Automatica 93 (2018) 75–82

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

A distributed Kalman filter with event-triggered communication and guaranteed stability[☆]

Giorgio Battistelli [a,*], Luigi Chisci [a], Daniela Selvi [b]

[a] Università di Firenze, Dipartimento di Ingegneria dell'Informazione (DINFO), Via di Santa Marta 3, 50139 Firenze, Italy
[b] IMT School for Advanced Studies Lucca, 55100 Lucca, Italy

ARTICLE INFO

Article history:
Received 23 June 2016
Received in revised form 8 September 2017
Accepted 29 December 2017

Keywords:
Distributed Kalman filtering
Sensor networks
Event-triggered communication
Sensor fusion

ABSTRACT

The paper addresses Kalman filtering over a peer-to-peer sensor network with a careful eye towards data transmission scheduling for reduced communication bandwidth and, consequently, enhanced energy efficiency and prolonged network lifetime. A novel consensus Kalman filter algorithm with event-triggered communication is developed by enforcing each node to transmit its local information to the neighbors only when this is considered as particularly significant for estimation purposes, in the sense that it notably deviates from the information that can be predicted from the last transmitted one. Further, it is proved how the filter guarantees stability (mean-square boundedness of the estimation error in each node) under network connectivity and system collective observability. Finally, numerical simulations are provided to demonstrate practical effectiveness of the distributed filter for trading off estimation performance versus transmission rate.

1. Introduction

Nowadays, *wireless sensor networks* (WSNs) are getting an ever increasing usage in a wide range of on-line monitoring tasks (e.g. navigation, tracking, environmental and power system monitoring, etc.) that require recursive estimation of the state of a linear or nonlinear dynamical system. Since the individual nodes of the sensor network are usually low-cost, battery-supplied devices with scarce energy resources, it becomes of paramount importance for networked state estimation to limit as much as possible data transmission which represents by far the most energy consuming node task.

In this respect, simple ways to limit the communication bandwidth are *periodic* and *random* transmission at a prescribed rate, whose effects on distributed Kalman filter stability and estimation performance are analyzed in Battistelli, Benavoli, and Chisci (2012b) for a *centralized network* wherein all nodes transmit their local data (either measurements or estimates) to a fusion center. It is natural to expect, however, that a data-driven transmission strategy could easily outperform periodic and random scheduling. These considerations motivate the growing interest towards the development of data-driven (or event-triggered) strategies for scheduling data communication. Interested readers are referred to Battistelli, Benavoli, and Chisci (2012a), Han et al. (2015), Marck and Sijs (2010), Shi, Chen, and Shi (2014), Shi, Johansson, and Qiu (2011), Shi, Shi, and Chen (2016), Sijs, Kester, and Noack (2014), Suh, Nguyen, and Ro (2007), Trimpe and D'Andrea (2014), and references therein, for an overview on the stability properties and performance achievable by these strategies on a centralized network. A great deal of work (Battistelli & Chisci, 2014; Battistelli, Chisci, Mugnai, Farina, & Graziano, 2015; Farina, Ferrari-Trecate, & Scattolini, 2010; Kamal, Farrell, & Roy-Chowdhury, 2013; Noack, Sijs, Reinhardt, & Hanebeck, 2016; Olfati-Saber, 2009; Stankovic, Stankovic, & Stipanovic, 2009; Ugrinovskii, 2013) has concerned distributed state estimation over a *peer-to-peer network* wherein there is no fusion center and each node (peer) operates in the same way and can only exchange data with a limited subset of neighbors. All these references, however, have considered the situation wherein each node broadcasts data to neighbors after each update of the local information. In this respect, recent work (Li, Jia, & Du, 2016; Li, Zhu, Chen, & Guan, 2012; Liu, Wang, He, & Zhou, 2015; Meng & Chen, 2014; Wu, Guo, & Yang, 2015; Yan, Zhang, Zhang, & Yang, 2013) has addressed distributed state estimation with event-triggered communication. In particular Liu et al. (2015) and Yan et al. (2013) proposed measurement-based transmission tests on the distance between the current and latest transmitted measurements and, respectively, on the innovation. Conversely Li et al. (2016, 2012), Meng and Chen (2014) and Wu et al. (2015)

[☆] The material in this paper was partially presented at the 19th International Conference on Information Fusion, 5–8 July 2016, Heidelberg, Germany. This paper was recommended for publication in revised form by Associate Editor Hideaki Ishii under the direction of Editor Christos G. Cassandras.

* Corresponding author.
E-mail addresses: giorgio.battistelli@unifi.it (G. Battistelli), luigi.chisci@unifi.it (L. Chisci), daniela.selvi@imtlucca.it (D. Selvi).

https://doi.org/10.1016/j.automatica.2018.03.005
0005-1098/

developed event-triggered distributed state estimators all relying on the consensus Kalman filter of Olfati-Saber (2009) but differing for the adopted triggering condition. This paper presents a novel event-triggered distributed state estimator based on a different consensus Kalman filtering approach (Battistelli & Chisci, 2014; Battistelli et al., 2015) as well as on a different transmission triggering condition which essentially requires that the local estimate and/or covariance of a given node be sufficiently far away from the ones that could be computed by neighbors, exploiting only the transmitted data. It is proved that the proposed distributed Kalman filter algorithm with event-triggered communication enjoys nice stability properties (i.e., mean-square boundedness of the state estimation error in all nodes) under minimal requirements of network connectivity and collective system observability extending, in a non trivial way, the results of Battistelli and Chisci (2014, 2016) and Battistelli et al. (2015) already available for the full transmission case. This paper extends preliminary work carried out in Battistelli, Chisci, and Selvi (2016) with the stability analysis.

2. Distributed state estimation setting

This paper addresses distributed state estimation (DSE) over a network in which each node can process local data as well as exchange data with neighbors. Further, some nodes (called sensor nodes) have also sensing capabilities, i.e., they can sense data from the environment. Notice that the presence of nodes without sensing capabilities serves only the purpose of improving network connectivity. In the sequel, the sensor network will be denoted as $(\mathcal{N}, \mathcal{A}, \mathcal{S})$ where: $\mathcal{N} = \{1, \ldots, N\}$ is the set of nodes; $\mathcal{A} \subseteq \mathcal{N} \times \mathcal{N}$ is the set of arcs (edges); $\mathcal{S} \subseteq \mathcal{N}$ is the subset of sensor nodes. A directed edge $(i, j) \in \mathcal{A}$ from node i to node j means that i can send messages to j, and we say that j is an out-neighbor of i and i is an in-neighbor of j. For each node $i \in \mathcal{N}$, $\mathcal{N}_i \subseteq \mathcal{N}$ will denote the set of its in-neighbors, i.e. $\mathcal{N}_i \stackrel{\triangle}{=} \{j : (j, i) \in \mathcal{A}\}$. The network does not contain self-loops so that $i \notin \mathcal{N}_i$.

The DSE problem can be formulated as follows. Each node $i \in \mathcal{N}$ must estimate at each time $k \in \mathbb{Z}^+ = \{0, 1, \ldots\}$ the state x_k of the dynamical system

$$x_{k+1} = Ax_k + w_k \tag{1}$$

given local measurements

$$y_k^i = C^i x_k + v_k^i , \qquad i \in \mathcal{S}, \tag{2}$$

and data received from all in-neighbors $j \in \mathcal{N}_i$.

It is assumed that w_k and v_k^i, $i \in \mathcal{S}$, are zero-mean white noises with positive definite covariance matrices Q and R^i, $i \in \mathcal{S}$, respectively. Further, the process disturbance and the measurement noises are supposed to be uncorrelated, i.e., $\mathbb{E}\{v_k^i \, w_\tau^\top\} = 0$, for any $k, \tau \in \mathbb{Z}^+$, and $i \in \mathcal{S}$.

In this setting, it was recently shown that there exist families of consensus-based DSE algorithms (Battistelli & Chisci, 2014; Battistelli et al., 2015; Kamal et al., 2013) able to guarantee stability of the estimation error in each network node under the minimal requirements of collective detectability and network connectivity. Generally speaking, these algorithms require that each node i transmits its local estimate $\hat{x}_{k|k}^i$ and covariance matrix $P_{k|k}^i$ to all its out-neighbors j such that $i \in \mathcal{N}_j$ at least once for each sampling interval. However, in many contexts, it is desirable to reduce data transmission as much as possible while preserving performance. The objective of this paper is precisely that of developing a strategy for controlling transmission in existing DSE algorithms, so that each node i selectively transmits only the most relevant data, without compromising stability properties.

To this end, let us introduce for each node i binary variables c_k^i such that $c_k^i = 1$ if node i transmits at time k or $c_k^i = 0$ otherwise. The focus is on data-driven (or event-triggered) transmission strategies in which the variable c_k^i is a function of the information currently available in node i and of the information most recently transmitted by node i.

3. Distributed Kalman-filtering with event-triggered communication

In this paper, we focus on a DSE algorithm wherein each node $i \in \mathcal{N}$ runs a local Kalman filter and then, in order to improve its local estimate, fuses the local information with the one received from its in-neighbors $j \in \mathcal{N}_i$. Concerning the local Kalman filter, it is convenient for the presentation of the algorithm to consider the *information form* of the Kalman filter recursion which, instead of the estimate $\hat{x}_{k|k}^i$ and of the covariance matrix $P_{k|k}^i$, propagates the *information matrix* $\Omega_{k|k}^i = (P_{k|k}^i)^{-1}$ and the *information vector* $q_{k|k}^i = \Omega_{k|k}^i \hat{x}_{k|k}^i$. Hereafter, the steps of the proposed DSE algorithm are described in some detail.

Correction: Let $(q_{k|k-1}^i, \Omega_{k|k-1}^i)$ denote the predicted *information pair* available in node i at time k. Then, for any sensor node $i \in \mathcal{S}$, the local information pair is updated by means of the standard Kalman filter correction step

$$q_{k|k}^i = q_{k|k-1}^i + (C^i)^\top (R^i)^{-1} y_k^i , \tag{3}$$

$$\Omega_{k|k}^i = \Omega_{k|k-1}^i + (C^i)^\top (R^i)^{-1} C^i . \tag{4}$$

In all the remaining nodes $i \in \mathcal{N} \setminus \mathcal{S}$, since no local measurement is available, we simply set $(q_{k|k}^i, \Omega_{k|k}^i) = (q_{k|k-1}^i, \Omega_{k|k-1}^i)$.

Information exchange: Notice preliminarily that, after the correction step, the currently available information is represented by the local posterior information pair $(q_{k|k}^i, \Omega_{k|k}^i)$. Let now n_k^i be the number of discrete time instants elapsed from the most recent transmission of node i, so that the most recently transmitted data is $(q_{k-n_k^i|k-n_k^i}^i, \Omega_{k-n_k^i|k-n_k^i}^i)$. Such data can be propagated in time by repeatedly applying the Kalman filter prediction step so as to obtain the information pair $(\bar{q}_k^i, \bar{\Omega}_k^i)$ (see Eqs. (13)–(14) in the prediction step below). Accordingly, $\bar{x}_k^i = (\bar{\Omega}_k^i)^{-1} \bar{q}_k^i$ represents a prediction of the system state based on the data most recently transmitted by node i.

Notice that $\left(\bar{q}_k^i, \bar{\Omega}_k^i\right)$ can be computed also by the out-neighbors of node i. Then, the idea is to selectively transmit only in case the discrepancy between $(q_{k|k}^i, \Omega_{k|k}^i)$ and $(\bar{q}_k^i, \bar{\Omega}_k^i)$ is large, which means that the data $(\bar{q}_k^i, \bar{\Omega}_k^i)$ currently computable by the out-neighbors of node i is no longer consistent with the data locally available in node i. More formally, the following event-triggered transmission strategy is adopted

$$c_k^i = \begin{cases} 0, & \text{if } \|\hat{x}_{k|k}^i - \bar{x}_k^i\|^2_{\Omega_{k|k}^i} \le \alpha \\ & \quad \text{and } \dfrac{1}{1+\beta}\, \Omega_{k|k}^i \le \bar{\Omega}_k^i \le (1+\delta)\, \Omega_{k|k}^i \\ 1, & \text{otherwise} \end{cases} \tag{5}$$

where α, β, and δ are positive scalars and, given a positive definite matrix M, $\| \cdot \|_M$ denotes the corresponding weighted Euclidean norm.

The three scalars α, β, and δ can be seen as design parameters which can be tuned so as to achieve a desired behavior in terms of transmission rate and performance. In particular, the transmission test in (5) is designed so as to ensure that, in the case of no transmission, the data $(\bar{q}_k^i, \bar{\Omega}_k^i)$ currently computable by the out-neighbors of node i are close to the data locally available in node i both in terms of mean and covariance.