Automatica 78 (2017) 202–210

Contents lists available at ScienceDirect

Automatica

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

Sensor selection for Kalman filtering of linear dynamical systems: Complexity, limitations and greedy algorithms[☆]

Haotian Zhang[a], Raid Ayoub[b], Shreyas Sundaram[c]

[a] Department of Electrical and Computer Engineering at the University of Waterloo, Waterloo, ON N2L 3G1, Canada
[b] Strategic CAD Labs, Intel Corporation, Hillsboro, OR 97124, USA
[c] School of Electrical and Computer Engineering at Purdue University, West Lafayette, IN 47907, USA

ARTICLE INFO

Article history:
Received 8 December 2015
Received in revised form
23 August 2016
Accepted 28 November 2016

Keywords:
Sensor selection
Kalman filters
Complexity
Greedy algorithms
Multi-sensor estimation

ABSTRACT

We consider the problem of selecting an optimal set of sensors to estimate the states of linear dynamical systems. Specifically, the goal is to choose (at design-time) a subset of sensors (satisfying certain budget constraints) from a given set in order to minimize the trace of the steady state *a priori* or *a posteriori* error covariance produced by a Kalman filter. We show that the *a priori* and *a posteriori* error covariance-based sensor selection problems are both NP-hard, even under the additional assumption that the system is stable. We then provide bounds on the worst-case performance of sensor selection algorithms based on the system dynamics, and show that greedy algorithms are optimal for a certain class of systems. However, as a negative result, we show that certain typical objective functions are not submodular or supermodular in general. While this makes it difficult to evaluate the performance of greedy algorithms for sensor selection (outside of certain special cases), we show via simulations that these greedy algorithms perform well in practice.

1. Introduction

One of the key problems in control system design is to select an appropriate set of actuators or sensors (either at design-time or at run-time) in order to achieve certain performance objectives (Van De Wal & De Jager, 2001). For the objective of estimating the state of a given linear Gauss–Markov system, there has been a growing literature in the past few years that studies how to dynamically select sensors at run-time to minimize certain metrics of the error covariance of the corresponding Kalman filter. This is known as the *sensor scheduling problem*, due to the fact that a different set of sensors can be chosen at each time-step (e.g., see Gupta, Chung, Hassibi, & Murray, 2006; Jawaid & Smith, 2015).

The design-time sensor selection problem (where the set of chosen sensors is not allowed to change over time) has been studied in various forms, including cases where the objective is to guarantee a certain structural property of the system (Pequito, Kar, & Aguiar, 2013), to optimize energy or information theoretic metrics (Krause, Singh, & Guestrin, 2008; Summers, Cortesi, & Lygeros, 2016), or to compute the optimal sensing matrix under a norm constraint (Belabbas, 2016).[1] Various sensor selection heuristics have also been proposed for estimation of *static* random variables (e.g., see Chepuri & Leus, 2015; Joshi & Boyd, 2009; Nordio, Tarable, Dabbene, & Tempo, 2015); however, the corresponding results do not directly translate to the case of estimating the (vector) state of dynamical systems.

In Dhingra, Jovanović, and Luo (2014), the authors studied the design-time actuator/sensor selection problem for continuous-time linear dynamical systems using the sparsity-promoting framework from Lin, Fardad, and Jovanovic (2013) and Polyak et al. (2013). For the sensor selection problem, the objective is to design a Kalman gain matrix to minimize the resulting H_2 norm from the noise to the predicted estimation error. Sparsity is achieved by adding a penalty function for non-zero columns of the gain

☆ This material is based upon work supported by a grant from Intel Corporation. The material in this paper was partially presented at the 54th IEEE Conference on Decision and Control, December 15–18, 2015, Osaka, Japan (Zhang et al., 2015). This paper was recommended for publication in revised form by Associate Editor Brett Ninness under the direction of Editor Torsten Söderström.

E-mail addresses: h223zhan@uwaterloo.ca (H. Zhang), raid.ayoub@intel.com (R. Ayoub), sundara2@purdue.edu (S. Sundaram).

http://dx.doi.org/10.1016/j.automatica.2016.12.025

[1] There have also been various recent studies of the dual design-time actuator placement problem (e.g., see Polyak, Khlebnikov, & Shcherbakov, 2013; Tzoumas, Rahimian, Pappas, & Jadbabaie, 2015).

matrix. In contrast to the formulation in Dhingra et al. (2014), in this paper, we directly focus on minimizing functions of the steady state error covariances of discrete-time Kalman filters, and impose a hard constraint on the set of sensors to be chosen.

In Tzoumas, Jadbabaie, and Pappas (2016, in press), the authors studied the design-time sensor selection problem for discrete-time linear time-varying systems over a finite horizon. They assumed that each sensor directly measures one component of the state vector, and the objective is either to minimize the estimation error with a cardinality constraint or to minimize the number of chosen sensors while guaranteeing a certain level of performance. Different from the formulation in Tzoumas et al. (2016, in press), we consider general measurement matrices and focus on minimizing the steady state estimation error of the Kalman filter.

In Yang et al. (2015), the authors considered the same problem as the one we considered here, namely the design-time sensor selection problem for Kalman filtering in discrete-time linear dynamical systems with hard constraints. They showed that the sensor selection problem can be expressed as a semidefinite program (SDP). However, the results in Yang et al. (2015) can only be applied to systems where the sensor noise terms are uncorrelated, and no theoretical guarantees were provided on the performance of the proposed heuristics.

In this paper, we consider the design-time sensor selection problem for optimal filtering of discrete-time linear dynamical systems. Specifically, we study the problem of choosing a set of sensors (under certain constraints) to optimize either the *a priori* or the *a posteriori* error covariance of the corresponding Kalman filter; we will refer to these problems as the priori and posteriori *Kalman filtering sensor selection (KFSS)* problems, respectively. Note that the priori KFSS problem is applicable for settings where a prediction of system states is needed and the posteriori KFSS problem is suitable for applications where the estimation can be conducted after receiving up-to-date measurements (Anderson & Moore, 1979).

Our contributions are threefold. First, we show that it is NP-hard to find the optimal solution of cost-constrained priori and posteriori KFSS problems, even under the assumption that the system is stable. It is often claimed in the literature that sensor selection problems are intractable (Huber, 2012; Joshi & Boyd, 2009); however, except for certain problems with utility or energy based cost functions (e.g., see Bian, Kempe, & Govindan, 2006; Tzoumas et al., 2015), to the best of our knowledge, there is still no explicit characterization of the complexity of the optimal-filtering based sensor selection problems considered in this paper.

Our second contribution is to provide insights into what factors of the system affect the performance of sensor selection algorithms by using the concept of the *sensor information matrix* (Huber, 2012). For the priori KFSS problem, we show that when the system is stable, the worst-case performance can be bounded by a parameter that depends only on the system dynamics matrix, and that the performance of a sensor selection algorithm cannot be arbitrarily bad if the system matrix is well conditioned, even under very large noise. For the posteriori KFSS problem, we show that for a given system, the worst-case performance of any selection of sensors can be upper-bounded in terms of the eigenvalues of the system noise covariance matrix and the corresponding sensor information matrix.

Since it is intractable to find the optimal selection of sensors in general, a reasonable tradeoff is to design appropriate approximation algorithms. In Jawaid and Smith (2015), the authors considered various cost functions for the (run-time) sensor scheduling problem. They showed that one of these considered cost functions is submodular while the others are not; for the submodular cost function, a certain greedy algorithm can be applied to obtain guaranteed performance. Greedy algorithms have also drawn much attention for other forms of sensor selection problems, e.g., see Krause et al. (2008), Shamaiah, Banerjee, and Vikalo (2010), Summers et al. (2016) and Tzoumas et al. (2016). Thus, our third contribution is the study of greedy algorithms for the priori and posteriori KFSS problems. We first show that greedy algorithms are optimal (with respect to the corresponding KFSS problems) for a certain class of systems. However, for general systems, as a negative result, we show that the cost functions of both the priori and posteriori KFSS problems (and the other cost functions studied in Jawaid & Smith, 2015) do not necessarily have certain modularity properties. This precludes the direct application of classical results from the theory of combinatorial optimization and implies that the underlying structures of the KFSS problems are different from the other types of sensor selection problems. Nevertheless, we show via simulations that greedy algorithms perform well in practice. Moreover, compared to the algorithms in Yang et al. (2015), the greedy algorithms provided in this paper can be applied to a more general class of systems (where the sensor noises are correlated), are more efficient and (in simulations) provide comparable performance. Preliminary versions of these results were presented in the conference paper Zhang, Ayoub, and Sundaram (2015).

The rest of the paper is organized as follows. In Section 2, we formulate the (design-time) sensor selection problems. In Section 3, we analyze the complexity of the priori and posteriori KFSS problems. In Section 4, we provide worst-case guarantees on the performance of sensor selection algorithms. In Section 5, we propose and study two greedy algorithms for sensor selection, and illustrate their performance and complexity in Section 6. We conclude in Section 7.

1.1. Notation and terminology

The set of integers, real numbers and complex numbers are denoted as $\mathbb{Z}$, $\mathbb{R}$ and $\mathbb{C}$, respectively. For a square matrix $M \in \mathbb{R}^{n\times n}$, let M^T, $\text{trace}(M)$, $\det(M)$, $\{\lambda_i(M)\}$ and $\{\sigma_i(M)\}$ be its transpose, trace, determinant, set of eigenvalues and set of singular values, respectively. The set of eigenvalues $\{\lambda_i(M)\}$ of M are ordered with nondecreasing magnitude (i.e., $|\lambda_1(M)| \geq \cdots \geq |\lambda_n(M)|$); the same order applies to the set of singular values $\{\sigma_i(M)\}$. A positive semi-definite matrix M is denoted by $M \succeq 0$ and $M \succeq N$ if $M - N \succeq 0$; the set of n by n positive semi-definite (resp. positive definite) matrices is denoted by $\mathbb{S}^n_+$ (resp. $\mathbb{S}^n_{++}$). The identity matrix with dimension n is denoted by $I_{n\times n}$. For a vector v, let $\text{diag}(v)$ be the diagonal matrix with diagonal entries being the elements of v; for a set of matrices $\{M_i\}_{i=1}^q$, let $\text{diag}(M_1, \ldots, M_q)$ be the block diagonal matrix with the ith diagonal block being M_i. For a random variable w, denote $\mathbb{E}[w]$ as its expectation.

2. Problem formulation

Consider the discrete-time linear system

$$x[k+1] = Ax[k] + w[k], \tag{1}$$

where $x[k] \in \mathbb{R}^n$ is the system state, $w[k] \in \mathbb{R}^n$ is a zero-mean white Gaussian noise process with $\mathbb{E}\left[w[k](w[k])^T\right] = W$ for all $k \in \mathbb{N}$, and $A \in \mathbb{R}^{n\times n}$ is the system dynamics matrix. We assume throughout that the pair $(A, W^{\frac{1}{2}})$ is stabilizable.

The set of sensors to be installed on the system must come from a given set $\mathcal{Q}$ consisting of q sensors. Each sensor $i \in \mathcal{Q}$ provides a measurement of the form

$$y_i[k] = C_i x[k] + v_i[k], \tag{2}$$

where $C_i \in \mathbb{R}^{s_i\times n}$ is the state measurement matrix for that sensor, and $v_i[k] \in \mathbb{R}^{s_i}$ is a zero-mean white Gaussian noise process. For convenience, we define $y[k] \triangleq \left[(y_1[k])^T \quad \cdots \quad (y_q[k])^T\right]^T$, $C \triangleq$