

Distributing the Kalman Filter for Large-Scale Systems

Usman A. Khan, *Student Member, IEEE*, José M. F. Moura, *Fellow, IEEE*

Department of Electrical and Computer Engineering

Carnegie Mellon University

5000 Forbes Ave, Pittsburgh, PA 15213

{ukhan, moura}@ece.cmu.edu

Ph: (412)268-7103 Fax: (412)268-3890

Abstract

This paper derives a *distributed* Kalman filter to estimate a sparsely connected, large-scale, n -dimensional, dynamical system monitored by a network of N sensors. Local Kalman filters are implemented on the (n_l -dimensional, where $n_l \ll n$) sub-systems that are obtained after spatially decomposing the large-scale system. The resulting sub-systems overlap, which along with an assimilation procedure on the local Kalman filters, preserve an L th order Gauss-Markovian structure of the centralized error processes. The information loss due to the L th order Gauss-Markovian approximation is controllable as it can be characterized by a divergence that decreases as $L \uparrow$. The order of the approximation, L , leads to a lower bound on the dimension of the sub-systems, hence, providing a criterion for sub-system selection. The assimilation procedure is carried out on the local error covariances with a distributed iterate collapse inversion (DICI) algorithm that we introduce. The DICI algorithm computes the (approximated) centralized Riccati and Lyapunov equations iteratively with only local communication and low-order computation. We fuse the observations that are common among the local Kalman filters using bipartite fusion graphs and consensus averaging algorithms. The proposed algorithm achieves full distribution of the Kalman filter that is coherent with the centralized Kalman filter with an L th order Gaussian-Markovian structure on the centralized error processes. Nowhere storage, communication, or computation of n -dimensional vectors and matrices is needed; only $n_l \ll n$ dimensional vectors and matrices are communicated or used in the computation at the sensors.

Index Terms

Large-scale systems, sparse matrices, distributed algorithms, matrix inversion, Kalman filtering, distributed estimation, iterative methods

This work was partially supported by the DARPA DSO Advanced Computing and Mathematics Program Integrated Sensing and Processing (ISP) Initiative under ARO grant # DAAD 19-02-1-0180, by NSF under grants # ECS-0225449 and # CNS-0428404, and by an IBM Faculty Award.

I. INTRODUCTION

Centralized implementation of the Kalman filter [1], [2], although possibly optimal, does not provide robustness and scalability when it comes to complex large-scale dynamical systems with their measurements distributed on a large geographical region. The reasons are twofold: (i) the large-scale systems are very high-dimensional, and thus extensive computations are required to implement the centralized procedure; and (ii) the span of the geographical region, over which the large-scale system is deployed or the physical phenomenon is observed, poses a large communication burden and thus, among other problems, adds latency to the estimation mechanism. To remove the difficulties posed by the centralized procedure, we propose a distributed estimation algorithm with a low order Kalman filter at each sensor. To account for the processing, communication, and limited resources at the sensors, the local Kalman filters involve computations and communications with local quantities only, i.e., vectors and matrices of low dimensions, $n_l \ll n$, where n is the dimension of the state vector—no sensor computes, communicates, or stores any n -dimensional quantity.

Much of the existing research on distributed Kalman filters focuses on sensor networks monitoring low dimensional systems, e.g., when multiple sensors mounted on a small number of robot platforms are used for target tracking [3], [4], [5]. This scenario addresses the problem of how to efficiently incorporate the distributed observations, which is also referred to in the literature as ‘data fusion,’ see also [6]. Data fusion for Kalman filters over arbitrary communication networks is discussed in [7], using consensus protocols in [8]. References [9], [4] incorporate packet losses, intermittent observations, and communication delays in the data fusion process. Although these solutions work well for low-dimensional dynamical systems, e.g., target tracking applications, they implement a replication of the n -dimensional Kalman filter at each sensor, hence, communicating and inverting $n \times n$ matrices locally, which, in general, is an $O(n^3)$ operation. In contrast, in the problems we consider, the state dimension, n , is very large, for example, in the range of 10^2 to 10^9 , and so the n th order replication of the global dynamics in the local Kalman filters is either not practical or not possible.

Kalman filters with reduced order models have been studied, in e.g., [10], [11] to address the computation burden posed by implementing n th order models. In these works, the reduced models are decoupled, which is sub-optimal as important coupling among the system variables is ignored. Furthermore, the network topology is either fully connected [10], or is close to fully connected [11], requiring long distance communication that is expensive. We are motivated by problems where the large-scale systems, although sparse, cannot be decoupled, and where, due to the sensor constraints, the communication and computation should both be local. Multi-level systems theory to derive a coordination algorithm among local Kalman filters is discussed in [12].

We present a distributed Kalman filter that addresses *both* the computation and communication challenges posed by complex large-scale dynamical systems, while preserving its coupled structure; in particular, nowhere in our distributed Kalman filter, do we require storage, communication, or computation of n -dimensional quantities. We briefly explain the key steps and approximations in our solution.

Spatial Decomposition of Complex Large-Scale Systems: To distribute the Kalman filter, we provide a spatial

decomposition of the complex large-scale dynamical system (of dimension n) that we refer to as *the overall system* into several, possibly many, local coupled dynamical systems (of dimension n_l , such that $n_l \ll n$) that we refer to as *sub-systems* in the following. The large-scale systems, we consider, are sparse and localized. Physical systems with such characteristics are described in section II-A. We exploit the underlying distributed and sparse structure of system dynamics that results from a spatio-temporal discretization of random fields. Dynamical systems with such structure are then spatially decomposed into sub-systems. The sub-systems overlap and thus the resulting local Kalman filters also overlap. This overlap along with an assimilation procedure on the local error covariances (of the local filters) preserve the structure of the centralized (approximated) error covariances. We preserve the coupling by applying the coupled state variables as inputs to the sub-systems. In contrast, an estimation scheme using local observers is addressed in [13], where the coupled states are also applied as inputs to the sub-systems, but, the error covariances are not assimilated and remain decoupled. Such scheme loses its coherence with the centralized procedure, as under arbitrary coupling, no structure of the centralized error covariance can be retained by the local filters.

Overlapping Dynamics at the Sub-systems: Bipartite Fusion Graphs: The sub-systems that we extract from the overall system overlap. This makes several state variables to be observed at different sub-systems. To fuse this shared information, we implement a fusion algorithm using bipartite fusion graphs, which we introduced in [14], and local average consensus algorithms [15], [8]. The interactions required by the fusion procedure are constrained to a small neighborhood (this may require multi-hop communication in the small neighborhood). The multi-hop communication requirements, in our case, arise because the dynamical field is distributed among local sub-systems and only local processing is carried out throughout the development. In contrast, as noted earlier, a distributed Kalman filter with single hop communication scheme is presented in [7], but, we reemphasize the fact the this single-hop communications is a result of replicating n th order Kalman filters at each sensor.

Assimilation of the Local Error Covariances—Distributed Iterate-Collapse Inversion (DICI) Algorithm: A key issue in distributing the Kalman filter is to maintain a reasonable coherence with the centralized estimation; this is to say that the local error covariances should approximate the centralized error covariance in a well-defined manner. If the local error covariances evolve independently at each sub-system they may lose any coherence with the centralized error covariance under arbitrary coupling. To address this issue, we employ a *cooperative assimilation procedure* among the local error covariances that preserves a Gauss-Markovian structure in the centralized error processes. This is equivalent to approximating¹ the inverse of the centralized error covariances (the information matrices) to be L -banded², as shown in [16]. The assimilation procedure is carried out with a distributed iterate-collapse inversion (DICI, pronounced die-see) algorithm, briefly introduced in [17], that has an iterate step and a collapse step.

¹A measure on the loss of optimality is provided in section II-C and it will be shown that this measure serves as a criterion for choosing the dimensions of the sub-systems.

²We refer to a matrix as an L -banded matrix ($L \geq 0$), if the elements outside the band defined by the L th upper and L th lower diagonal are 0.

We use the Information filter, [18], [11], format of the Kalman filter where the information matrices (inverse of the error covariances) are iterated at each time step. We introduce the DICI algorithm that provides an iterative matrix inversion procedure to obtain the local error covariances from the local information matrices. The inverses of the local information matrices are assimilated among the sub-systems such that the local error covariances, obtained after the assimilation, preserve the Gauss-Markovian structure of the centralized error covariances. Iterative matrix inversion can also be carried out using the distributed Jacobi algorithm [19] where the computational complexity scales linearly with the dimension, n , of the overall system. In contrast, the computational complexity of the DICI algorithms is independent of n . In addition, the error process of the DICI algorithm is bounded above by the error process of the distributed Jacobi algorithm. We show the convergence of the iterate step of the DICI algorithm analytically and resort to numerical simulations to show the convergence of its collapse step.

In summary, the spatial decomposition of the complex large-scale systems, fusion algorithms for fusing observations, the DICI algorithm to assimilate the local error covariance combine to give a robust, scalable, and distributed implementation of the Kalman filter.

We describe the rest of the paper. Section II covers the discrete-time models, centralized Information filters, and centralized L -banded Information filters. Section III covers the model distribution step. We introduce the local Information filters in section III-B along with the necessary notation. Section IV gives the observation fusion step of the local Information filters, and section V presents the distributed iterate collapse inversion (DICI) algorithm. The filter step of the local Information filters is provided in section VI, and the prediction step of the local Information filters is provided in section VII. We conclude the paper with results in section VIII and conclusions in section IX. Appendix I discusses the L -banded inversion theorem, [20].

II. BACKGROUND

In this section, we motivate the type of applications and large-scale dynamical systems of interest to us. The context is that of a time-varying random field governed by partial differential equations (PDEs); these systems can also be generalized to arbitrary dynamical systems belonging to a particular structural class, as we elaborate in Subsection II-A. To fix notation, we then present the centralized version of the Information filter.

A. Global Model

Many physical phenomenon [21], [22], [23], [24], [25], [26], e.g., ocean/wind circulation and heat/wave equations, can be broadly characterized by a PDE of the Navier-Stokes type. These are highly non-linear and different regimens arise from different assumptions. Non-linear approximation models are commonly used for prediction (predictive models). For data assimilation, i.e., combining models with measured data, e.g., satellite altimetry data in ocean models, it is unfeasible to use non-linear models; rather linearized approximations (dynamical linearization) are employed. Here, our goal is to motivate how discrete linear models occur that exhibit a sparse and localized structure that we use to distribute the model in Section III. Hence, we take a very simplistic example and consider

elliptical operators, given by,

$$\mathcal{L} = \alpha \frac{\partial^2}{\partial \rho_x^2} + \beta \frac{\partial^2}{\partial \rho_y^2}, \quad (1)$$

where ρ_x and ρ_y represent the horizontal and vertical dimensions, respectively, and α, β are constants pertinent to the specific application. The spatial model for the continuous-time physical phenomenon (e.g., heat, wave, or, wind), x_t , with a continuous-time random noise term, u_t , can now be written as,

$$\mathcal{L}x_t = u_t, \quad (2)$$

which reduces to a Laplacian equation, $\Delta x_t = 0$ ($\Delta = \partial^2/\partial \rho_x^2 + \partial^2/\partial \rho_y^2$), or to a Poisson equation, $-\Delta x_t = u_t$, for appropriate choices of the constants. We spatially discretize the general model in (2) on an $M \times J$ uniform mesh grid, where the standard 2nd order difference equation approximation to the 2nd order derivative is

$$\frac{\partial^2 x_t}{\partial \rho_x^2} \sim x_{i+1,j} - 2x_{i,j} + x_{i-1,j}, \quad \frac{\partial^2 x_t}{\partial \rho_y^2} \sim x_{i,j+1} - 2x_{i,j} + x_{i,j-1}, \quad (3)$$

where x_{ij} is the value of the random field, x_t , at the ij -th location in the $M \times J$ grid. With this approximation, the Laplace equation (say), leads to the individual equations of the form

$$x_{i,j} - \frac{1}{4}(x_{i,j-1} + x_{i,j+1} + x_{i-1,j} + x_{i+1,j}) = 0. \quad (4)$$

Let \mathbf{I} be an identity matrix and let \mathbf{A} be a tridiagonal matrix with zeros on the main diagonal and ones on the upper and lower diagonals, then (4), $\forall i, j \in M \times J$ (with appropriate constants for modeling the elliptical operator in (2)), can be collected in the linear system of equations and (2) can be spatially discretized in the form

$$\mathbf{F}_c \mathbf{x}_t = \mathbf{b}_c + \mathbf{G}_c \mathbf{u}_t, \quad (5)$$

where the vector \mathbf{b}_c collects the boundary conditions, the term $\mathbf{G}_c \mathbf{u}_t$ controls where the random noise is added and

$$\mathbf{x}_t = \left[(\mathbf{x}_t^1)^T, \dots, (\mathbf{x}_t^M)^T \right]^T \quad (6)$$

and $\mathbf{x}_t^i = [x_{i1}, \dots, x_{iJ}]^T$. The matrix \mathbf{F}_c is given by

$$\mathbf{F}_c = \begin{bmatrix} \mathbf{B} & \mathbf{C} & & \\ \cdot & \cdot & \cdot & \\ & \mathbf{C} & \mathbf{B} & \end{bmatrix} = \mathbf{I} \otimes \mathbf{B} + \mathbf{A} \otimes \mathbf{C}, \quad (7)$$

where $\mathbf{B} = \mu \mathbf{I} + \beta_h \mathbf{A}$, and $\mathbf{C} = \beta_v \mathbf{I}$, the constants μ, β_h, β_v are in terms of α, β in (1), and \otimes is the Kronecker product [27].

If we further include time dependence, e.g., by discretizing the diffusion equation, we get

$$\dot{\mathbf{x}}_t = \mathbf{F}_c \mathbf{x}_t - \mathbf{b}_c - \mathbf{G}_c \mathbf{u}_t, \quad (8)$$

where $\dot{\mathbf{x}}_t = d\mathbf{x}_t/dt$. We again employ the standard first difference approximation of the first derivative with Δ_T being the sampling interval and k being the discrete-time index, and, write the above equation as

$$\mathbf{x}_{k+1} = (\mathbf{I} + \Delta_T \mathbf{F}_c) \mathbf{x}_k - \Delta_T \mathbf{b}_c - \Delta_T \mathbf{G}_c \mathbf{u}_k, \quad k \geq 0 \quad (9)$$

which can be generalized as the following discrete-time model (where for simplicity of the presentation and without loss of generality we drop the term \mathbf{b}_c in the sequel)

$$\mathbf{x}_{k+1} = \mathbf{F}\mathbf{x}_k + \mathbf{G}\mathbf{u}_k. \quad (10)$$

In the above model, $\mathbf{x}_k \in \mathbb{R}^n$ is the state vector, $\mathbf{x}_0 \in \mathbb{R}^n$ are the initial state conditions, $\mathbf{F} \in \mathbb{R}^{n \times n}$ is the model matrix, $\mathbf{u}_k \in \mathbb{R}^{j'}$ is the state noise vector and $\mathbf{G} \in \mathbb{R}^{n \times j'}$ is the state noise matrix.

Here, we note that the model matrix, \mathbf{F} , is perfectly banded in case of PDEs as in (7). We can relax this to sparse and localized, i.e., the coupling among the states decays with distance (in an appropriate measure). Besides random fields, the discrete-space-time model (10) with sparse and localized structure also occurs, e.g., in image processing, the dynamics at a pixel depends on neighboring pixel values [28], [29]; in power systems, the power grid models, under certain assumptions, exhibit banded structures, [30], [31], [32]. Systems that are sparse but not localized can be converted to sparse and localized by using matrix bandwidth reduction algorithms [33]. We further assume that the overall system is coupled, irreducible, and globally observable.

Let the system described in (10) be monitored by a network of N sensors. Observations at sensor l and time k are

$$\mathbf{y}_k^{(l)} = \mathbf{H}_l \mathbf{x}_k + \mathbf{w}_k^{(l)}, \quad (11)$$

where $\mathbf{H}_l \in \mathbb{R}^{p_l \times n}$ is the local observation matrix for sensor l , p_l is the number of simultaneous observations made by sensor l at time k , and $\mathbf{w}_k^{(l)} \in \mathbb{R}^{p_l}$ is the local observation noise. In the context of the systems we are interested in, it is natural to assume that the observations are localized. These local observations at sensor l may be, e.g., the temperature or height at location l or an average of the temperatures or heights at l and neighboring locations.

We stack the observations at all N sensors in the sensor network to get the global observation model as follows. Let p be the total number of observations at all the sensors. Let the global observation vector, $\mathbf{y}_k \in \mathbb{R}^p$, the global observation matrix, $\mathbf{H} \in \mathbb{R}^{p \times n}$, and the global observation noise vector, $\mathbf{w}_k \in \mathbb{R}^p$, be

$$\mathbf{y}_k = \begin{bmatrix} \mathbf{y}_k^{(1)} \\ \vdots \\ \mathbf{y}_k^{(N)} \end{bmatrix}, \quad \mathbf{H} = \begin{bmatrix} \mathbf{H}_1 \\ \vdots \\ \mathbf{H}_N \end{bmatrix}, \quad \mathbf{w}_k = \begin{bmatrix} \mathbf{w}_k^{(1)} \\ \vdots \\ \mathbf{w}_k^{(N)} \end{bmatrix}. \quad (12)$$

Then the global observation model is given by

$$\mathbf{y}_k = \mathbf{H}\mathbf{x}_k + \mathbf{w}_k. \quad (13)$$

We adopt standard assumptions on the statistical characteristics of the noise. The state noise sequence, $\{\mathbf{u}_k\}_{k \geq 0}$, the observation noise sequence, $\{\mathbf{w}_k\}_{k \geq 0}$, and the initial conditions, \mathbf{x}_0 , are independent, Gaussian, zero-mean, with

$$\mathbb{E}[\mathbf{u}_\ell \mathbf{u}_\tau^H] = \mathbf{Q} \delta_{\ell\tau} \text{ and } \mathbb{E}[\mathbf{w}_\ell \mathbf{w}_\tau^H] = \mathbf{R} \delta_{\ell\tau}, \text{ and } \mathbb{E}[\mathbf{x}_0 \mathbf{x}_0^H] = \mathbf{S}_0, \quad (14)$$

where the superscript H denotes the Hermitian, the Kronecker delta $\delta_{\ell\tau} = 1$, if and only if $\ell = \tau$, and zero otherwise. Since the observation noise at different sensors is independent, we can partition the global observation noise covariance matrix, \mathbf{R} , with $\mathbf{R}_l \in \mathbb{R}^{p_l \times p_l}$ being the local observation noise covariance matrix at sensor l , as

$$\mathbf{R} = \text{blockdiag}[\mathbf{R}_1, \dots, \mathbf{R}_N]. \quad (15)$$

For the rest of the presentation, we consider time-invariant models, specifically, the matrices, \mathbf{F} , \mathbf{G} , \mathbf{H} , \mathbf{Q} , \mathbf{R} , are time-invariant. The discussion, however, is not limited to either zero-mean initial condition or time-invariant models and generalizations to the time-variant models will be added as we proceed.

B. Centralized Information Filter

Let $\check{\mathbf{S}}_{k|k}$ and $\check{\mathbf{S}}_{k|k-1}$ be (filtered and prediction) error covariances, and their inverses be the information matrices, $\check{\mathbf{Z}}_{k|k}$ and $\check{\mathbf{Z}}_{k|k-1}$. Let $\hat{\mathbf{x}}_{k|k}$ and $\hat{\mathbf{x}}_{k|k-1}$ be the filtered estimate and the predicted estimate of the state vector, \mathbf{x}_k , respectively. We have the following relations.

$$\check{\mathbf{S}}_{k|k} = \check{\mathbf{Z}}_{k|k}^{-1} \quad (16)$$

$$\check{\mathbf{S}}_{k|k-1} = \check{\mathbf{Z}}_{k|k-1}^{-1} \quad (17)$$

$$\hat{\mathbf{z}}_{k|k-1} = \check{\mathbf{Z}}_{k|k-1} \hat{\mathbf{x}}_{k|k-1} \quad (18)$$

$$\hat{\mathbf{z}}_{k|k} = \check{\mathbf{Z}}_{k|k} \hat{\mathbf{x}}_{k|k} \quad (19)$$

Define the n -dimensional global observation variables as

$$\mathbf{i}_k = \mathbf{H}^T \mathbf{R}^{-1} \mathbf{y}_k, \quad (20)$$

$$\mathcal{I} = \mathbf{H}^T \mathbf{R}^{-1} \mathbf{H}, \quad (21)$$

and the n -dimensional local observation variables at sensor l as

$$\mathbf{i}_{l,k} = \mathbf{H}_l^T \mathbf{R}_l^{-1} \mathbf{y}_k^{(l)}, \quad (22)$$

$$\mathcal{I}_l = \mathbf{H}_l^T \mathbf{R}_l^{-1} \mathbf{H}_l. \quad (23)$$

When the observations are distributed among the sensors, see (11), the CIF can be implemented by collecting all the sensor observations at a central location; or with observation fusion by realizing that the global observation variables in (20)–(21), can be written as, see [3], [11], [7],

$$\begin{aligned} \mathbf{i}_k &= \mathbf{H}^T \mathbf{R}^{-1} \mathbf{y}_k \\ &= \mathbf{H}_1^T \mathbf{R}_1^{-1} \mathbf{y}_k^{(1)} + \cdots + \mathbf{H}_N^T \mathbf{R}_N^{-1} \mathbf{y}_k^{(N)} \\ &= \sum_{l=1}^N \mathbf{i}_{l,k}. \end{aligned} \quad (24)$$

Similarly,

$$\mathcal{I} = \sum_{l=1}^N \mathcal{I}_l. \quad (25)$$

The *filter step* of the CIF is

$$\check{\mathbf{Z}}_{k|k} = \check{\mathbf{Z}}_{k|k-1} + \sum_{l=1}^N \mathcal{I}_l, \quad k \geq 0 \quad (26a)$$

$$\hat{\mathbf{z}}_{k|k} = \hat{\mathbf{z}}_{k|k-1} + \sum_{l=1}^N \mathbf{i}_{l,k}, \quad k \geq 0. \quad (26b)$$

The *prediction step* of the CIF is

$$\check{\mathbf{Z}}_{k|k-1} = \check{\mathbf{S}}_{k|k-1}^{-1} = (\mathbf{F}\check{\mathbf{Z}}_{k-1|k-1}^{-1}\mathbf{F}^T + \mathbf{G}\mathbf{Q}\mathbf{G}^T)^{-1}, \quad k \geq 1, \quad \check{\mathbf{Z}}_{0|-1} = \check{\mathbf{S}}_0^{-1}, \quad (27a)$$

$$\hat{\mathbf{z}}_{k|k-1} = \check{\mathbf{Z}}_{k|k-1} \left(\mathbf{F}\check{\mathbf{Z}}_{k-1|k-1}^{-1}\hat{\mathbf{z}}_{k-1|k-1} \right), \quad k \geq 1, \quad \hat{\mathbf{z}}_{0|-1} = \mathbf{0}. \quad (27b)$$

C. Centralized L -Banded Information filters

To avoid the $O(n^3)$ computations of the global quantities in (27), e.g., the inversion, $\check{\mathbf{Z}}_{k-1|k-1}^{-1}$, we approximate the information matrices, $\check{\mathbf{Z}}_{k|k}$ and $\check{\mathbf{Z}}_{k|k-1}$, to be L -banded matrices, $\mathbf{Z}_{k|k}$ and $\mathbf{Z}_{k|k-1}$. We refer to the CIF with this approximation as the centralized L -banded Information filter (CLBIF). This approach is studied in [34], where the information loss between $\check{\mathbf{Z}}$ and \mathbf{Z} , is given by the divergence (see [34])

$$\text{Divergence}(\check{\mathbf{Z}}, \mathbf{Z}) = \frac{1}{2} \left\| \check{\mathbf{Z}}^{-\frac{T}{2}} (\mathbf{Z} - \check{\mathbf{Z}}) \mathbf{Z}^{-\frac{1}{2}} \right\|_F^2 \leq \frac{1}{2} \left(\sum_i \lambda_i^{-\frac{1}{2}}(\check{\mathbf{Z}}) \right)^2 \left(\sum_i \lambda_i^{-\frac{1}{2}}(\mathbf{Z}) \right)^2 \left\| \mathbf{Z} - \check{\mathbf{Z}} \right\|_F^2, \quad (28)$$

where $\|\cdot\|_F$ is the Frobenius norm and $\lambda_i(\mathbf{Z})$ is the i th eigenvalue of the matrix \mathbf{Z} .

This approximation on the information matrices is equivalent to approximating the Gaussian error processes,

$$\epsilon_{k|k} = \mathbf{x}_k - \hat{\mathbf{x}}_{k|k}, \quad \epsilon_{k|k-1} = \mathbf{x}_k - \hat{\mathbf{x}}_{k|k-1}, \quad (29)$$

to be Gauss-Markovian of L th order [16]. Reference [20] presents an algorithm to derive the approximation that is optimal in Kullback-Leibler or maximum entropy sense in the class of all L -banded matrices approximating the inverse of the error covariance matrix. In the sequel, we assume this optimal L -banded approximation.

The CLBIF (with the L -banded information matrices, $\mathbf{Z}_{k|k}$ and $\mathbf{Z}_{k|k-1}$) is given by the *filter step* in (26a)–(26b) and the *prediction step* in (27a)–(27b), where the optimal information matrices, $\check{\mathbf{Z}}_{k|k}$ and $\check{\mathbf{Z}}_{k|k-1}$, are replaced by their L -banded approximations. The algorithms in [20], [35] reduce the computational complexity of the CLBIF to $O(n^2)$ but the resulting algorithm is still centralized and deals with the n -dimensional state. To distribute the CLBIF, we start by distributing the global model (10)–(13) in the following section.

III. SPATIAL DECOMPOSITION OF COMPLEX LARGE-SCALE SYSTEMS

Instead of implementing CLBIF based on the global model, we implement local Information filters at the subsystems obtained by spatially decomposing the overall system. Subsection III-A deals with this decomposition by exploiting the sparse and localized structure of the model matrix, \mathbf{F} . In the following, we devise a sub-system at each sensor monitoring the system. This can be applied to general model matrices, \mathbf{F} , but is only practical if \mathbf{F} is sparse and localized as we explain below.

A. Reduced Models at Each Sensor

This subsection shows how to distribute the global model (10) and (13), in order to get the reduced order subsystems. We illustrate the procedure with a simple example that reflects our assumptions on the dynamical system

structure. Consider a five dimensional system with the global dynamical model

$$\begin{aligned} \mathbf{x}_{k+1} &= \begin{bmatrix} f_{11} & f_{12} & 0 & 0 & 0 \\ f_{21} & f_{22} & 0 & f_{24} & 0 \\ f_{31} & 0 & f_{33} & 0 & 0 \\ 0 & 0 & f_{43} & 0 & f_{45} \\ 0 & 0 & 0 & f_{54} & f_{55} \end{bmatrix} \mathbf{x}_k + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & g_{32} \\ 0 & 0 \\ g_{51} & 0 \end{bmatrix} \mathbf{u}_k \\ &= \mathbf{F}\mathbf{x}_k + \mathbf{G}\mathbf{u}_k. \end{aligned} \quad (30)$$

The system has two external noise sources $\mathbf{u}_k = [u_{1k}, u_{2k}]^T$. We monitor this system with $N = 3$ sensors, having scalar observations, $y_k^{(l)}$, at each sensor l . The global observation vector, \mathbf{y}_k , stacks the local observations, $y_k^{(l)}$, and is

$$\begin{aligned} \mathbf{y}_k &= \begin{bmatrix} y_k^{(1)} \\ y_k^{(2)} \\ y_k^{(3)} \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 \end{bmatrix} \mathbf{x}_k + \begin{bmatrix} w_k^{(1)} \\ w_k^{(2)} \\ w_k^{(3)} \end{bmatrix} \\ &= \mathbf{H}\mathbf{x}_k + \mathbf{w}_k, \end{aligned} \quad (31)$$

where $\mathbf{H} = [\mathbf{H}_1^T \ \mathbf{H}_2^T \ \mathbf{H}_3^T]^T$. We distribute global model of equations (30) and (31) in the following Subsections.

1) *Graphical Representation using System Digraphs*: A system digraph visualizes the dynamical interdependence of the system. A system digraph, [13], $\mathcal{J} = [V, E]$, is a directed graphical representation of the system, where $V = X \cup U$ is the vertex set consisting of the states, $X = \{x_i\}_{i=1, \dots, n}$, and the noise inputs, $U = \{u_i\}_{i=1, \dots, j}$. The interconnection matrix, E , is the binary representation (having a 1 for each non-zero entry) of the model matrix, \mathbf{F} , and the state noise matrix, \mathbf{G} , concatenated together. The interconnection matrix, E , for the system in (30) is,

$$E = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 \end{bmatrix}. \quad (32)$$

The system digraph is shown in Figure 1(a).

2) *Reduced Models from the Cut-point Sets*: We have $N = 3$ sensors monitoring the system through the observation model (31). We associate to each sensor l a *cut-point set*, $V^{(l)}$, where $V^{(l)} \subseteq X$. If a sensor observes a linear combination of several states, its cut-point set, $V^{(l)}$, includes all these states; see [36] for the definition of cut-point sets, and algorithms to find all cut-point sets and a minimal cut-point set, if it exists. The cut-point sets select the local states involved in the local dynamics at each sensor. From (31), the cut-point sets³ are shown in

³For simplicity of the presentation, we chose here that each state variable is observed by at least one sensor. If this is not true, we can easily account for this by extending the cut-point sets, $V^{(l)}$, to $\bar{V}^{(l)}$, such that

$$\bigcup_{l=1}^N \bar{V}^{(l)} = X. \quad (33)$$

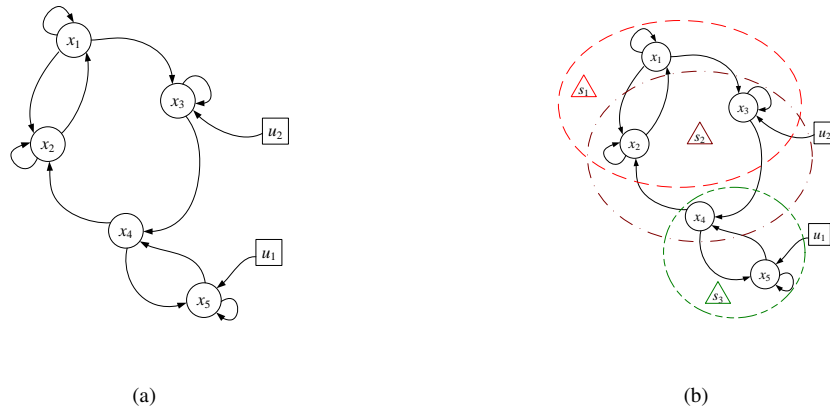


Fig. 1. System Digraph and cut-point sets: (a) Digraph representation of the 5 dimensional system, (30)–(31). The circles represent the states, \mathbf{x} , and the squares represent the input noise sources, \mathbf{u} . (b) The cut-point sets associated to the 3 sensors (Δ) are shown by the dashed circles.




Fig. 2. Partitioning of the global model matrix, \mathbf{F} , into local model matrices, $\mathbf{F}^{(l)}$, and the local internal input matrices, $\mathbf{D}^{(l)}$, shown for sensor 1 and sensor 3, from the example system, (30)–(31).

Figure ?? where we have the following cut-point set, e.g., at sensor 1,

$$V^{(1)} = \{x_1, x_2, x_3\}. \quad (34)$$

Dimension of the sub-systems: The local states at sensor l , i.e., the components of the local state vector, $\mathbf{x}_k^{(l)}$, are the elements in its associated cut-point set, $V^{(l)}$. The dimension of the local Kalman filter implemented at sensor l is now n_l . The set of n_l -dimensional local Kalman filters will give rise, as will be clear later, to an L -banded centralized Information matrix with $L = \min(n_1, \dots, n_N)$. The loss in the optimality as a function of L is given by the divergence (28). Hence, for a desired level of performance, i.e., for a fixed L , we may need to extend (include additional states in a cut-point set) the cut-point sets, $V^{(l)}$, to $V_L^{(l)}$, such that

$$n_l = |V_L^{(l)}| \geq L, \quad \forall l, \quad (35)$$

where $|\cdot|$ when applied to a set denotes its cardinality. This procedure of choosing an L based on a certain desired performance gives a lower bound on the dimensions of the sub-systems.

Coupled States as Inputs: The directed edges coming into a cut-point set are the inputs required by the reduced model. In the context of our running illustration (30)–(31), we see that the local state vector for sensor 1 is, $\mathbf{x}_k^{(1)} = [x_{1,k}, x_{2,k}, x_{3,k}]^T$, and the inputs to the local model consist of a subset of the state set, X , (at sensor s_1 , $x_{4,k}$ is the input coming from sensor s_2) and a subset of the noise input set, U , ($u_{2,k}$ at sensor s_1).

Local models: For the local model at sensor l , we collect the states required as input in a local internal input vector, $\mathbf{d}_k^{(l)}$ (we use the word *internal* to distinguish from the externally applied inputs), and the noise sources required as input in a local noise input vector, $\mathbf{u}_k^{(l)}$. We collect the elements from \mathbf{F} corresponding to the local state vector, $\mathbf{x}_k^{(l)}$, in a local model matrix, $\mathbf{F}^{(l)}$. Similarly, we collect the elements from \mathbf{F} corresponding to the local internal input vector, $\mathbf{d}_k^{(l)}$, in a local internal input matrix, $\mathbf{D}^{(l)}$, and the elements from \mathbf{G} corresponding to the local noise input vector, $\mathbf{u}_k^{(l)}$, in a local state noise matrix, $\mathbf{G}^{(l)}$. Figure 2 shows this partitioning for sensors s_1 and s_3 . We have the following local models for (30).

$$\begin{aligned} \mathbf{x}_{k+1}^{(1)} &= \begin{bmatrix} f_{11} & f_{12} & 0 \\ f_{21} & f_{22} & 0 \\ f_{31} & 0 & f_{33} \end{bmatrix} \mathbf{x}_k^{(1)} + \begin{bmatrix} 0 \\ f_{24} \\ 0 \end{bmatrix} x_{4,k} + \begin{bmatrix} 0 \\ 0 \\ g_{32} \end{bmatrix} u_{2,k}, \\ &= \mathbf{F}^{(1)} \mathbf{x}_k^{(1)} + \mathbf{D}^{(1)} \mathbf{d}_k^{(1)} + \mathbf{G}^{(1)} \mathbf{u}_k^{(1)} \end{aligned} \quad (36)$$

$$\begin{aligned} \mathbf{x}_{k+1}^{(2)} &= \begin{bmatrix} f_{22} & 0 & f_{42} \\ 0 & f_{33} & 0 \\ 0 & f_{43} & 0 \end{bmatrix} \mathbf{x}_k^{(2)} + \begin{bmatrix} f_{21} & 0 \\ f_{31} & 0 \\ 0 & f_{45} \end{bmatrix} \begin{bmatrix} x_{1,k} \\ x_{5,k} \end{bmatrix} + \begin{bmatrix} 0 \\ g_{32} \\ 0 \end{bmatrix} u_{2,k}, \\ &= \mathbf{F}^{(2)} \mathbf{x}_k^{(2)} + \mathbf{D}^{(2)} \mathbf{d}_k^{(2)} + \mathbf{G}^{(2)} \mathbf{u}_k^{(2)} \end{aligned} \quad (37)$$

$$\begin{aligned} \mathbf{x}_{k+1}^{(3)} &= \begin{bmatrix} 0 & f_{45} \\ f_{54} & f_{55} \end{bmatrix} \mathbf{x}_k^{(3)} + \begin{bmatrix} f_{43} \\ 0 \end{bmatrix} x_{3,k} + \begin{bmatrix} 0 \\ g_{51} \end{bmatrix} u_{1,k}, \\ &= \mathbf{F}^{(3)} \mathbf{x}_k^{(3)} + \mathbf{D}^{(3)} \mathbf{d}_k^{(3)} + \mathbf{G}^{(3)} \mathbf{u}_k^{(3)} \end{aligned} \quad (38)$$

We may also capture the above extraction of the local states by the cut-point sets, with the following procedure. Let the total number of states in the cut-point set at sensor l , $V^{(l)}$, be n_l . Let \mathbf{T}_l be an $n_l \times n$ selection matrix, such that it selects n_l states in the cut-point set, $V^{(l)}$, from the entire state vector, \mathbf{x}_k , according to the following relation,

$$\mathbf{x}_k^{(l)} = \mathbf{T}_l \mathbf{x}_k. \quad (39)$$

For sensor 1, the selection matrix, \mathbf{T}_1 , is

$$\mathbf{T}_1 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}. \quad (40)$$

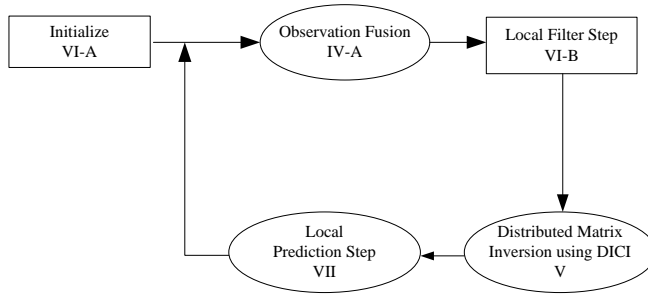


Fig. 3. Block Diagram for the LIFs: Steps involved in the LIF implementation. The ovals represent the steps that require local communication.

We establish a *reduced* local observation matrix, $\mathbf{H}^{(l)}$, by retaining the terms corresponding to the local state vector, $\mathbf{x}_k^{(l)}$, from the local observation matrix, \mathbf{H}_l . We may write

$$\mathbf{H}^{(l)} = \mathbf{H}_l \mathbf{T}_l^\#, \quad (41)$$

where ‘#’ denotes the pseudo-inverse of the matrix. In the context of the running illustration, the reduced local observation matrix $\mathbf{H}^{(1)} = [1, 1, 1]$ is obtained from the local observation matrix $\mathbf{H}_1 = [1, 1, 1, 0, 0]$. Note that \mathbf{H}_l picks the states from the global state vector, \mathbf{x}_k , whereas $\mathbf{H}^{(l)}$ picks the states from the local state vector, $\mathbf{x}_k^{(l)}$. The reduced local observation models are given by

$$\mathbf{y}_k^{(l)} = \mathbf{H}^{(l)} \mathbf{x}_k^{(l)} + \mathbf{w}_k^{(l)}. \quad (42)$$

We now make some additional comments. For simplicity of the explanation, we refer to our running example, (30)–(31). We note that the reduced models at the sensors overlap, as shown by the overlapping cut-point sets in Figure ???. Due to this overlap, observations corresponding to the shared states are available at multiple sensors that should be fused. We further note that the reduced model (36) at sensor 1 is coupled to the reduced model (37) at sensor 2 through the state $x_{4,k}$. The state $x_{4,k}$ at sensor 1 does not appear in the local state vector, i.e., $x_{4,k} \notin \mathbf{x}_k^{(1)}$. But it is still required as an internal input at sensor 1 to preserve the global dynamics. Hence, sensor 2 communicates the state $x_{4,k}$, which appears in its local state vector, i.e., $x_{4,k} \in \mathbf{x}_k^{(2)}$, to sensor 1. Hence at an arbitrary sensor l , we derive the reduced model to be

$$\mathbf{x}_{k+1}^{(l)} = \mathbf{F}^{(l)} \mathbf{x}_k^{(l)} + \mathbf{D}^{(l)} \mathbf{d}_k^{(l)} + \mathbf{G}^{(l)} \mathbf{u}_k^{(l)}. \quad (43)$$

Since the value of the state itself is unknown, sensor 2 communicates its estimate, $\hat{x}_{4,k|k}^{(2)}$, to sensor 1. This allows sensor 1 to complete its local model and preserve global dynamics, thus, taking into account the coupling between the local reduced-order models. This process is repeated at all sensors. Hence, the local internal input vector, $\mathbf{d}_k^{(l)}$, is replaced by its estimate, $\hat{\mathbf{d}}_{k|k}^{(l)}$. It is worth mentioning here that if the dynamics were time-dependent, i.e., the matrices, \mathbf{F} , \mathbf{G} , \mathbf{H} , change with time, k , then the above decomposition procedure will have to be repeated at each k . This may result into a different communication topology over which the sensors communicate at each k .

$$\mathbf{S} = \begin{pmatrix} s_{11} & s_{12} & s_{13} & s_{14} & s_{15} \\ s_{12} & s_{22} & s_{23} & s_{24} & s_{25} \\ s_{13} & s_{23} & s_{33} & s_{34} & s_{35} \\ s_{14} & s_{24} & s_{34} & s_{44} & s_{45} \\ s_{15} & s_{25} & s_{35} & s_{45} & s_{55} \end{pmatrix} = \begin{pmatrix} z_{11} & z_{12} & 0 \\ z_{12} & z_{22} & z_{23} \\ z_{23} & z_{33} & z_{34} \\ 0 & z_{34} & z_{44} & z_{45} \\ 0 & z_{45} & z_{45} & z_{55} \end{pmatrix}^{-1} = \mathbf{Z}^{-1}$$

Fig. 4. Relationship between the global error covariance matrices, \mathbf{S} and their inverses, the global information matrices, \mathbf{Z} , with $L = 1$ -banded approximation on \mathbf{Z} . The figure also shows how the local matrices, $\mathbf{S}^{(l)}$ and $\mathbf{Z}^{(l)}$, constitute their global counterparts. Since this relation holds for both the estimation and prediction matrices, the figure removes the subscripts.

B. Local Information Filters

To distribute the estimation of the global state vector, \mathbf{x}_k , we implement local Information filters (LIFs) at each sensor l , which are based on the sensor-based reduced models (43) and (42). Each LIF computes local objects (matrices and vectors of dimension n_l) which are then fused (if required) by exchanging information among the neighbors. Some of the update procedures are iterative. Although, no centralized knowledge of the estimation of the global state exists, the union of the local state vector represents, in a distributed way, the knowledge that exists in a centralized fashion in the CLBIF. In most applications nowhere in the network is there the need for this centralized knowledge.

The LIFs consist of initial conditions, a local filter step (including observation fusion and distributed matrix inversion) and a local prediction step (including estimate fusion), see Figure 3. These steps are presented in the next four sections. To proceed with the next sections, we provide notation in the following Subsection.

1) *Notation:* We use the notation that the superscript (l) refers to a local reduced-order variable ($n_l \times 1$ vector or $n_l \times n_l$ matrix) at sensor l to define the reduced observation vector, $\mathbf{i}_k^{(l)}$, and the reduced observation matrix, $\mathcal{I}^{(l)}$, as

$$\mathbf{i}_k^{(l)} = (\mathbf{H}^{(l)})^T \mathbf{R}_l^{-1} \mathbf{y}_k^{(l)}, \quad (44)$$

$$\mathcal{I}^{(l)} = (\mathbf{H}^{(l)})^T \mathbf{R}_l^{-1} \mathbf{H}^{(l)}. \quad (45)$$

The local error covariance matrices, $\mathbf{S}_{k|k}^{(l)}$ and $\mathbf{S}_{k|k-1}^{(l)}$, are the overlapping diagonal submatrices of the global error covariance matrices, $\mathbf{S}_{k|k}$ and $\mathbf{S}_{k|k-1}$. Let $\mathbf{Z}_{k|k}^{(l)}$ and $\mathbf{Z}_{k|k-1}^{(l)}$ be the local information matrices. These local information matrices are overlapping diagonal submatrices of the global L -banded information matrices, $\mathbf{Z}_{k|k}$ and $\mathbf{Z}_{k|k-1}$. These local matrices overlap because the reduced sensor-based models (43) have overlapping state vectors, $\mathbf{x}_k^{(l)}$. Figure 4 captures the relationship between the local error covariance matrices and the local information matrices given by (16) and (17).

IV. OVERLAPPING REDUCED MODELS

After the model distribution step introduced in section III, the reduced models among the sensors may share states, as shown by the overlapped cut-point sets in Figure ???. Since the sensors sharing the states have independent




Fig. 5. Bipartite Fusion graph, \mathcal{B} , is shown for the example system, (30)–(31).

observations of the shared states, observations corresponding to the shared states should be fused. We present observation fusion in subsection IV-A with the help of bipartite fusion graphs, [14].

A. Observation Fusion

Equations (24) and (25) show that the observation fusion is equivalent to adding the corresponding n –dimensional local observation variables, (22)–(23). In CLBIF, we implement this fusion directly because each local observation variable in (22)–(23) corresponds to the full n –dimensional state vector, \mathbf{x}_k . Since the n_l –dimensional *reduced* observation variables, (44)–(45), correspond to different local state vectors, $\mathbf{x}_k^{(l)}$, they cannot be added directly.

To achieve observation fusion, we introduce the following undirected bipartite fusion graph⁴, \mathcal{B} . Let $S_N = \{s_1, \dots, s_N\}$ be the set of sensors and X be the set of states. The vertex set of the bipartite fusion graph, \mathcal{B} , is $S_N \cup X$. We now define the edge set, $E_{\mathcal{B}}$, of the fusion graph, \mathcal{B} . The sensor s_i is connected to the state variable x_j , if s_i observes the state variable x_j . In other words, we have an edge between sensor s_i and state variable x_j , if the local observation matrix, \mathbf{H}_i , at sensor s_i , contains a non-zero entry in its j th column. Figure 5 shows the bipartite graph for the example system in (30)–(31).

We now further provide some notation on the communication topology. Let \mathcal{G} denote the sensor communication graph that determines how the sensors communicate among themselves. Let $\mathcal{K}(l)$ be the subgraph of \mathcal{G} , that contains the sensor l and the sensors directly connected to the sensor l . For the j th state, x_j , let \mathcal{G}_j be the induced subgraph of the sensor communication graph, \mathcal{G} , such that it contains the sensors having the state x_j in their reduced models.

The set of vertices in \mathcal{G}_j come directly from the bipartite fusion graph, \mathcal{B} . For example, from Figure 5, we see that \mathcal{G}_1 contains s_1 as a single vertex, \mathcal{G}_2 contains s_1, s_2 as vertices, and so on. States having more than one sensor connected to them in the bipartite fusion graph, \mathcal{B} , are the states for which fusion is required, since we have multiple observations for that state. Furthermore, Figure 5 also gives the vertices in the associated subgraphs, \mathcal{G}_j , over which the fusion is to be carried out.

With the help of the above discussion, we establish the fusion of the reduced observation variables, (44)–(45). The reduced model at each sensor involves n_l state variables, and each element in the $n_l \times 1$ reduced observation vector, $\mathbf{i}_k^{(l)}$, corresponds to one of these states, i.e., each entry in $\mathbf{i}_k^{(l)}$ has some information about its corresponding

⁴A bipartite graph is a graph whose vertices can be divided into two disjoint sets X and S_N , such that, every edge connects a vertex in X to a vertex in S_N , and, there is no edge between any two vertices of the same set, [37].

state variable. Let the entries of the $n_l \times 1$ reduced observation vector, $\mathbf{i}_k^{(l)}$, at sensor l , be subscripted by the n_l state variables modeled at sensor l . In the context of the example given by system (30)–(31), we have

$$\mathbf{i}_k^{(1)} = \begin{bmatrix} i_{k,x_1}^{(1)} \\ i_{k,x_2}^{(1)} \\ i_{k,x_3}^{(1)} \end{bmatrix}, \quad \mathbf{i}_k^{(2)} = \begin{bmatrix} i_{k,x_2}^{(2)} \\ i_{k,x_3}^{(2)} \\ i_{k,x_4}^{(2)} \end{bmatrix}, \quad \mathbf{i}_k^{(3)} = \begin{bmatrix} i_{k,x_4}^{(3)} \\ i_{k,x_5}^{(3)} \end{bmatrix}. \quad (46)$$

For each state x_j , the observation fusion is carried out on the sensors attached to this state in the bipartite fusion graph, \mathcal{B} . The fused observation vectors denoted by $\mathbf{i}_{f,k}^{(l)}$ are given by

$$\mathbf{i}_{f,k}^{(1)} = \begin{bmatrix} i_{k,x_1}^{(1)} \\ i_{k,x_2}^{(1)} + i_{k,x_2}^{(2)} \\ i_{k,x_3}^{(1)} + i_{k,x_3}^{(2)} \end{bmatrix}, \quad \mathbf{i}_{f,k}^{(2)} = \begin{bmatrix} i_{k,x_2}^{(2)} + i_{k,x_2}^{(1)} \\ i_{k,x_3}^{(2)} + i_{k,x_3}^{(1)} \\ i_{k,x_4}^{(2)} + i_{k,x_4}^{(3)} \end{bmatrix}, \quad \mathbf{i}_{f,k}^{(3)} = \begin{bmatrix} i_{k,x_4}^{(3)} + i_{k,x_4}^{(2)} \\ i_{k,x_5}^{(3)} \end{bmatrix}. \quad (47)$$

Generalizing to the arbitrary sensor l , we may write the entry, $i_{f,k,x_j}^{(l)}$, corresponding to x_j in the fused observation vector, $\mathbf{i}_{f,k}^{(l)}$ as

$$i_{f,k,x_j}^{(l)} = \sum_{s \in \mathcal{G}_j} i_{k,x_j}^{(s)}, \quad (48)$$

where $i_{k,x_j}^{(s)}$ is the entry corresponding to x_j in the reduced observation vector at sensor s , $\mathbf{i}_k^{(s)}$.

Since the communication network on \mathcal{G}_j may not be all-to-all, an iterative weighted averaging algorithm [15], is used to compute the fusion in (48) over arbitrarily connected communication networks with only local communication (the communication may be multi-hop). A similar procedure on the pairs of state variables and their associated subgraphs can be implemented to fuse the reduced observation matrices, $\mathcal{I}^{(l)}$. Since, we assume the observation model to be stationary (\mathbf{H} and \mathbf{R} are time-independent), the fusion on the reduced observation matrix, $\mathcal{I}^{(l)}$, is to be carried out only once and can be an offline procedure. If that is not the case and \mathbf{H} and \mathbf{R} are time dependent fusion on \mathcal{I} has to be repeated at each time, k .

Here, we assume that the communication is fast enough so that the consensus algorithm can converge, see [38] for a discussion on distributed Kalman filtering based on consensus strategies. The convergence of the consensus algorithm is shown to be geometric and the convergence rate can be increased by optimizing the weight matrix for the consensus iterations using semidefinite programming (SDP) [39]. The communication topology of the sensor network can also be improved to increase the convergence speed of the consensus algorithms [40], [41], [42].

Here, we comment on estimate fusion. Since we fused the observations concerning the shared states among the sensors, one may ask if it is required to carry out fusion of the estimates of the shared states. It turns out that consensus on the observations leads to consensus on the estimates. This will become clear with the introduction of the local filter and the local prediction step of the LIFs, therefore, we defer the discussion on estimate fusion to section VII-C.

V. DISTRIBUTED MATRIX INVERSION WITH LOCAL COMMUNICATION

In this section, we discuss the fusion of the error covariances. Consider the example model (30)–(31), when we employ LIFs on the distributed models (36)–(38). The local estimation information matrices, $\mathbf{Z}_{k|k}^{(1)}$, $\mathbf{Z}_{k|k}^{(2)}$, and $\mathbf{Z}_{k|k}^{(3)}$,

correspond to the overlapping diagonal submatrices of the global 5×5 estimation information matrix, $\mathbf{Z}_{k|k}$, see Figure 4, with $L = 1$ -banded assumption on $\mathbf{Z}_{k|k}$. It will be shown (section VII-A) that the local prediction information matrix, $\mathbf{Z}_{k+1|k}^{(l)}$, is a function of the local error covariance matrices, $\mathbf{S}_{k|k}^{(l)}$, and hence we need to compute $\mathbf{S}_{k|k}^{(l)}$ from the local filter information matrices, $\mathbf{Z}_{k|k}^{(l)}$, which we get from the local filter step (section VI). As can be seen from Figure 4 and (16), for these local submatrices,

$$\mathbf{S}^{(l)} \neq \left(\mathbf{Z}^{(l)}\right)^{-1}. \quad (49)$$

Collecting all the local information matrices, $\mathbf{Z}_{k|k}^{(l)}$, at each sensor and then carrying out an $n \times n$ matrix inversion is *not* a practical solution for large-scale systems (where n may be large), because of the large communication overhead and $O(n^3)$ computational cost. Using the L -banded structure on the global estimation information matrix, $\mathbf{Z}_{k|k}$, we present below a distributed iterate collapse inversion overrelaxation (DICI-OR, pronounced *die-see-O-R*) algorithm⁵. For notational convenience, we disregard the time indices in the following development.

We present a generalization of the centralized Jacobi overrelaxation (JOR) algorithm to solve matrix inversion in section V-A and show that the computations required in its distributed implementation scale linearly with the dimension, n , of the system. We then present the DICI-OR algorithm and show that it is independent of the dimension, n , of the system.

A. Centralized Jacobi Overrelaxation (JOR) Algorithm

The centralized JOR algorithm for vectors [19] solves a linear system of n equations iteratively, by successive substitution. It can be easily extended to get the centralized JOR algorithm for matrices that solves

$$\mathbf{ZS} = \mathbf{T}, \quad (50)$$

for the unknown matrix, \mathbf{S} , where the matrices \mathbf{Z} and \mathbf{T} are known. Let $\mathbf{M} = \text{diag}(\mathbf{Z})$, for some $\gamma > 0$

$$\mathbf{S}_{t+1} = \left((1 - \gamma)\mathbf{I}_{n \times n} + \gamma\mathbf{M}^{-1}(\mathbf{M} - \mathbf{Z})\right)\mathbf{S}_t + \gamma\mathbf{M}^{-1}\mathbf{T}, \quad (51)$$

converges to \mathbf{S} , and is the centralized JOR algorithm for matrices, solving n coupled linear systems of equations (50), where γ is sometimes called a relaxation parameter [19]. Putting $\mathbf{T} = \mathbf{I}_{n \times n}$, we can solve for $\mathbf{ZS} = \mathbf{I}_{n \times n} \Rightarrow \mathbf{S} = \mathbf{Z}^{-1}$, and, if \mathbf{Z} is known, the following iterations converge to \mathbf{Z}^{-1} ,

$$\mathbf{S}_{t+1} = \mathbf{P}_\gamma \mathbf{S}_t + \gamma\mathbf{M}^{-1}, \quad (52)$$

where the multiplier matrix, \mathbf{P}_γ , is defined as

$$\mathbf{P}_\gamma = (1 - \gamma)\mathbf{I}_{n \times n} + \gamma\mathbf{M}^{-1}(\mathbf{M} - \mathbf{Z}). \quad (53)$$

⁵It is worth mentioning here that the DICI algorithm (for solving $\mathbf{ZS} = \mathbf{I}_{n \times n}$, with SPD L -banded matrix $\mathbf{Z} \in \mathbb{R}^{n \times n}$ and the $n \times n$ identity matrix $\mathbf{I}_{n \times n}$) is neither a direct extension nor a generalization of (block) Jacobi or Gauss-Seidel type iterative algorithms (that solve a vector version, $\mathbf{Zs} = \mathbf{b}$ with $\mathbf{s}, \mathbf{b} \in \mathbb{R}^n$, of $\mathbf{ZS} = \mathbf{I}_{n \times n}$, see [19], [43], [44], [45], [46].) Using the Jacobi or Gauss-Seidel type iterative schemes for solving $\mathbf{ZS} = \mathbf{I}$ is equivalent to solving n linear systems of equations, $\mathbf{Zs} = \mathbf{b}$; hence, the complexity scales linearly with n . Instead the DICI algorithm employs a non-linear collapse operator that exploits the structure of the inverse, \mathbf{S} , of a symmetric positive definite L -banded matrix, \mathbf{Z} , which makes its complexity independent of n .

The Jacobi algorithm can now be considered as a special case of the JOR algorithm with $\gamma = 1$.

1) *Convergence*: Let \mathbf{S} be the stationary point of the iterations in (52), and let $\tilde{\mathbf{S}}_{t+1}$ denote the iterations of (52). It can be shown that the error process, $\tilde{\mathbf{E}}_{t+1}$, for the JOR algorithm is

$$\tilde{\mathbf{E}}_{t+1} = \mathbf{P}_\gamma \tilde{\mathbf{E}}_t, \quad (54)$$

which decays to zero if $\|\mathbf{P}_\gamma\|_2 < 1$, where $\|\cdot\|_2$ denotes the spectral norm of a matrix. The JOR algorithm (52) converges for all symmetric positive definite matrices, \mathbf{Z} , for sufficiently small $\gamma > 0$, see [19], an alternate convergence proof is provided in [47] via convex M -matrices, whereas, convergence for parallel asynchronous team algorithms is provided in [48]. Since, the information matrix, \mathbf{Z} , is the inverse of an error covariance matrix; the information matrix, \mathbf{Z} , is symmetric positive definite by definition and the JOR algorithm always converges. Plugging $\gamma = 1$ in (52) gives us the centralized Jacobi algorithm for matrices, which converges for all diagonally dominant matrices, \mathbf{Z} , see [19]. We can further write the error process as,

$$\tilde{\mathbf{E}}_{t+1} = \mathbf{P}_\gamma^{t+1} (\mathbf{S}_0 - \mathbf{S}), \quad (55)$$

where the matrix \mathbf{S}_0 is the initial condition. The spectral norm of the error process can be bounded by

$$\|\tilde{\mathbf{E}}_{t+1}\|_2 \leq |\lambda_{\max}(\mathbf{P}_\gamma)|^{t+1} \|\mathbf{S}_0 - \mathbf{S}\|_2, \quad (56)$$

where λ_{\max} is the maximum eigenvalue (in magnitude) of the multiplier matrix, \mathbf{P}_γ . The JOR algorithm is centralized as it requires the complete $n \times n$ matrices involved. This requires global communication and an n th order computation at each iteration of the algorithm. We present below its distributed implementation.

2) *Distributed JOR algorithm*: We are interested in the local error covariances that lie on the L -band of the matrix $\mathbf{S} = \mathbf{Z}^{-1}$. Distributing the JOR (in addition to [19], distributed Jacobi and Gauss Seidel type iterative algorithms can also be found in [43], [45], [47]) algorithm (52) directly to compute the L -band of \mathbf{S} gives us the following equations for the ij -th element, s_{ij} , in \mathbf{S}_{t+1} ,

$$s_{ij,t+1} = \mathbf{p}_i \mathbf{s}_t^j, \quad i \neq j, \quad (57)$$

$$s_{ij,t+1} = \mathbf{p}_i \mathbf{s}_t^j + m_{ii}^{-1}, \quad i = j, \quad (58)$$

where: the row vector, \mathbf{p}_i , is the i th row of the multiplier matrix, \mathbf{P}_γ ; the column vector, \mathbf{s}_t^j , is the j th column of the matrix, \mathbf{S}_t ; and the scalar element, m_{ii} , is the i th diagonal element of the diagonal matrix, \mathbf{M} . Since the matrix \mathbf{Z} is L -banded, the multiplier matrix, \mathbf{P}_γ , in (53) is also L -banded. The i th row of the multiplier matrix, \mathbf{P}_γ , contains non-zeros at most at $2L + 1$ locations in the index set, $\kappa = \{i - L, \dots, i, \dots, i + L\}$. These non-zero elements pick the corresponding elements with indices in the index set, κ , in the j th column, \mathbf{s}_t^j , of \mathbf{S}_t . Due to the L -bandedness of the multiplier matrix, \mathbf{P}_γ , the JOR algorithm can be easily distributed with appropriate communication with the neighboring sensors.

A major drawback of the distributed JOR algorithm is that at each sensor the computation requirements scale linearly with the dimension, n , of the system. This can be seen by writing out the iteration in (57), e.g., for an

L -banded element such that $|i - j| = L$. In the context of Figure 4, we can write $s_{45,t+1}$ from (57) as

$$s_{45,t+1} = p_{43}s_{35,t} + p_{44}s_{45,t} + p_{45}s_{55,t} \quad (59)$$

The element $s_{35,t}$ does not lie in the L -band of \mathbf{S} , and, hence, does not belong to any local error covariance matrix, $\mathbf{S}^{(l)}$. Iterating on it using (57) gives

$$s_{35,t+1} = p_{32}s_{25,t} + p_{33}s_{35,t} + p_{34}s_{45,t}. \quad (60)$$

The computation in (60), involves $s_{25,t}$, iterating on which, in turn, requires another off L -band element, $s_{15,t}$, and so on. Hence, a single iteration of the algorithm, although distributed and requiring only local communication, sweeps the entire rows in \mathbf{S} and the computation requirements scale linearly with n . We now present a solution to this problem.

B. Distributed Iterate Collapse Inversion Overrelaxation (DICI-OR) Algorithm

In this section, we present the distribute iterate collapse inversion overrelaxation (DICI-OR) algorithm. The DICI-OR algorithm is divided into two steps: (i) an iterate step and (ii) a collapse step. The **iterate step** can be written, in general, for the ij -th element that lies in the L -band ($|i - j| \leq L$) of the matrix \mathbf{S}_{t+1} as

$$s_{ij,t+1} = \mathbf{p}_i \mathbf{S}_t^j, \quad i \neq j, |i - j| \leq L, \quad (61)$$

$$s_{ij,t+1} = \mathbf{p}_i \mathbf{S}_t^j + m_{ii}^{-1}, \quad i = j, |i - j| \leq L, \quad (62)$$

where the symbols are defined as in section V-A.2.

As we explained before in section V-A.2, the implementation of (61)–(62) requires non L -banded elements that, in turn, require more non L -banded elements. To address this problem we introduce a **collapse step**. We assume that \mathbf{S}_t is the inverse of an L -banded matrix and use the results⁶ in [35], to compute a non L -band element (s_{ij} such that $|i - j| > L$) from the L -band elements (s_{ij} such that $|i - j| \leq L$). In general, a non L -band element in a matrix whose inverse is $L = 1$ -banded can be written as

$$s_{ij} = s_{i,j-1} s_{i+1,j-1}^{-1} s_{i+1,j}, \quad |i - j| > L \quad (63)$$

which gives us the collapse step. In the context of Figure 4, instead of iterating on s_{35} as in (60), we employ the collapse step,

$$s_{35,t} = s_{34,t} s_{44,t}^{-1} s_{45,t}, \quad (64)$$

that prevents us from iterating further on the non L -banded elements.

⁶If \mathbf{S} is the inverse of an L -banded matrix, then the submatrices that do not lie in the L -band of \mathbf{S} , can be computed from the submatrices that lie in the L -band of \mathbf{S} , [35]. So, to compute the inverse, $\mathbf{S} = \mathbf{Z}^{-1}$, of an L -banded matrix, \mathbf{Z} , we just compute the submatrices that lie in the L -band of its inverse, \mathbf{S} ; the remaining submatrices are derived from these using the expressions given in [35].

The initial conditions of the DIC-OR algorithm are given by

$$\mathbf{P}_\gamma^{(l)} = (1 - \gamma)\mathbf{I}_{n \times n}^{(l)} + \gamma \left(\mathbf{M}^{(l)} \right)^{-1} \left(\mathbf{M}^{(l)} - \mathbf{Z}^{(l)} \right), \quad (65)$$

$$\mathbf{S}_0^{(l)} = \left(\mathbf{Z}^{(l)} \right)^{-1} \quad (66)$$

Note that equations (65)–(66) do not require any communication and can be computed at each sensor directly from the local information matrix, $\mathbf{Z}^{(l)}$. This is because the matrix \mathbf{M} is diagonal, and its local submatrix, $\mathbf{M}^{(l)}$ in (65), is the exact inverse of the matrix formed by the diagonal elements of $\mathbf{Z}^{(l)}$.

We refer to equations (61)–(63), combined with (65)–(66), and appropriate communication from neighboring sensors as the DIC-OR algorithm. The DIC-OR algorithm can be easily extended to $L > 1$. The only step to take care of is the collapse step, since (63) holds only for $L = 1$. The appropriate formulae to replace (63), when $L > 1$, are provided in [35]. The computation requirements for the DIC algorithm are independent of n and DIC provides a scalable implementation of the matrix inversion problem. The DIC algorithm (without the overrelaxation parameter, γ) can be obtained from the DIC-OR algorithm by setting $\gamma = 1$.

1) *Convergence of the DIC-OR algorithm:* The iterate and the collapse step of the DIC algorithm over the entire sensor network can be combined in matrix form as follows.

$$\text{Iterate Step:} \quad \mathbf{S}_{t+1} = \mathbf{P}_\gamma \bar{\mathbf{S}}_t + \mathbf{M}^{-1}, \quad |i - j| \leq L \quad (67)$$

$$\text{Collapse Step:} \quad \bar{\mathbf{S}}_{t+1} = \zeta(\mathbf{S}_{t+1}), \quad |i - j| > L \quad (68)$$

The operator $\zeta(\cdot)$ is the collapse operator; it takes an arbitrary symmetric positive definite matrix and converts it to a symmetric positive definite matrix whose inverse is L -banded by using the results in [35]. The DIC algorithm is a composition of the linear iterate operator, \mathbf{P}_γ given in (53), followed by the collapse operator, $\zeta(\cdot)$ given in (63) for $L = 1$ and in [35] for $L > 1$. Combining (67) and (68) summarizes the DIC algorithm as,

$$\bar{\mathbf{S}} = \zeta \left(\mathbf{P}_\gamma \left(\bar{\mathbf{S}} + (\mathbf{P}_\gamma \mathbf{M})^{-1} \right) \right). \quad (69)$$

We define a composition map, $\Upsilon : \Xi \mapsto \Xi$, where $\Xi \subset \mathbb{R}^{n \times n}$ is the set of all symmetric positive definite matrices, as $\Upsilon \doteq \zeta \circ \mathbf{P}_\gamma$. To prove the convergence of the DIC-OR algorithm, we are required to show that the composition map, Υ , is a contraction map under some norm, that we choose to be the spectral norm $\|\cdot\|_2$, i.e., for $\alpha \in [0, 1)$,

$$\|\Upsilon(\mathbf{X}_\Xi) - \Upsilon(\mathbf{Y}_\Xi)\|_2 \leq \alpha \|\mathbf{X}_\Xi - \mathbf{Y}_\Xi\|_2, \quad \forall \mathbf{X}_\Xi, \mathbf{Y}_\Xi \in \Xi \quad (70)$$

The convergence of the iterate step of the DIC-OR algorithm is based on the iterate operator, \mathbf{P}_γ , which is proved to be a contraction map in [19]. For the convergence of the collapse operator, ζ , we resort to a numerical procedure and show, in the following, that (70) is a contraction by simulating (70) 1.17×10^6 times.

For the simulations, we generate $n \times n$ matrices, \mathbf{X}_{rand} , with i.i.d. normally distributed elements and get, $\mathbf{X}_{\text{sym}} = \mathbf{X}_{\text{rand}} + \mathbf{X}_{\text{rand}}^T$. We eigen-decompose $\mathbf{X}_{\text{sym}} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$. We replace $\mathbf{\Lambda}$ with a diagonal matrix, $\mathbf{\Lambda}_\Xi$, whose diagonal elements are drawn from a uniform distribution in the interval $(0, 10]$. This leads to a random symmetric positive definite matrix,

$$\mathbf{X}_\Xi = \mathbf{V}\mathbf{\Lambda}_\Xi\mathbf{V}^T.$$




Fig. 6. Histogram of α : Simulations of the quotient of (71) are performed 1.17×10^6 times and the results are provided as a histogram.

For $n = 100$ and L a random integer between 1 and $n/2 = 50$, we compute, by Monte Carlo simulations, the quotient of (70)

$$\frac{\|\Upsilon(\mathbf{X}_{\Xi}) - \Upsilon(\mathbf{Y}_{\Xi})\|_2}{\|\mathbf{X}_{\Xi} - \mathbf{Y}_{\Xi}\|_2}. \quad (71)$$

The number of trials is 1.17×10^6 . The histogram of the values of α , in (71), (with 1000 bins) is plotted in Figure 6. The maximum value of α found in these 1.17×10^6 simulations is 0.9955 and the minimum value is 0.1938. Since $\alpha \in (0, 1)$, i.e., strictly less than 1, we assume that (70) is numerically verified.

2) *Error Bound for the DICI-OR algorithm:* Let the matrix produced by the DICI-OR algorithm at the $t + 1$ -th iteration be $\hat{\mathbf{S}}_{t+1}$. The error process in the DICI-OR algorithm is given by

$$\hat{\mathbf{E}}_{t+1} = \hat{\mathbf{S}}_{t+1} - \mathbf{S}. \quad (72)$$

Claim: The spectral norm of the error process, $\|\hat{\mathbf{E}}_{t+1}\|_2$, of the DICI-OR algorithm is *bounded above* by the spectral norm of the error process, $\|\tilde{\mathbf{E}}_{t+1}\|_2$, of the JOR algorithm. Since, the JOR algorithm always converges for symmetric positive definite matrices, \mathbf{Z} , we deduce that the DICI-OR algorithm converges.

We verify this claim numerically by Monte Carlo simulations. The number of trials is 4490, and we compute the error process, $\hat{\mathbf{E}}_{t+1}(K)$, of the DICI-OR algorithm and the error process, $\tilde{\mathbf{E}}_{t+1}(K)$, of the JOR algorithm. We choose the relaxation parameter, γ , to be 0.1. In Figure ??, Figure ??, and Figure ??, we show the following,

$$\begin{aligned} & \max_K \left(\|\tilde{\mathbf{E}}_{t+1}(K)\|_2 - \|\hat{\mathbf{E}}_{t+1}(K)\|_2 \right), \\ & \min_K \left(\|\tilde{\mathbf{E}}_{t+1}(K)\|_2 - \|\hat{\mathbf{E}}_{t+1}(K)\|_2 \right), \\ & \text{mean}_K \left(\|\tilde{\mathbf{E}}_{t+1}(K)\|_2 - \|\hat{\mathbf{E}}_{t+1}(K)\|_2 \right), \end{aligned}$$

respectively, against the number of iterations of the JOR and DICI-OR algorithm. Since all the three figures show that the max, min, and the mean of the difference of the spectral norm of the two error processes, $\|\tilde{\mathbf{E}}_{t+1}(K)\|_2 -$

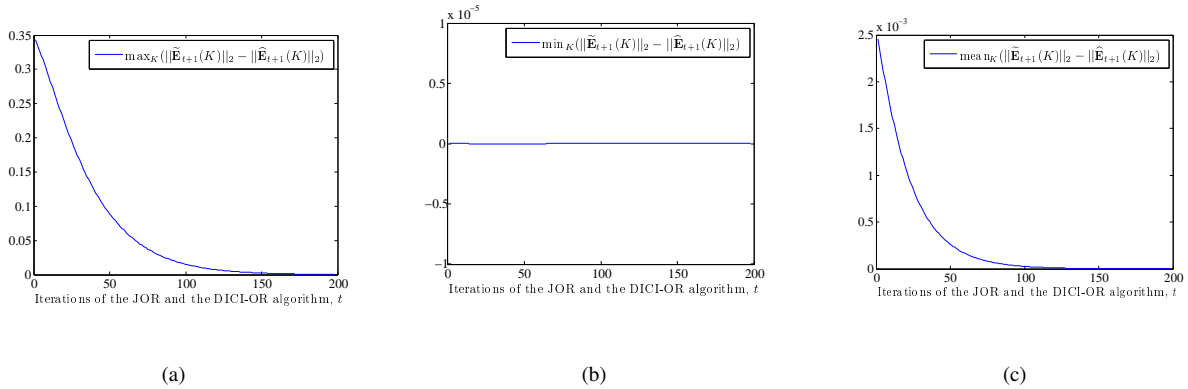


Fig. 7. Simulation for the error bound of the DICl-OR algorithm.

$\|\hat{\mathbf{E}}_{t+1}(K)\|_2$, is always ≥ 0 , we deduce that (using equation (56))

$$\begin{aligned} \|\hat{\mathbf{E}}_{t+1}\|_2 &\leq \|\tilde{\mathbf{E}}_{t+1}\|_2, \\ &\leq |\lambda_{\max}(\mathbf{P}_\gamma)|^{t+1} \|(\mathbf{S}_0 - \mathbf{S})\|_2. \end{aligned} \quad (73)$$

This verifies our claim numerically and provides us an upper bound on the spectral norm of the error process of the DICl algorithm.

VI. LOCAL INFORMATION FILTERS: INITIAL CONDITIONS AND LOCAL FILTER STEP

The initial conditions and the local filter step of the LIFs are presented in the next subsections.

A. Initial Conditions

The initial condition on the local predictor is

$$\hat{\mathbf{z}}_{0|-1}^{(l)} = \mathbf{0}. \quad (74)$$

Since the local information matrix and the local error covariances are not the inverse of each other, (49), we obtain the initial condition on the prediction information matrix by using the L -banded inversion theorem [20], provided in appendix I. This step may require a local communication step further elaborated in section VII-A.

$$\mathbf{Z}_{0|-1}^{(l)} \xleftarrow{L\text{-Banded Inversion Theorem}} \mathbf{S}_0^{(l)} \quad (75)$$

B. Local Filter Step

In this section, we present the local filter step of the LIFs. The local filter step is given by

$$\mathbf{Z}_{k|k}^{(l)} = \mathbf{Z}_{k|k-1}^{(l)} + \mathcal{I}_f^{(l)}, \quad (76a)$$

$$\hat{\mathbf{z}}_{k|k}^{(l)} = \hat{\mathbf{z}}_{k|k-1}^{(l)} + \mathbf{i}_{f,k}^{(l)}, \quad (76b)$$

where $\mathcal{I}_f^{(l)}$ and $\mathbf{i}_{f,k}^{(l)}$ denote the fused observation variables. Fusion of the observations is presented in section IV-A. The distribution of the addition operation, '+', in (26) is straightforward in (76). Recall that the observation fusion, (48), is carried out using the iterative weighted averaging algorithm. The asymptotic convergence of this iterative algorithm is guaranteed under certain conditions, e.g., connected sensor communication sub-graph, \mathcal{G}_j , see [15], for further details. Hence, with the required assumptions on the sub-graph, \mathcal{G}_j , the observation fusion algorithm, (48), asymptotically converges, and hence (with a slight abuse of notation),

$$\bigcup_{l=1}^N \mathbf{i}_{f,k}^{(l)} \rightarrow \mathbf{i}_k \text{ and } \bigcup_{l=1}^N \mathcal{I}_f^{(l)} \rightarrow \mathcal{I}. \quad (77)$$

The above notation implies that the local fused information variables, $\mathcal{I}_f^{(l)}$ and $\mathbf{i}_{f,k}^{(l)}$, when combined over the entire sensor network, asymptotically converge to the global information variables, \mathcal{I} and \mathbf{i}_k . This, in turn, implies that the local filter step of the LIFs asymptotically converges to the global filter step, (26), of the CLBIF.

Once the local filter step is completed, the DIC algorithm is employed on the local information matrices, $\mathbf{Z}_{k|k}^{(l)}$ obtained from (76a), to convert them into the local error covariance matrices, $\mathbf{S}_{k|k}^{(l)}$. Finally, to convert the estimates in the information domain, $\widehat{\mathbf{z}}_{k|k}^{(l)}$, to the estimates in the Kalman filter domain, $\widehat{\mathbf{x}}_{k|k}^{(l)}$, we specialize the DIC algorithm to matrix-vector product (19).

VII. LOCAL INFORMATION FILTERS: LOCAL PREDICTION STEP

This section presents the distribution of the global prediction step, (27), into the local prediction step at each LIF. This section requires the results of the DIC algorithm for the L -banded matrices, introduced in section V.

A. Computing the local prediction information matrix, $\mathbf{Z}_{k|k-1}^{(l)}$

Because of the coupled local dynamics of the reduced sensor-based models, each sensor may require that some of the estimated states be communicated as internal inputs, $\widehat{\mathbf{d}}_{k|k}^{(l)}$, to its LIF, as shown in (43). These states are the directed edges into each cut-point set in Figure ???. Hence, the error associated to a local estimation procedure is also influenced by the error associated to the neighboring estimation procedure, from where the internal inputs are being communicated. This dependence is true for all sensors and is reflected in the local prediction error covariance matrix, $\mathbf{S}_{k|k-1}^{(l)}$, as it is a function of the global estimation error covariance matrix, $\mathbf{S}_{k-1|k-1}$. Equation (78) follows from (27a) after expanding (27a) for each diagonal submatrix, $\mathbf{S}_{k|k-1}^{(l)}$, in $\mathbf{S}_{k|k}$.

$$\mathbf{S}_{k|k-1}^{(l)} = \mathbf{F}_l \mathbf{S}_{k-1|k-1} \mathbf{F}_l^T + \mathbf{G}^{(l)} \mathbf{Q}^{(l)} \mathbf{G}^{(l)T}. \quad (78)$$

The matrix, $\mathbf{F}_l = \mathbf{T}_l \mathbf{F}$, (the matrix \mathbf{T}_l is introduced in (39)) is an $n_l \times n$ matrix, which relates the state vector, \mathbf{x}_k , to the local state vector, $\mathbf{x}_k^{(l)}$. Figure 2 shows that the matrix, \mathbf{F}_l , is further divided into $\mathbf{F}^{(l)}$ and $\mathbf{D}^{(l)}$. With this sub-division of \mathbf{F}_l , the first term on the right hand side of (78), $\mathbf{F}_l \mathbf{S}_{k-1|k-1} \mathbf{F}_l^T$, can be expanded, and (78) can be written as

$$\begin{aligned} \mathbf{S}_{k|k-1}^{(l)} &= \mathbf{F}^{(l)} \mathbf{S}_{k-1|k-1} \mathbf{F}^{(l)T} + \mathbf{F}^{(l)} \mathbf{S}_{k-1|k-1}^{x^{(l)d^{(l)}}} \mathbf{D}^{(l)T} \\ &+ \left(\mathbf{F}^{(l)} \mathbf{S}_{k-1|k-1}^{x^{(l)d^{(l)}}} \mathbf{D}^{(l)T} \right)^T + \mathbf{D}^{(l)} \mathbf{S}_{k-1|k-1}^{d^{(l)d^{(l)}}} \mathbf{D}^{(l)T} + \mathbf{G}^{(l)} \mathbf{Q}^{(l)} \mathbf{G}^{(l)T}, \end{aligned} \quad (79)$$

where

$\mathbf{S}_{k-1|k-1}^{(l)}$ is the local error covariance matrix, which is available from (76a) and the DIC1 algorithm at sensor l ;
 $\mathbf{S}_{k-1|k-1}^{d^{(l)}d^{(l)}}$ is the local error covariance matrix, which is available from (76a) and the DIC1 algorithm at the sensors having the states, $\mathbf{d}_k^{(l)}$, in their reduced models;

$\mathbf{S}_{k-1|k-1}^{x^{(l)}d^{(l)}}$ is the error cross correlation between the local state vector, $\mathbf{x}_k^{(l)}$, and the local internal input vector, $\mathbf{d}_k^{(l)}$.

The non L -banded entries in this matrix can be computed from the equation (64), see [35]. Since the model matrix, \mathbf{F} , is sparse, we do not need the entire error covariance matrix, $\mathbf{S}_{k-1|k-1}$, only certain of its submatrices. Since the model matrix, \mathbf{F} , is localized, long-distance communication is not required, and the submatrices are available at the neighboring sensors.

Once we have calculated the local prediction error covariance matrix, $\mathbf{S}_{k|k-1}^{(l)}$, we realize (49) and compute the local prediction information matrix, $\mathbf{Z}_{k|k-1}^{(l)}$, using the L -banded Inversion Theorem (see [20] and appendix I).

$$\mathbf{Z}_{k|k-1}^{(l)} \xleftarrow{\text{L-Banded Inversion Theorem}} \mathbf{S}_{k|k-1}^{(l)}. \quad (80)$$

From (86) in appendix I, to calculate the local prediction information matrix, $\mathbf{Z}_{k|k-1}^{(l)}$, we only need the $\mathbf{S}_{k|k-1}^{(l)}$ from sensor ' l ' and from some additional neighboring sensors. Hence $\mathbf{Z}_{k|k-1}^{(l)}$ is again computed with only local communication and n_l th order computation.

B. Computing the local predictor, $\hat{\mathbf{z}}_{k|k-1}^{(l)}$

We illustrate the computation of the local predictor, $\hat{\mathbf{z}}_{k|k-1}^{(3)}$, for the 5-dimensional system, (30)–(31), with $L = 1$. The local predictor, $\hat{\mathbf{z}}_{k|k-1}^{(3)}$, at sensor 3 follows from the global predictor, (27b), and is given by

$$\hat{\mathbf{z}}_{k|k-1}^{(3)} = \mathbf{Z}_{k|k-1}^{(3)} \left(\mathbf{F}^{(3)} \hat{\mathbf{x}}_{k-1|k-1}^{(3)} + \mathbf{D}^{(3)} \hat{\mathbf{d}}_{k-1|k-1}^{(3)} \right) + \begin{bmatrix} z_{34} \left(f_{31} \hat{x}_{1,k-1|k-1}^{(1)} + f_{33} \hat{x}_{3,k-1|k-1}^{(2)} \right) \\ 0 \end{bmatrix}, \quad (81)$$

where z_{34} is the only term arising due to the $L = 1$ -banded (tridiagonal) assumption on the prediction information matrix, $\mathbf{Z}_{k|k-1}$. Note that $f_{31} \hat{x}_{1,k-1|k-1}^{(1)} + f_{33} \hat{x}_{3,k-1|k-1}^{(2)}$ is a result of $\mathbf{f}_3 \hat{\mathbf{x}}_{k-1|k-1}$, where \mathbf{f}_3 is the third row of the model matrix, \mathbf{F} . A model matrix with a localized and sparse structure ensures that $\mathbf{f}_3 \hat{\mathbf{x}}_{k-1|k-1}$ is computed from a small subset of the estimated state vector, $\hat{\mathbf{x}}_{k-1|k-1}^{(\mathcal{Q})}$, communicated by a subset $\mathcal{Q} \subseteq \mathcal{K}(l)$ of the neighboring sensors, which are modeling these states in their reduced models. This may require multi-hop communication.

Generalizing, the local predictor in the information domain, $\hat{\mathbf{z}}_{k|k-1}^{(l)}$, is given by

$$\hat{\mathbf{z}}_{k|k-1}^{(l)} = \mathbf{Z}_{k|k-1}^{(l)} \left(\mathbf{F}^{(l)} \hat{\mathbf{x}}_{k-1|k-1}^{(l)} + \mathbf{D}^{(l)} \hat{\mathbf{d}}_{k-1|k-1}^{(l)} \right) + f_1 \left(\mathbf{Z}_{k|k-1}^{(\mathcal{V})}, \mathbf{F}^{(\mathcal{V})}, \hat{\mathbf{x}}_{k-1|k-1}^{(\mathcal{Q})} \right) \quad (82)$$

for some $\mathcal{V}, \mathcal{Q} \subseteq \mathcal{K}(l)$, where $f_1(\cdot)$ is a linear function and depends on L .

C. Estimate Fusion

We present the following fact.

Fact: Let m denote the iterations of the consensus algorithm that is employed to fuse the observations. As $m \rightarrow \infty$, the local estimates, $\hat{\mathbf{z}}_{k|k}^{(l)}$, in (76b) also reach a consensus on the estimates of the shared states.

It is straightforward to note that as $m \rightarrow \infty$ we have a consensus on the estimates (of the shared states) in the local filter step (76b), if we have a consensus on the local predictors (of the shared states), $\widehat{\mathbf{z}}_{k|k-1}$. To show the consensus on the local predictors, we refer back to our illustration and write the local predictors for sensor 2 as follows,

$$\widehat{\mathbf{z}}_{k|k-1}^{(2)} = \mathbf{Z}_{k|k-1}^{(2)} \left(\mathbf{F}^{(2)} \widehat{\mathbf{x}}_{k-1|k-1}^{(2)} + \mathbf{D}^{(2)} \widehat{\mathbf{d}}_{k-1|k-1}^{(2)} \right) + \begin{bmatrix} z_{12} \left(f_{11} \widehat{x}_{1,k-1|k-1}^{(1)} + f_{12} \widehat{x}_{2,k-1|k-1}^{(2)} \right) \\ 0 \\ z_{45} \left(f_{54} \widehat{x}_{4,k-1|k-1}^{(2)} + f_{55} \widehat{x}_{5,k-1|k-1}^{(3)} \right) \end{bmatrix}. \quad (83)$$

The predictor for the shared state $x_{4,k}$ can now be extracted from (81) and (83) and can be verified to be the following for $l = 2, 3$.

$$\widehat{z}_{4,k|k-1}^{(l)} = z_{34} f_{31} \widehat{x}_{1,k-1|k-1}^{(1)} + (z_{34} f_{33} + z_{44} f_{43}) \widehat{x}_{3,k-1|k-1}^{(2)} + z_{45} f_{54} \widehat{x}_{4,k-1|k-1}^{(2)} + (z_{44} f_{45} + z_{45} f_{55}) \widehat{x}_{5,k-1|k-1}^{(3)} \quad (84)$$

The elements z_{ij} belong to the prediction information matrix, which is computed using the DIC1 algorithm and the L -banded inversion theorem. It is noteworthy that the DIC1 algorithm is not a consensus algorithm and thus the elements z_{ij} are the same across the sensor network at any iteration of the DIC1 algorithm. With a consensus on the local predictors, the iterations of the consensus algorithm on the observations lead to a consensus on the shared estimates.

VIII. RESULTS

A. Summary of the LIFs

We summarize the distributed local Information filters. The initial conditions are given by (74) and (75). Observation fusion is carried out using (48). The fused observation variables, $\mathbf{i}_{f,k}^{(l)}$ and $\mathbf{I}_{f,k}^{(l)}$, are then employed in the local filter step, (76a) and (76b), to obtain the local information matrix and the local estimator, $\widehat{\mathbf{Z}}_{k|k}^{(l)}$ and $\mathbf{z}_{k|k}^{(l)}$, respectively. We then implement the DIC1 algorithm (61)–(62) and (63) to compute the local error covariance matrix, $\mathbf{S}_{k|k}^{(l)}$, from the local information matrix, $\mathbf{Z}_{k|k}^{(l)}$. The DIC1 algorithm is again employed to compute the local estimates in the Kalman filter domain, $\widehat{\mathbf{x}}_{k|k}^{(l)}$, from the local estimator, $\widehat{\mathbf{z}}_{k|k}^{(l)}$, as a special case. Finally the local prediction step is completed by computing the local prediction error covariance matrix, $\widehat{\mathbf{S}}_{k|k-1}^{(l)}$, the local prediction information matrix, $\widehat{\mathbf{Z}}_{k|k-1}^{(l)}$, and, the local predictor, $\widehat{\mathbf{z}}_{k|k-1}^{(l)}$, from (79), (80), and (82), respectively.

B. Simulations

We simulate a $n = 100$ -dimensional system with $N = 10$ sensors monitoring the system. Figures ?? and ?? show the non-zero elements (chosen at random) of the model matrix, \mathbf{F} , such that $\|\mathbf{F}\|_2 = 1$. The model matrix in Figure ?? is $L = 20$ -banded. The model matrix in Figure ?? is $L = 36$ -banded that is obtained by employing the reverse Cuthill-McKee algorithm [33] for bandwidth reduction of a sparse random \mathbf{F} . The non-zeros (chosen at random as $\text{Normal}(0, 1)$) of the global observation matrix, \mathbf{H} , are shown in Figure 9. The l th row of the global observation matrix, \mathbf{H} , is the local observation matrix, \mathbf{H}_l , at sensor l . Distributed Kalman filters are implemented on (i) \mathbf{F} in ?? and \mathbf{H} in Figure 9; and (ii) \mathbf{F} in ?? and \mathbf{H} in Figure 9. The trace of the error covariance matrix,




Fig. 8. (a & b) Non-zero elements (chosen at random) of 100×100 , $L = 20$ -banded (Figure ??) and $L = 36$ -banded (Figure ??) model matrices, \mathbf{F} , such that $\|\mathbf{F}\|_2 = 1$. (c & d) Distributed Kalman filter is implemented on the model matrices in Figure ??-?? and the global observation matrix, \mathbf{H} (Figure 9), in Figure ??-??. The expectation operator in the trace (on horizontal axis) is simulated over 1000 Monte Carlo trials.




Fig. 9. Global observation matrix, \mathbf{H} . The non-zero elements (chosen at random) are shown. There are $N = 10$ sensors, where the l th row of \mathbf{H} corresponds to the local observation matrix, \mathbf{H}_l , at sensor l . The overlapping states (for which fusion is required) can be seen as the overlapping portion of the rows.




Fig. 10. Performance of the DICl algorithm as a function of the number of DICl iterations, t .

$\mathbf{S}_{k|k}$, is simulated for different values of L in $[n, 1, 2, 5, 10, 15, 20]$ and the plots are shown (after averaging over 1000 Monte Carlo trials) in Figure ?? for case (i); and in Figure ?? for case (ii). The stopping criteria for the DICl algorithm and the consensus algorithm are such that the deviation in their last 10 iterations is less than 10^{-5} . In both Figure ?? and Figure ??, $\text{tr}(\mathbf{S}_{k|k})$ represents the trace of the solution of the Riccati equation in the CIF (no approximation).

With 1000 Monte Carlo trials, we further simulate the trace of the error covariance, $\text{tr}(\mathbf{S}_{k|k})$, for case (ii) and $L = 20$ -banded approximations as a function of the number of iterations, t , of the DICl-OR algorithm. We compare this with (a) the simulation obtained from the $O(n^3)$ direct inverse of the the error covariance (with $L = 20$ -banded approximation on its inverse); and (b) $\text{tr}(\mathbf{S}_{k|k})$, trace of the solution of the Riccati equation of the CIF (no approximation). We choose $t = [1, 10, 30, 100, 200]$ for the DICl algorithm and show the results in Figure 10. As $t \uparrow$, the curves we obtain from the DICl algorithm get closer to the curve we obtain with the direct inverse.

The simulations confirm the following: (i) The LIFs asymptotically track the results of the CLBIF, see Figure 10; (ii) We verify that as $L \uparrow$, the performance is virtually indistinguishable from that of the CIF, as pointed out in [35]; this is in agreement with the fact that the approximation is optimal in Kullback-Leibler sense, as shown in [20]. Here, we also point out that, as we increase L the performance increases, but, we pay a price in terms of the communication cost, as we may have to communicate in a larger neighborhood. (iii) If we increase the number of Monte Carlo trials the variations reduce, and the filters eventually follow the solution of the Riccati equation, $\text{tr}(\mathbf{S}_{k|k})$; (iv) The curve in Figure 10 with $t = 1$ shows the decoupled LIFs, when the global error covariance is treated as a block diagonal matrix. This is unstable as we do not fuse the error covariances and the individual sub-systems are not observable. We next discuss the computational advantage of the LIFs over some of the existing methods.

C. Complexity

We regard the multiplication of two $n \times n$ matrices as an $O(n^3)$ operation, inversion of an $n \times n$ matrix also as an $O(n^3)$ operation and multiplication of an $n \times n$ matrix and an $n \times 1$ vector as an $O(n^2)$ operation. For all of the following, we assume N sensors monitoring the global system, (10).

1) *Centralized Information Filter, CIF*: This is the case where each sensor sends its local observation vector to a centralized location or a fusion center, where the global observation vector is then put together. The fusion center then implements the CIF, with an $O(n^3)$ computation complexity for each time k , so we have the complexity as $O(n^3k)$, with inordinate communication requirements (back and forth communication between the sensors and the central location).

2) *Information Filters with Replicated Models at each Sensor and Observation Fusion*: In this scheme, the local observation vectors are not communicated to the central location, but are fused over an all-to-all communication network, [3], or an arbitrary network, [7], that requires an iterative algorithm. Computational complexity at each sensor is $O(n^3k)$ in [3]. In [7], let the complexity of the iterative algorithm be $O(\varphi(n))$ and let t_φ be the number of iterations required by the iterative algorithm to converge. Each sensor implements a CIF after fusing the observations. For each time index, k , each sensor requires $O(n^3)$ operations plus the operations required for the iterative algorithms, which are $O(\varphi(n)t_\varphi)$; so the total computation complexity is $O((n^3 + \varphi(n)t_\varphi)k)$ at each sensor. Communication requirements are global in [3], and local in [7].

3) *Distributed Kalman Filters: Local Information Filters, LIFs*: The distributed Kalman filter presented in this paper has three iterative algorithms. In all the other steps the computation complexity is dominated by $O(n_l^3)$, where $n_l \ll n$. Let t_o be the iterations required by the weighted averaging algorithm, where at each step of the iterative algorithm the computations are dominated by $O(n_l^2)$. Let t_{J_1} be the iterations required by the DIC1 algorithm for vectors, where at each step of the iterative algorithm the computations are dominated by an $O(L^2)$ operations. Let t_{J_2} be the iterations required by the DIC1 algorithm, where at each step of the iterative algorithm the computations are dominated by $O(L^4)$ operations. Recalling that $L \sim n_l$ from section III, the total computation complexity is $O((n_l^3 + n_l^2t_o + n_l^2t_{J_1} + n_l^4t_{J_2})k)$. Let $t_\infty = \max(t_o, t_{J_1}, t_{J_2})$, then the computation complexity is bounded by $O(n_l^4t_\infty k)$ at each sensor for the LIFs, which is much smaller than the computational cost of the solutions in VIII-C.1 and VIII-C.2. The communication requirement in the LIFs may be multi-hop but is always constrained in a neighborhood because of the structural assumptions on the model matrix, \mathbf{F} .

IX. CONCLUSIONS

In conclusion, we presented a *distributed* implementation of the Kalman filter for sparse large-scale systems monitored by sensor networks. In our solution, the communication, computing, and storage is local and distributed across the sensor network, no single sensor processes n -dimensional vectors or matrices, where n , usually a large number, is the dimension of the state vector representing the random field. We achieve this by solving three linked problems: (1) *Distributing* the global state variable model representing the random field into sensor based reduced-order models. These reduced-order models are obtained using a graph-theoretic model distribution

REFERENCES

- [1] R. Kalman, "A new approach to linear filtering and prediction problems," *Trans. of the ASME - Journal of Basic Engineering*, vol. 82, no. 2, pp. 35–45, 1960.
- [2] R. Kalman and R. Bucy, "New results in linear filtering and prediction theory," *ASME Journal of Basic Engineering*, vol. 83, pp. 95–108, 1961.
- [3] B. Rao and H. Durrant-Whyte, "Fully decentralized algorithm for multisensor Kalman filtering," *IEE Proceedings-Control Theory and Applications*, vol. 138, pp. 413–420, Sep. 1991.
- [4] V. Saligrama and D. Castanon, "Reliable distributed estimation with intermittent communications," in *45th IEEE Conference on Decision and Control*, San Diego, CA, Dec. 2006, pp. 6763–6768.
- [5] T. Chung, V. Gupta, J. Burdick, and R. Murray, "On a decentralized active sensing strategy using mobile sensor platforms in a network," in *43rd IEEE Conference on Decision and Control*, Paradise Island, Bahamas, Dec. 2004, vol. 2, pp. 1914–1919.
- [6] H. Hashemipour, S. Roy, and A. Laub, "Decentralized structures for parallel Kalman filtering," *IEEE Trans. on Automatic Control*, vol. 33, no. 1, pp. 88–94, Jan. 1988.
- [7] R. Olfati-Saber, "Distributed Kalman filters with embedded consensus filters," in *44th IEEE Conference on Decision and Control*, Seville, Spain, Dec. 2005, pp. 8179 – 8184.
- [8] R. Olfati-Saber and J. Shamma, "Consensus filters for sensor networks and distributed sensor fusion," in *44th IEEE Conference on Decision and Control*, Seville, Spain, Dec. 2005, pp. 6698 – 6703.
- [9] B. Sinopoli, L. Schenato, M. Franceschetti, K. Poolla, M. Jordan, and S. Sastry, "Kalman filter with intermittent observations," *IEEE Trans. on Automatic Control*, vol. 49, no. 9, pp. 1453–1464, Sep. 2004.
- [10] T. Berg and H. Durrant-Whyte, "Model distribution in decentralized multi-sensor data fusion," Tech. Rep., University of Oxford, 1990.
- [11] A. Mutambara, *Decentralized estimation and control for multisensor systems*, CRC Press, Boca Raton, FL, 1998.
- [12] A. R. M. Noton, "Two-level form of the kalman filter," *IEEE Transactions on Automatic Control*, vol. 16, no. 2, pp. 128–133, Apr. 1971.
- [13] D. D. Šiljak, *Decentralized control of complex systems*, Academic Press Inc., Boston, MA, 1991.
- [14] Usman A. Khan and José M. F. Moura, "Distributed Kalman filters in sensor networks: Bipartite fusion graphs," in *15th IEEE Workshop on Statistical Sig. Proc.*, Madison, WI, Aug. 26–29 2007, pp. 700–704.
- [15] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," *Systems and Controls Letters*, vol. 53, no. 1, pp. 65–78, Apr. 2004.
- [16] N. Balram and José M. F. Moura, "Noncausal Gauss Markov random fields: Parameter structure and estimation," *IEEE Trans. on Information Theory*, vol. 39, no. 4, pp. 1333–1355, Jul. 1993.
- [17] Usman A. Khan and José M. F. Moura, "Distributed Iterate-Collapse inversion (DICI) algorithm for L -banded matrices," in *33rd International Conference on Acoustics, Speech, and Signal Processing*, Las Vegas, NV, Mar. 30–Apr. 04 2008, accepted for publication.
- [18] B. Anderson and J. Moore, *Optimal filtering*, Prentice Hall, Englewood Cliffs, NJ, 1979.
- [19] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and distributed computation*, Prentice Hall, Englewood Cliffs, NJ, 1989.
- [20] A. Kavcic and José M. F. Moura, "Matrices with banded inverses: Inversion algorithms and factorization of Gauss-Markov processes," *IEEE Trans. on Information Theory*, vol. 46, no. 4, pp. 1495–1509, Jul. 2000.
- [21] R. Nash Jr. and S. Jordan, "Statistical geodesy: An engineering perspective," *Proceedings of the IEEE*, vol. 66, no. 5, pp. 532–550, May 1978.
- [22] W. Ledsham and D. Staelin, "An extended Kalman-Bucy filter for atmospheric temperature profile retrieval with a passive microwave sounder," *Journal of Applied Meteorology*, vol. 17, pp. 1023–1033, Jul. 1978.
- [23] R. Brammer, "Estimation of the ocean geoid near the Blake Escarpment using Geos-3 satellite altimetry data," *Journal of Geophysical Research*, vol. 84, no. B8, pp. 3843–3852, Jul. 1979.
- [24] J. Galantowicz, D. Entekhabi, and E. Njoku, "Tests of sequential data assimilation for retrieving profile soil moisture and temperature from observed L-band radiobrightness," *IEEE Trans. on Geoscience and Remote Sensing*, vol. 37, no. 4, pp. 1860–1870, Jul. 1999.
- [25] M. Buehner and P. Malanotte-Rizzoli, "Reduced-rank Kalman filters applied to an idealized model of the wind-driven ocean circulation," *Journal of Geophysical Research*, vol. 108, no. C6, pp. 3192, 2003.
- [26] M. Buehner, P. Malanotte-Rizzoli, A. Busalacchi, and T. Inui, "Estimation of the tropical Atlantic circulation from altimetry data using a reduced-rank stationary Kalman filter," *Interhemispheric water exchanges in the Atlantic Ocean, Elsevier Oceanographic Series*, vol. 68, no. 9, pp. 49–92, 2003.

- [27] A. Graham, *Kronecker Products and Matrix Calculus with Applications*, Ellis Horwood, Chichester, UK, 1981.
- [28] J. Brailean and A. Katsaggelos, "Simultaneous recursive displacement estimation and restoration of noisy-blurred image sequences," *IEEE Trans. on Image Processing*, vol. 4, no. 9, pp. 1236–1251, Sep. 1995.
- [29] F. Khellah, P. Fieguth, M. Murray, and M. Allen, "Statistical processing of large image sequences," *IEEE Trans. on Image Processing*, vol. 14, no. 1, pp. 80–93, Jan. 2005.
- [30] A. Bergen and D. Hill, "A structure preserving model for power system stability analysis," *IEEE Trans. on Power Apparatus and Systems*, vol. PAS-100, no. 1, pp. 25–35, Jan. 1981.
- [31] M. Ilic, E. Allen, J. Chapman, C. King, J. Lang, and E. Litvinov, "Preventing future blackouts by means of enhanced electric power systems control: From complexity to order," *Proceedings of the IEEE*, vol. 93, no. 11, pp. 1920–1941, Nov. 2005.
- [32] Marija D. Ilić, Le Xie, Usman A. Khan, and José M. F. Moura, "Modelling future cyber-physical energy systems," in *IEEE Power Engineering Society General Meeting*, Pittsburgh, PA, Jul. 2008, submitted.
- [33] E. Cuthill J. McKee, "Reducing the bandwidth of sparse symmetric matrices," in *Proceedings of the 24th National Conference*, New York, 1969, pp. 157–172.
- [34] T. Chin, W. Karl, and A. Willsky, "Sequential filtering for multi-frame visual reconstruction," *Special Issue on Multidimensional Signal Processing*, vol. 28, no. 3, pp. 311–333, 1992.
- [35] A. Asif and José M. F. Moura, "Inversion of block matrices with L-block banded inverse," *IEEE Trans. on Sig. Proc.*, vol. 53, no. 2, pp. 630–642, Feb. 2005.
- [36] H. Zhang, José M. F. Moura, and B. Krogh, "Estimation in sensor networks: A graph approach," in *4th International Symposium on Information Processing in Sensor Networks*, Los Angeles, CA, Apr. 2005, pp. 203–209.
- [37] Béla Bollobás, *Modern Graph Theory*, Springer, New York, NY, 1998.
- [38] R. Carli, A. Chiuso, L. Schenato, and S. Zampieri, "Distributed kalman filtering based on consensus strategies," Tech. Rep., Department of Information Engineering, University of Padova, 2007.
- [39] L. Xiao and S. Boyd, "Optimally control of spatially distributed systems," in *Workshop on large scale nonlinear and semidefinite programming*, Waterloo, Canada, May 2004.
- [40] S. Aldosari and José M. F. Moura, "Distributed detection in sensor networks: Connectivity graph and small world networks," in *39th Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, Nov. 2005, pp. 230–234.
- [41] S. Kar and José M. F. Moura, "Topology for global average consensus," in *40th Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, Oct. 2006.
- [42] S. Kar and José M. F. Moura, "Sensor networks with random links: Topology design for distributed consensus," Submitted to IEEE Trans. on Signal Processing, http://arxiv.org/PS_cache/arxiv/pdf/0704/0704.0954v1.pdf, Apr. 2007.
- [43] J. Garloff, "Block methods for the solution of linear interval equations," *SIAM Journal on Matrix Analysis and Applications*, vol. 11, no. 1, pp. 89–106, Jan. 1990.
- [44] A. I. Zečević and D. D. Šiljak, "A block-parallel newton method via overlapping epsilon decompositions," *SIAM Journal on Matrix Analysis and Applications*, vol. 15, no. 3, pp. 824–844, 1994.
- [45] D. D. Šiljak and A. I. Zečević, "Overlapping block-iterative methods for solving algebraic equations," *Journal of Difference Equations and Applications*, vol. 1, pp. 125–136, 1995.
- [46] V. Delouille, R. Neelamani, and R. Baraniuk, "Robust distributed estimation using the embedded subgraphs algorithm," *IEEE Trans. on Sig. Proc.*, vol. 54, pp. 2998–3010, Aug. 2006.
- [47] D. M. Stipanovic and D. D. Šiljak, "Jacobi and gauss-seidel iterations for polytopic systems: convergence via convex m-matrices," *Reliable Computing*, vol. 6, no. 2, pp. 123–137, May 2000.
- [48] B. Baran, E. Kaszkurewicz, and A. Bhaya, "Parallel asynchronous team algorithms: convergence and performance analysis," *IEEE Transactions on Parallel and Distributed Systems*, vol. 7, pp. 677–688, Jul. 1996.