $\mathbf{I}_{i,k} \triangleq \mathbf{H}_{i,k}^T \mathbf{R}_{i,k}^{-1} \mathbf{H}_{i,k}$ from every sensor and add this information to the local information state estimate. Notice, only the new information is communicated; knowledge of teammate sensor models is not needed, since the sensor model $\mathbf{H}_{i,k}$ is incorporated into the new information $\mathbf{i}_{i,k}$ and $\mathbf{I}_{i,k}$.

For a fully connected network, the distributed information filter is exactly equivalent to a centralized version (and hence centralized Kalman filter [15]). However, when the interaction topology is not fully connected, issues arise in communicating the observations, and the distributed version is no longer equivalent to the centralized version (*cf.* [10]–[12]). In this paper, we use consensus filters to address the issue of communicating new observations. Exactly how this is done is discussed in Sec. III.

B. Consensus Filter

Consensus algorithms are decentralized methods for a team of agents to agree on specific *consensus states*. In a consensus filter, each agent exchanges information with neighboring agents and not the entire team. Over time the agents reach an agreement (or consensus) concerning the consensus state [2], [3]. Furthermore, *average* consensus occurs when the final consensus state is the average of the initial values [16], [17].

Before presenting the consensus algorithm used in this paper, some graph theory terminology is needed. At any discrete-time instant τ, the communication topology between n agents can be described by the graph $G[\tau] = (V, E[\tau])$ where $V = \{1, \ldots, n\}$ is the vertex set and $E[\tau] \subseteq V \times V$ is the edge set. The ij^{th} element of the *adjacency matrix* $\mathbf{A}[\tau]$ of graph $G[\tau]$ is $A_{ij}[\tau] = 1$ if $i \neq j$ and the edge $(j, i) \in E[\tau]$, otherwise $A_{ij}[\tau] = 0$. From the adjacency matrix one can construct the graph's Laplacian matrix $\mathbf{L}[\tau]$ as

$$L_{ij}[\tau] = \begin{cases} -A_{ij}[\tau] & \text{if } i \neq j, \\ \sum_{j=1, j\neq i}^{n} A_{ij}[\tau] & \text{if } i = j\ . \end{cases} \quad (6)$$

In the consensus algorithm employed in this paper, each agent in the network maintains a local copy of the consensus state $\xi_i \in \mathbb{R}^m$. Each agent i updates ξ_i using its neighbors' consensus states according to the rule

$$\xi_i[\tau+1] = \xi_i[\tau] - \frac{1}{d_\tau} \sum_{i=1}^{n} A_{ij}[\tau](\xi_i[\tau] - \xi_j[\tau]) \quad (7)$$

where $d_\tau \in [d_{max,\tau}, \infty)$ and $d_{max,\tau}$ denotes the maximal degree of $G[\tau]$. The consensus protocol given in Eq. (7) was chosen by the need for discrete-time average consensus [17]. A nice distributed scheme to choose the weights d_τ is to use *Metropolis weights* [18].

After arranging the local information states into the vector $\boldsymbol{\xi}[\tau] = [\xi_1^T[\tau], \ldots, \xi_n^T[\tau]]^T$, the update can be written as

$$\boldsymbol{\xi}[\tau+1] = (\Psi[\tau] \otimes \mathbf{I})\boldsymbol{\xi}[\tau] \quad (8)$$

where,

$$\Psi[\tau] = \mathbf{I} - \frac{1}{d_\tau}\mathbf{L}[\tau]\ , \quad (9)$$

$\mathbf{I}$ is the appropriate size identity matrix, and $\otimes$ denotes the matrix Kronecker product. Note that $\Psi[\tau]$ defined in this manner is a stochastic matrix.

We desire to ensure average consensus. We know that (7) achieves average consensus when the following conditions hold [17]:

1) $G[\tau]$ $\tau = 1, \ldots, \infty$ are bidirectional graphs, and
2) there exists a $T \geq 0$ such that for every interval $[\tau, \tau + T]$ the union of the graphs over the time interval is connected.

Exactly why average consensus is desired will be discussed in the next section.

III. Information Consensus Filter

We now present the *information consensus filter* (ICF). The ICF uses consensus protocol (7) on the information state and information matrix in a decentralized information filter. The consensus filter addresses the issue of communicating new information throughout the network. The ICF is a distributed filter, where each agent maintains a local information filter. In this paper, we assume that communication and prediction updates are synchronized in the network.

There are three steps in the ICF: prediction, local measurement update, and consensus update. The first (not necessarily sequential) step of the ICF is the consensus update. Here, each agent updates its local information state and matrix using (7). The second step of the ICF is the measurement update, where each agent i fuses only the local observations $\mathbf{i}_{i,k}$ and $\mathbf{I}_{i,k}$. These are added to the local information state and matrix instead of the global information $\mathbf{i}_k$ and $\mathbf{I}_k$ (see (3) and (4)). The third step of the ICF is the local prediction step, and this step is exactly the information filter prediction step given in (1) and (2).

The local ICF is summarized in Algorithm 1. In this algorithm, τ is the time index for the consensus protocol, and $T_p \in \mathbb{Z}^+$ is the time interval between prediction updates. The time index τ is faster than k; one time step $k-1 \to k$ is equivalent to T_p time steps of the consensus time index $\tau \to \tau + T_p$. This notation allows unit time intervals in both τ and k making the following analysis less complicated.

Algorithm 1 Information Consensus Filter

Initialization (for node i):

$$\hat{\mathbf{y}}_i = \mathbf{y}[0] \qquad \mathbf{Y} = \mathbf{Y}[0]$$
$$\tau = 1 \qquad \tau_p = \tau + T_p$$

loop {Local iteration on node i}

1: Consensus Update

$$\hat{\mathbf{y}}_i \Leftarrow \hat{\mathbf{y}}_i - \frac{1}{d_k}\sum_{i=1}^{n} A_{ij}[\tau](\hat{\mathbf{y}}_i - \hat{\mathbf{y}}_j) \tag{10}$$

$$\mathbf{Y}_i \Leftarrow \mathbf{Y}_i - \frac{1}{d_k}\sum_{i=1}^{n} A_{ij}[\tau](\mathbf{Y}_i - \mathbf{Y}_j) \tag{11}$$

$$\tau \Leftarrow \tau + 1$$

2: **if** new observations are taken **then**

Measurement Update

$$\hat{\mathbf{y}}_i \Leftarrow \hat{\mathbf{y}}_i + \mathbf{i}_i \qquad \mathbf{Y}_i \Leftarrow \mathbf{Y}_i + \mathbf{I}_i \tag{12}$$

3: **if** time for a predication step (*i.e.,* $\tau = \tau_p$) **then**

Prediction Step

$$\mathbf{M}_i = \mathbf{F}_k^{-T}\mathbf{Y}_i\mathbf{F}_k^{-1} \quad \text{and} \quad \mathbf{Y}_{i,\text{tmp}} \Leftarrow \mathbf{Y}_i \tag{13}$$
$$\mathbf{Y}_i \Leftarrow \mathbf{M}_i - \mathbf{M}_i(\mathbf{M}_i + \mathbf{Q}_k^{-1})^{-1}\mathbf{M}_i \tag{14}$$
$$\hat{\mathbf{y}}_i \Leftarrow \mathbf{Y}_i\mathbf{F}_k\mathbf{Y}_{i,\text{tmp}}^{-1}\hat{\mathbf{y}}_i \tag{15}$$
$$\tau_p = \tau + T_p$$

end loop

IV. Accuracy of Fusion Process

Sensor fusion consisting of consensus filters applied directly to the information state and matrix will yield unbiased and conservative local estimates. By conservative we mean $\mathbf{Y}_i^{-1} \geq \bar{\mathbf{Y}}_i^{-1}$, that is, $\mathbf{Y}_i^{-1}$ less the true error covariance $\bar{\mathbf{Y}}_i^{-1} = E\left\{(\mathbf{Y}_i^{-1}\hat{\mathbf{y}}_i - \mathbf{x})(\mathbf{Y}_i^{-1}\hat{\mathbf{y}}_i - \mathbf{x})^{-T})\right\}$ is positive semi-definite.

To evaluate the first and second order statistics of local estimates, we will need the following Lemma:

Lemma 1. *Suppose a network of n agents, where each agent has the conservative and unbiased estimate of the state $\mathbf{x}$, parameterized in information space as $\hat{\mathbf{y}}_i$ and $\mathbf{Y}_i$ for $i = 1, \ldots, n$. The convex combination of the agents' local estimates, given by*

$$\hat{\mathbf{y}}_{(n)} \triangleq \omega_1\hat{\mathbf{y}}_1 + \cdots + \omega_n\hat{\mathbf{y}}_n, \tag{16}$$
$$\mathbf{Y}_{(n)} \triangleq \omega_1\mathbf{Y}_1 + \cdots + \omega_n\mathbf{Y}_n\ , \tag{17}$$

where $\sum_{j=1}^{n}\omega_i = 1$ and $\omega_i \in [0,1]$ yields the unbiased and conservative estimate $\hat{\mathbf{x}}_{(n)} = \mathbf{Y}_{(n)}^{-1}\hat{\mathbf{y}}_{(n)}$ with error covariance matrix estimate $\mathbf{P}_{(n)} = \mathbf{Y}_{(n)}^{-1}$ for any choice of ω_i such that $\sum_{j=1}^{n}\omega_i = 1$.

The proof of Lemma 1 is given in Ref. [19].

Assume the ICF has just completed Prediction Step 3, and the prior estimates at each node are unbiased, conservative, and equal, *e.g.,* $\hat{\mathbf{y}}_{j,k|k-1} = \hat{\mathbf{y}}_{i,k|k-1}$ and $\mathbf{Y}_{j,k|k-1} = \mathbf{Y}_{i,k|k-1}$ for all i, j. Because of this equality, the consensus update in Step 1 ((10) and (11)) has no effect. Now assume each node i makes the observation $\mathbf{i}_{i,k}$ (and $\mathbf{I}_{i,k}$) and locally fuses this observation yielding $\hat{\mathbf{y}}_{i,k|k}^{\tau_0}$ and $\mathbf{Y}_{i,k|k}^{\tau_0}$, for all $i = 1, \ldots, n$, where τ_0 indicates the time of the consensus filter when the observations are fused. The local estimates, which are not necessarily equivalent, are the best estimates given the local observation. Also, the local estimate $\hat{\mathbf{x}}_i^{\tau_0} = (\mathbf{Y}_i^{\tau_0})^{-1}\hat{\mathbf{y}}_i^{\tau_0}$ is not equivalent to a hypothetical centralized filter that fuses the observations from every node.

To evaluate the statistics of the local estimates, we look at three situations in the next three sections: 1) the consensus filters converge before the next prediction update, 2) the consensus filters converge after the next prediction update, and 3) the consensus filters do not converge. Each of the next three sections shows that the ICF produces unbiased and conservative local estimates. In each section, we compare exactly how conservative the local covariance matrix estimates are compared to that of a hypothetical centralized filter. This comparison with a centralized filter shows how much confidence is lost due to the consensus filters in the ICF.

A. Consensus Filters Converge Before Prediction

Node i has the estimate $\hat{\mathbf{y}}_{i,k|k}$ with its respective information matrix. We assume that the ICF now iterates enough so that the consensus filters converge before the next prediction step (*i.e.,* $T_p \gg 1$). Assuming the conditions for average consensus are satisfied, each agent would have the estimate

$$\hat{\mathbf{y}}_{i,k|k} = \hat{\mathbf{y}}_{i,k|k-1} + \frac{1}{n}\sum_{j=1}^{n}\mathbf{i}_{j,k} \tag{18}$$

$$\mathbf{Y}_{i,k|k} = \mathbf{Y}_{i,k|k-1} + \frac{1}{n}\sum_{j=1}^{n}\mathbf{I}_{j,k}\ , \tag{19}$$

where the covariance estimate of the local error is $\mathbf{Y}_{i,k|k}^{-1}$. By Lemma 1, the local estimate $\hat{\mathbf{x}}_{i,k|k} = \mathbf{Y}_{i,k|k}^{-1}\hat{\mathbf{y}}_{i,k|k}$ is unbiased and conservative. Notice that (19) differs from the hypothetical centralized IF (Eq. (4)) by the scale factor $\frac{1}{n}$. The ICF scales the new information by the inverse of the size of the network. Compared to a centralized filter fusing each agent's measurements, the ICF is "less confident" about the observations by an amount proportional to inverse of the network size.

B. Consensus Filters Converge After Prediction

Following the ICF measurement update (Step 3), each node has the local information state $\hat{\mathbf{y}}_{i,k|k} = \hat{\mathbf{y}}_{i,k|k-1} + \mathbf{i}_{i,k}$ and information matrix $\mathbf{Y}_{i,k|k} = \mathbf{Y}_{i,k|k-1} + \mathbf{I}_{i,k}$. Each node performs the prediction step (Step 4) of the ICF yielding the local information state $\hat{\mathbf{y}}_{i|k+1|k}$ and matrix $\mathbf{Y}_{i|k+1|k}$. We now assume that the ICF iterates sufficiently long for the consensus filters to converge before the next measurement update or prediction step. The converged information state and matrix at node i will be denoted respectively by

$$\hat{\mathbf{y}}^{c}_{i,k+1|k} = \frac{1}{n}\sum_{i=1}^{n}\hat{\mathbf{y}}_{i,k+1|k} \tag{20}$$

$$\mathbf{Y}^{c}_{i,k+1|k} = \frac{1}{n}\sum_{i=1}^{n}\mathbf{Y}_{i,k+1|k}\ , \tag{21}$$

where $\mathbf{Y}_{i,k+1|k} = \mathbf{M}_{i,k+1} - \mathbf{M}_{i,k+1}\Sigma^{-1}_{i,k+1}\mathbf{M}_{i,k+1}$ and $\hat{\mathbf{y}}_{i,k+1|k} = \mathbf{Y}_{i,k+1|k}\mathbf{F}_k\mathbf{Y}^{-1}_{i,k|k}\hat{\mathbf{y}}_{i,k|k}$. From Lemma 1, the local estimate, $\hat{\mathbf{x}}^{c}_{i,k+1|k} = (\mathbf{Y}^{c}_{i,k+1|k})^{-1}\hat{\mathbf{y}}^{c}_{i,k+1|k}$ with covariance estimate $(\mathbf{Y}^{c}_{i,k+1|k})^{-1}$, is unbiased and conservative, since it is the convex combination of n unbiased and conservative information states and matrices with coefficients $\omega_i = \frac{1}{n}$, for all $i = 1, \ldots, n$.

The loss in confidence of the ICF compared to a hypothetical centralized filter is less intuitive in this case due to the non-linear (information matrix) time update. The loss is given by $\mathbf{Y}_{loss} = \mathbf{Y}_{k+1|k} - \mathbf{Y}^{c}_{i,k+1|k}$ (where again $\mathbf{Y}^{c}_{i,k+1|k} = \mathbf{Y}^{c}_{j,k+1|k}$ for all i and j) which is

$$\begin{aligned}\mathbf{Y}_{loss} &= \mathbf{M}_{k+1} - \frac{1}{n}\sum_{i=1}^{n}\mathbf{M}_{i,k+1}\\ &- \mathbf{M}_{k+1}\Sigma^{-1}_{k+1}\mathbf{M}_{k+1} + \frac{1}{n}\sum_{i=1}^{n}\mathbf{M}_{i,k+1}\Sigma^{-1}_{i,k+1}\mathbf{M}_{i,k+1}.\end{aligned} \tag{22}$$

Assuming the centralized prior is equivalent to the local priors, $\mathbf{Y}_{k|k-1} = \mathbf{Y}_{i,k|k-1}$ for all $i = 1, \ldots, n$, then the first line in (22) shows that there is a loss from averaging that scales the new information by the inverse of the size of the network, since

$$\mathbf{M}_{k+1} = \mathbf{F}^{-T}_{k+1}(\mathbf{Y}_{k|k-1} + \mathbf{I}_k)\mathbf{F}^{-1}_{k+1} \tag{23}$$

$$\frac{1}{n}\sum_{i=1}^{n}\mathbf{M}_{i,k+1} = \mathbf{F}^{-T}_{k+1}(\mathbf{Y}_{k|k-1} + \frac{1}{n}\mathbf{I}_k)\mathbf{F}^{-1}_{k+1}. \tag{24}$$

The second line in (22) affects the loss from correlation incurred between local measurements as they pass through the common process model [9]. This correlation is not taken into account in the ICF.

C. Consensus Filters Do Not Converge

We now address the more practical situation where the consensus filters do not converge in between measurement updates. What can we say about the local intermediate estimates? In this scenario, we can show that the local estimates are again unbiased and conservative. Define, the transition matrix

$$\Phi[\tau, \tau_0] \triangleq \Psi[\tau]\cdots\Psi[\tau_0]\ .$$

Assuming no prediction step, at time τ_u, the information state and information matrix are, respectively,

$$\hat{\mathbf{y}}^{\tau_u}_{i,k|k} = \sum_{j=1}^{n}\Phi_{ij}[\tau_u, \tau_0]\hat{\mathbf{y}}^{\tau_0}_{j,k|k} \tag{25}$$

$$\mathbf{Y}^{\tau_u}_{i,k|k} = \sum_{j=1}^{n}\Phi_{ij}[\tau_u, \tau_0]\mathbf{Y}^{\tau_0}_{j,k|k} \tag{26}$$

where the superscripts on $\hat{\mathbf{y}}^{\tau_u}_{i,k|k}$ and $\mathbf{Y}^{\tau_u}_{i,k|k}$ indicate the consensus filter time index. Since $\Phi[\tau_u, \tau_0]$ is a stochastic matrix, we know from Lemma 1 that the local estimate is unbiased and conservative for any τ_u.

Break up node i's information matrix at time τ_u into two different parts: one part is the information it has in common with the rest of the network $\mathbf{Y}^{\tau_u}_{\cap,k|k}$, and the second part is everything else $\mathbf{Y}^{\tau_u}_{i\setminus\cap,k|k}$. We denote these divisions by

$$\mathbf{Y}^{\tau_u}_{\cap,k|k} = \mathbf{Y}^{\tau_0}_{i,k|k-1} + \sum_{j=1}^{n}\min\{\Phi_j\}\mathbf{I}_{j,k},$$

$$\mathbf{Y}^{\tau_u}_{i\setminus\cap,k|k} = \mathbf{Y}^{\tau_u}_{i,k|k} - \mathbf{Y}^{\tau_u}_{\cap,k|k}\ ,$$

where $\min\{\Phi_j\} = \min\{\Phi_{1j}[\tau_u, \tau_0], \ldots, \Phi_{nj}[\tau_u, \tau_0]\}$ and, again, we assume the local priors are equal, *i.e.,* $\mathbf{Y}^{\tau_0}_{i,k|k-1} = \mathbf{Y}^{\tau_0}_{j,k|k-1}$.

Each agent performs a prediction step, which is followed by a number of consensus filter iterations. This yields

$$\begin{aligned}\mathbf{Y}^{\tau}_{i,k+1|k} &= \sum_{j=1}^{n}\Phi_{ij}[\tau, \tau_u]\mathbf{M}^{\tau_u}_{j,k+1}\\ &\quad - \sum_{j=1}^{n}\Phi_{ij}[\tau, \tau_u]\mathbf{M}^{\tau_u}_{j,k+1}\left(\boldsymbol{\Sigma}^{\tau_u}_{j,k+1}\right)^{-1}\mathbf{M}^{\tau_u}_{j,k+1}\end{aligned} \tag{27}$$

where

$$\mathbf{M}^{\tau_u}_{i,k+1} = \mathbf{F}^{-T}_{k+1}\mathbf{Y}^{\tau_u}_{i,k|k}\mathbf{F}^{-1}_{k+1} \tag{28}$$

$$\boldsymbol{\Sigma}^{\tau_u}_{i,k+1} = \mathbf{M}^{\tau_u}_{i,k+1} + \mathbf{Q}^{-1}_{k}\ . \tag{29}$$

The first term on the right hand side of (27) is

$$\begin{aligned}&\sum_{j=1}^{n}\Phi_{ij}[\tau, \tau_u]\mathbf{M}^{\tau_u}_{i,k+1}\\ &\quad = \mathbf{F}^{-T}_{k+1}\left(\mathbf{Y}^{\tau_u}_{\cap,k|k} + \sum_{j=1}^{n}\Phi_{ij}[\tau, \tau_u]\mathbf{Y}^{\tau_u}_{j\setminus\cap,k|k}\right)\mathbf{F}^{-1}_{k+1}\ ,\end{aligned}$$

which is similar to the average in Eq. (24). Here, rather than an average, there is a weighted average on the uncommon information $\mathbf{Y}^{\tau_u}_{j\setminus\cap,k|k}$. This first term indicates that the ICF is conservative compared to the centralized filter by an amount determined by the *weighted average*.

The uncommon information $\mathbf{Y}^{\tau_u}_{j\setminus\cap,k|k}$ is not mutually independent for $j = 1, \ldots, n$, and it is not taken into account in the second term in (27):

$$\sum_{j=1}^{n}\Phi_{ij}[\tau, \tau_u]\mathbf{M}^{\tau_u}_{i,k+1}\left(\boldsymbol{\Sigma}^{\tau_u}_{i,k+1}\right)^{-1}\mathbf{M}^{\tau_u}_{i,k+1}\ . \tag{30}$$

Furthermore, there is an additional dependency incurred as these values pass through the common process model [20]. These two dependencies are not known in general decentralized networks and are one of the causes of the track-to-track problem discussed previously. These dependencies are not taken into account in the ICF. Similarly, the CI fusion method will be subject to this same problem. A similar phenomenon should occur with the KCF; however, since it is a state space method, the exact effect is not straightforward.

From the preceding work, it can be seen that the ICF filter estimates are always unbiased and conservative, assuming the priors are unbiased and conservative. Furthermore, the conservative nature of the ICF comes from averaging, and nodes in the ICF are "less-confident" about their local estimate by an amount governed by a weighted average, which, assuming the conditions for average consensus are satisfied (*cf.* Section II-B), approaches a true average asymptotically.

V. Simulations

A. Unbiased

We show that the ICF estimate is unbiased. Both a centralized (hypothetical) information filter and the ICF are examined. The filter estimate, itself, is a random variable because of the prior, the measurements, and the process model. We sample the filter 2000 times using $n = 30$. The prior estimate used to initialize the filters is $\mathbf{x} \sim \mathcal{N}([3,1]^T, \text{diag}(2.5,3))$ where the mean $[3,1]^T$ is the actual initial location. The process model is given by $\mathbf{F} = [1,1;0,1]$ with a covariance $\mathbf{Q} = .1[.5,1]^T[.5,1]$. The observations are one dimensional with a randomly generated model $\mathbf{H}$ and the measurement covariance is $\mathbf{R} = 1$ for every agent. For both examples, there is only one measurement step and one prediction step. The consensus filter in the ICF iterates 10 times prior to prediction and then 10 times after the prediction step (*i.e.,* $T_p = 10$). There is no measurement update after the prediction step.

In Figure 1, the gray dots each represent one realization of the filter. Figure 1a plots samples from the global (hypothetical) information filter. Figure 1b plots samples from the ICF. The solid black line in both of these figures is the 50 percent confidence region as indicated from the global information filter covariance $\mathbf{Y}^{-1}$ (IF). The dashed solid black line is the 50 percent confidence region indicated from one node's ICF covariance estimate $\mathbf{Y}_i^{-1}$.

Notice that the actual distribution of the ICF filter estimate is close to the global information filter distribution as seen from the spread of the dots. However, because of the consensus update, the local agents believe the covariance to be more conservative, as seen from the dashed black line being wider than the spread of dots. Remember, this discrepancy comes because each agent thinks the measurement covariance is scaled by the network size, and because there is unaccounted correlation between agent's information (*cf.* Section IV-C).

B. Information Loss

In Figure 2 we compare the actual differences in information for four scenarios: 1) the global (hypothetical) information filter (IF), 2) the ICF where the consensus filters converge prior to prediction (ICF 1), 3) the ICF where the consensus filters converge after prediction (ICF 2), and 4) the ICF with $T_p = 1$ (ICF). Since $T_p = 1$, there is one consensus iteration for every measurement and prediction step. The only differences in the model parameters, between here and those in the last simulation, is $\mathbf{R} = 10$ for each of the $n = 3$ agents. We plot the marginalized (1,1) element of agent 1's information matrix. The first point on each line is the prior (equal for all agents). The second point is after measurement update and one consensus update. The third point is after a prediction step. The last point is after another consensus update (this point is irrelevant for the IF).

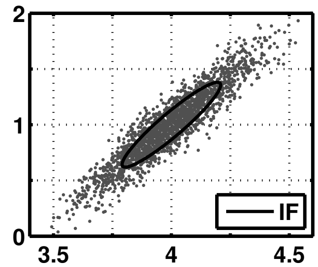

(a) Centralized IF

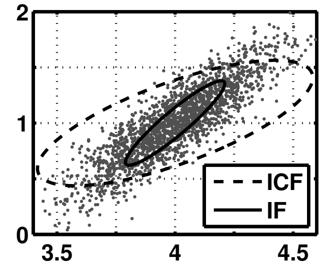

(b) ICF

Fig. 1: Here we sample the centralized IF and the ICF. Each gray dot is one sample. The black solid line is the 50% confidence region for the centralized IF. The dashed black line is this 50% confidence region as believed by a local agent.

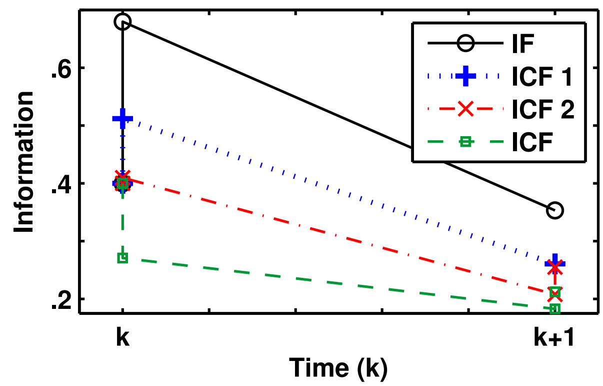

Fig. 2: Here we plot the information $\mathbf{Y}$ as it evolves through the ICF. IF is the centralized information filter. ICF 1 is the ICF filter where the consensus filter converge prior to prediction. The ICF 2 is where they converge after prediction. The ICF is the true filter where the consensus filters have not yet converged. It is plotted in k's time units.

The plot given is typical. The largest to smallest information as believed by the filters is the IF, ICF 1, ICF 2, and ICF. However, as we saw in the previous section we know that the covariance of the actual ICF estimate is closer to the IF covariance.

C. Complete system

We ran a simulation for $\tau = 200$ time steps with $T_p = 25$. The state observed is actually three dimensional, however, we only plot the first dimension in Figures 3 and 4. The black dots indicate the true state being tracked. Initially, the (true) prior mean is $[10,-3,5]^T$. Each agent's observations are scalar with a random measurement model. At every time step the consensus filter runs. The prediction step runs at every $\tau = \{25, 50, 75, \dots\}$, and the measurement update runs at every $\tau = \{26, 51, 76, \dots\}$. There are $n = 5$ agents with

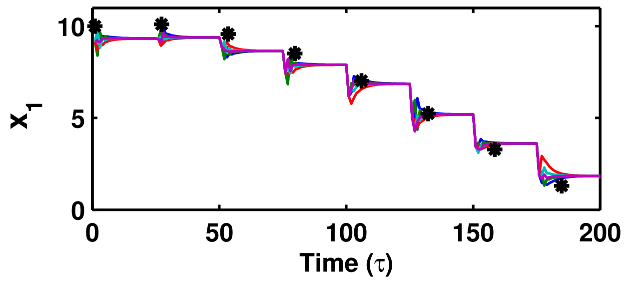

Fig. 3: The ICF in action. The black dot is the true state being estimated and each line represents one local agent's estimate. It is plotted in the time units of τ.

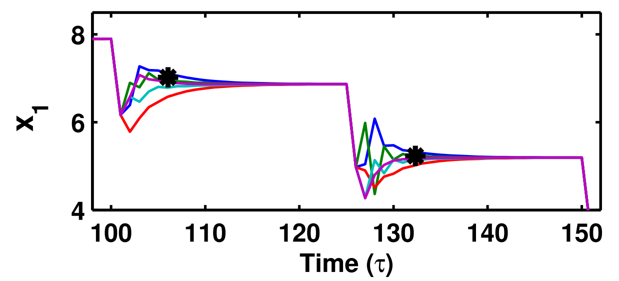

Fig. 4: Zoomed in look at Figure 3.

static bidirectional communication topology that is sparse yet connected. In Figures 3 and 4 there is one line for each agent's estimates; the black marks indicate the true state value.

In this scenario there is sufficient time for the consensus filters to converge prior to the next prediction step. If the network were larger this would most likely not be the case. The ICF is able to accurately track the state.

VI. Conclusions and Future Work

In this paper we have presented a novel distributed dynamic-state estimation method. We have presented a statistical analysis of this method. The method is unbiased and conservative. The actual covariance of local estimates in this filter is close to the centralized filter. However, the local agent's belief of this uncertainty is quite different from actual variance of the estimate.

One possible improvement to the ICF is finding a way to recover the centralized "hypothetical" estimate in a distributed manner. This actually is a problem in finding the true error covariance of the estimate and not the estimate itself, since the estimation is comparable to the centralized filter. If the network size were known, simply scaling the new information by the network size would cause the local estimates to be over-confident before the consensus filters converged. Because of the non-linear (in terms of the information matrix) prediction step, resolving the issue after the filters converge is not possible without extra information.

Also, there is a need for asynchronous local agents. Both asynchronous communication and prediction steps are needed. The asynchronous communication should easily be handled using the knowledge that consensus filters converge if the network topology is connected over a union of finite topologies [17], [21], [22]. To handle asynchronous prediction steps, the Kalman (information) filter delayed data association problem should be a good starting point [23], [24].

References

[1] L. Xiao, S. Boyd, and S. Lall, "A scheme for robust distributed sensor fusion based on average consensus," *Information Processing in Sensor Networks (IPSN)*, pp. 63–70, 2005.

[2] W. Ren, R. Beard, and E. Atkins, "Information consensus in multi-vehicle cooperative control," *Control Systems Magazine, IEEE*, vol. 27, no. 2, pp. 71–82, April 2007.

[3] R. Olfati-Saber, J. Fax, and R. Murray, "Consensus and cooperation in networked multi-agent systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, Jan. 2007.

[4] R. Olfati-Saber, "Distributed Kalman filter with embedded consensus filters," *Decision and Control and European Control Conference (CDC-ECC). 44th IEEE Conference on*, pp. 8179–8184, Dec. 2005.

[5] ——, "Distributed Kalman filtering for sensor networks," *Proceedings of the IEEE Conference on Decision and Control*, pp. 5492–5498, Dec. 2007.

[6] R. Olfati-Saber and N. Sandell, "Distributed tracking in sensor networks with limited sensing range," *Proceedings of the American Control Conference*, pp. 3157–3162, June 2008.

[7] S. Stankovic, M. Stankovic, and D. Stipanovic, "A consensus based overlapping decentralized estimator in lossy networks: Stability and denoising effects," *Proceedings of the American Control Conference*, pp. 4364–4369, June 2008.

[8] R. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering (ASME)*, vol. 82D, pp. 35–45, 1960.

[9] Y. Bar-Shalom, "On the track-to-track correlation problem," *IEEE Trans. Autom. Control*, vol. 26, no. 2, pp. 571–572, Apr 1981.

[10] S. Utete and H. Durrant-Whyte, "Reliability in decentralised data fusion networks," *Multisensor Fusion and Integration for Intelligent Systems, 1994. IEEE International Conference on MFI '94.*, pp. 215–221, Oct 1994.

[11] ——, "Routing for reliability in decentralised sensing networks," *American Control Conference, 1994*, vol. 2, pp. 2268–2272 vol.2, June-1 July 1994.

[12] S. Grime and H. Durrant-Whyte, "Data fusion in decentralized sensor networks," *Control Engineering Practice*, vol. 2, no. 5, pp. 849–863, 1994.

[13] A. G. O. Mutambara, *Decentralized Estimation and Control for Multisensor Systems*. Boca Raton, FL, USA: CRC Press, Inc., 1998.

[14] D. Nicholson, C. Lloyd, S. Julier, and J. Uhlmann, "Scalable distributed data fusion," *Proceedings of the International Conference on Information Fusion*, vol. 1, pp. 630–635 vol.1, 2002.

[15] P. S. Maybeck, *Stochastic Models, Estimation and Control*. Academic Press Inc., 1979, vol. 1.

[16] R. Olfati-Saber and R. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Trans. Autom. Control*, vol. 49, no. 9, pp. 1520–1533, Sept. 2004.

[17] D. Kingston and R. Beard, "Discrete-time average-consensus under switching network topologies," *Proceedings of the American Control Conference*, pp. 3551–3556, June 2006.

[18] S. Boyd, P. Diaconis, and L. Xiao, "Fastest Mixing Markov Chain on a Graph," *SIAM REVIEW*, vol. 46, no. 4, pp. 667–689, 2004.

[19] D. W. Casbeer, "Decentralized estimation in a multi-static UAV radar tracking system," Ph.D. dissertation, Brigham Young University, 2009.

[20] Y. Bar-Shalom, "On the track-to-track correlation problem," *IEEE Trans. Autom. Control*, vol. 26, no. 2, pp. 571–572, Apr 1981.

[21] L. Moreau, "Stability of multiagent systems with time-dependent communciation links," *IEEE Trans. Autom. Control*, vol. 50, no. 2, pp. 169–182, 2005.

[22] W. Ren and R. W. Beard, "Consensus seeking in multiagent systems under dynamically changing interaction topologies," *IEEE Trans. Autom. Control*, vol. 50, no. 5, pp. 655–661, May 2005.

[23] Y. Bar-Shalom and X. Li, *Multitarget-multisensor Tracking: Principles and Techniques*. Artech House, Norwood, 1995.

[24] E. Nettleton and H. Durrant-Whyte, "Delayed and asequent data in decentralised sensing networks," in *Sensor Fusion and Decentralised Control in Robotic Systems IV, Proc. SPIE*, vol. 4571, 2001, pp. 1–9.