Automatica 95 (2018) 63–72

Contents lists available at ScienceDirect

Automatica

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

A distributed Kalman filtering algorithm with fast finite-time convergence for sensor networks☆

Zongze Wu [a], Minyue Fu [b,a,*], Yong Xu [a], Renquan Lu [a]

[a] School of Automation, Guangdong University of Technology, and Guangdong Key Laboratory of IoT Information Technology, Guangzhou 510006, China
[b] School of Electrical Engineering and Computer Science, The University of Newcastle, NSW 2308, Australia

ARTICLE INFO

Article history:
Received 24 January 2017
Received in revised form 9 December 2017
Accepted 17 April 2018

Keywords:
Distributed estimation
Distributed Kalman filtering
Sensor networks
Maximum likelihood estimation
Weighted least-squares estimation
Distributed field estimation

ABSTRACT

This paper proposes a new distributed algorithm for Kalman filtering. It is assumed that a linear discrete-time dynamic system is monitored by a network of sensors with some being active and some idle. The goal of distributed state estimation is to devise a distributed algorithm such that each node can independently compute the optimal state estimate by using its local measurements and information exchange with its neighbours. The proposed algorithm applies to acyclic network graphs (i.e., tree graphs) with fast finite-time convergence, but is also applicable to cyclic graphs by combining it with a distributed loop removal algorithm. The proposed algorithm enjoys low complexities, robustness against transmission adversaries and asynchronous implementability. The proposed distributed algorithm also applies to maximum likelihood estimation and weighted least-squares estimation, as special cases. With simple modifications, the proposed algorithm also applies to an important problem in signal processing called distributed field estimation.

1. Introduction

Distributed Kalman filtering (DKF) has been an extremely active research topic for over a decade, in response to rapid development and vast deployment of low-cost sensors and sensor networks. The technical challenge is how to migrate the well-established central (or traditional) Kalman filtering (KF) approach (Anderson & Moore, 1979; Kalman, 1960) to complex large-scale dynamic systems with measurements distributed over a large geographical area (Khan & Moura, 2008). Available DKF algorithms are already abundant. For example, a one-step prediction algorithm was introduced in Zhou (2013); a distributed iterate-collapse inversion algorithm in conjunction with a bipartite fusion graph was introduced in Khan and Moura (2008) for spatially sparse systems; distributed fusion estimation was proposed in Chen, Zhang, and Yu (2014a, b) and Chen, Zhang, Yu, Hu, and Song (2015); DKF using quantized information were studied in Li, Kar, Alsaadi, Dobaie, and Cui (2015), Riberio, Giannakis, and Roumeliotis (2006) and Song, Yu, and Zhang (2014); a DKF design using the well-known gossip protocol was given in Li, Kar, Moura, Poor, and Cui (2015); DKF using diffusion strategies was studied in Cattivelli and Sayed (2010) and Hu, Xie, and Zhang (2012); DKF with out-of-sequence measurements was treated in Shen, Song, Zhu, and Luo (2009), and fusion-centre based DKF designs were shown in Song, Xu, and Zhu (2014) and Xu, Song, Luo, and Zhu (2012). See Mahmoud and Khalid (2013) for a recent survey on DKF. Related works also include distributed maximum likelihood estimation (Zhao & Nehorai, 2007) and distributed weighted least-squares estimation (Marelli & Fu, 2015).

Numerous applications of DKF have been reported in the literature, ranging from environmental monitoring to surveillance, detection, tracking and object classification. Target tracking using a sensor network over a large geographical area is an active research topic recently, and DKF has been shown to play an important role in this application (Medeiros, Park, & Kak, 2008; Zhou, Fang, & Hong, 2013). In fact, this line of research can be traced back to Durrant-Whyte and Rao (1991) in 1991 and Regazzoni (1994) in 1994. DKF also finds wide applications in plant-wide control systems (Vadigepalli & Doyle, 2003), stochastic nonlinear systems with communication delays and packet losses (Wang, Fang, & Liu, 2015), clock synchronization for sensor networks (Luo & Wu, 2013), wireless sensor networks (Riberio & Giannakis, 2006), spatial estimation (Cortés, 2009), power networks (Kanna, Dini, Xia, Hui, & Mandic, 2015; Roshany-Yamchi et al., 2013; Sun, Fu, Wang, & Zhang, 2015; Tai, Lin, Fu, & Sun, 2013). We will also show in this paper a novel

☆ The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Bert Tanner under the direction of Editor Christos G. Cassandras.

* Corresponding author at: School of Electrical Engineering and Computer Science, The University of Newcastle, NSW 2308, Australia.
E-mail addresses: zzwu@scut.edu.cn (Z. Wu), minyue.fu@newcastle.edu.au (M. Fu), yxu@hdu.edu.cn (Y. Xu), rqlu@gdut.edu.cn (R. Lu).

https://doi.org/10.1016/j.automatica.2018.05.012

application of DKF in distributed field estimation where a sensor network is used to estimate the parameters of a physical field over a large geographical area, a problem with vast applications on its own (see, e.g., Martinez, 2010; Talarico, Schmid, Alkhweldi, & Valenti, 2014; Wang, Ishwar, & Saligrama, 2008).

One common approach to DKF is to use an average consensus strategy as introduced in Xiao and Boyd (2004); see Carli, Chiuso, Schenato, and Zampieri (2008), Das and Moura (2015), Kar and Moura (2011), Song, Yu et al. (2014), Xu et al. (2012), and Zhou et al. (2013) for examples of this approach. The main shortcomings of this approach include (1) usually only asymptotic convergence is guaranteed, meaning that an infinite number of iterations is required in theory; (2) stopping criteria are difficult to give for practical applications; (3) only sub-optimal estimates are usually given. Another common problem with many existing DKF algorithms is that a fusion centre is required, which means that they are not fully distributed. Approaches of this kind include Chen et al. (2014a, b, 2015), Shen et al. (2009), Song, Xu et al. (2014) and Xu et al. (2012). We also note that fusion centre is also commonly used in the so-called parallel Kalman filtering where local measurements are used to produce local estimates that are then fused together in a fusion centre; see, e.g., Hashemipour, Sumit, and Laub (1988). The need for fully distributed DKF with good optimality properties and low computational, communicational and storage complexities is urgent. These features are essential to make DKF scalable to large-scale sensor networks.

In this paper, we consider a sensor network used to detect, monitor and track targets within the geographical area covered by the network. Each target is modelled as a linear dynamic system, suitable for describing the motion of a moving target or changes in time-varying parameters. The goal of this paper is to devise a distributed algorithm that allows us to estimate the state of each target system. The algorithm needs to have low complexities per node (in terms of communication, computation and storage) so that it is *scalable* to large-scale sensor networks. For convenience, we mainly consider the tracking of a single target system, as the tracking of multiple targets can be achieved by a multiple number of single target trackers. We consider the scenario where the target system is measured by a small subset of active sensors at each time instant, while the rest of the sensors are idle. The set of active sensors is allowed to vary for each time. The use of idle nodes is motivated by applications where only a small fraction of sensors are able to measure information for a particular system. For example, a surveillance network may be responsible to monitor certain types of targets over a large geographical area. A moving target may be observed only by a small number of sensors around it, but the information about it may need to be distributed among the whole network so that the target can be tracked as it traverses within the network.

Following the standard (central) KF approach, we also divide the DKF problem into two steps, a maximum likelihood estimation (MLE) step and a one-step forward prediction step. The core technical issue is how to carry out the MLE step in a fully distributed manner. For this, we propose a fully distributed maximum likelihood estimation (DMLE) algorithm. Under the assumption that the communication graph for the sensor network is acyclic (i.e. it is a tree graph), the algorithm delivers the same (optimal) estimate as given by a central MLE algorithm, but with the advantage of fast convergence. That is, the proposed DMLE algorithm converges in a finite number of iterations (this number equals the maximum diameter of the graph). For sensor networks with a cyclic graph, we can apply a distributed depth-first-search (DFS) algorithm to convert the given graph to a spanning tree before applying the DMLE algorithm. We will show that the proposed DKF algorithm, which is based on the aforementioned DMLE algorithm, enjoys a number of nice properties, including low computational, communicational and storage complexities, robustness against transmission loss and delay and asynchronous implementability. The algorithm can run in either a point-to-point communication mode between neighbouring sensors (for better privacy) or in a broadcast mode (for lower communication burden).

We emphasize that the above target tracking setting is used to motivate the proposed DKF algorithm. Other applications of the proposed DKF algorithm include cascading failure monitoring in a power network, vehicle tracking in a transportation network, fire monitoring in a forest, etc. In particular, we will study the application to distributed field estimation in detail. To help illustrate the proposed DKF, two simulation examples will be given, one on single and multiple target tracking and one on distributed field estimation.

The rest of the paper is organized as follows: Section 2 formulates the DKF problem; Section 3 gives the proposed DKF algorithm; Section 4 details a number of properties of this algorithm; Section 5 discusses its application to distributed field estimation; Section 6 demonstrates the algorithm via some simulation examples; and Section 7 concludes the paper.

2. Problem formulation

Consider a dynamic model for a target system described by

$$x(t+1) = A(t)x(t) + w(t), \tag{1}$$

where $t = 0, 1, \ldots$ is the time index, $x(t) \in \mathcal{R}^n$ is the state, $w(t) \in \mathcal{R}^n$ is a zero-mean i.i.d. Gaussian noise with covariance $W(t) \geq 0$, and $A(t) \in \mathcal{R}^{n\times n}$ is the (possibly) time-varying transition matrix. The initial state $x(0)$ is an independent Gaussian variable with mean x_0 and covariance $P_0 > 0$.

The target system is measured by a network of sensors which can be represented by a graph $\mathcal{G}(t) = \{\mathcal{V}(t), \mathcal{E}(t)\}$ with a set of nodes $\mathcal{V}(t)$ and a set of edges $\mathcal{E}(t) \subset \{(i,j) : i \neq j, i, j \in \mathcal{V}(t)\}$. The set $\mathcal{V}(t) = \mathcal{I}(t) \cup \mathcal{S}(t)$, where $\mathcal{S}(t)$ is a set of sensing nodes with measurements and $\mathcal{I}(t)$ is an idle set consisting of nodes without measurements. Each node $i \in \mathcal{S}(t)$ has a measurement:

$$y_i(t) = C_i(t)x(t) + v_i(t), \tag{2}$$

where $y_i(t) \in \mathcal{R}^{r_i}$, $C_i(t) \in \mathcal{R}^{r_i\times n}$ is the (possibly) time-varying measurement vector, $v_i(t)$ is the measurement noise which is a zero-mean i.i.d. Gaussian noise with covariance $R_i(t) > 0$. Stacking up all the measurements, we get

$$y(t) = C(t)x(t) + v(t). \tag{3}$$

The covariance of $v(t)$ is $R(t) = \mathrm{diag}\{R_i(t)\}$. It is assumed throughout the paper that the system with (1) and (3) is observable.

We assume that $\mathcal{G}(t)$ is undirected and acyclic. A graph is called *undirected* if each edge is undirected. A graph is called *acyclic* if it is *connected* and has no loops, i.e., it is a *tree graph*. We will show how to deal with cyclic graphs later. Denote by $\mathcal{N}_i(t)$ the set of neighbouring nodes connected to node i, and denote by $|\mathcal{N}_i(t)|$ the cardinality of $\mathcal{N}_i(t)$. We assume that $|\mathcal{N}_i(t)| \ll |\mathcal{V}(t)|$ for each $i \in \mathcal{V}(t)$. We denote by $d(t)$ the diameter of the graph $\mathcal{G}(t)$, which is the length of the longest path between two nodes.

For notational convenience, we will suppress the time dependence of the system parameters $(A(t), C(t), W(t), R(t))$ as well as those of the network graph $(\mathcal{G}(t), \mathcal{V}(t), \mathcal{I}(t), \mathcal{S}(t), \mathcal{E}(t), \mathcal{N}_i(t), d(t))$. In addition, with some abuse of notation, $i \in \mathcal{G}$ means $\in \mathcal{V}$, and $(i,j) \in \mathcal{G}$ means $(i,j) \in \mathcal{E}$.

It is well known (Anderson & Moore, 1979) that the optimal state estimation, in the maximum likelihood sense, is given by the celebrated Kalman filter (KF), which we will call *central Kalman filter*. Given the system model (1) and the measurement model