Automatica 85 (2017) 100–112

Contents lists available at ScienceDirect

Automatica

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

Optimal capacity allocation for sampled networked systems[☆]

Xudong Chen [a], Mohamed-Ali Belabbas [b], Tamer Başar [b]

[a] The ECEE Department, University of Colorado at Boulder, United States
[b] Coordinated Science Laboratory, University of Illinois at Urbana–Champaign, United States

ARTICLE INFO

Article history:
Received 15 May 2016
Received in revised form 12 April 2017
Accepted 23 June 2017

Keywords:
Sampled networked systems
Mean squared error estimation
Optimal capacity allocation
Algebraic Riccati equations
Capacity filtration

ABSTRACT

We consider the problem of estimating the states of weakly coupled linear systems from sampled measurements. We assume that the total capacity available to the sensors to transmit their samples to a network manager in charge of the estimation is bounded above, and that each sample requires the same amount of communication. Our goal is then to find an optimal allocation of the capacity to the sensors so that the time-averaged estimation error is minimized. We show that when the total available channel capacity is large, this resource allocation problem can be recast as a strictly convex optimization problem, and hence there exists a unique optimal allocation of the capacity. We further investigate how this optimal allocation varies as the available capacity increases. In particular, we show that if the coupling among the subsystems is weak, then the sampling rate allocated to each sensor is nondecreasing in the total sampling rate, and is strictly increasing if and only if the total sampling rate exceeds a certain threshold.

1. Introduction

This paper addresses situations in which a network manager is tasked with estimating the state of an ensemble of weakly interconnected linear systems. For the estimation to be performed, the systems send sampled measurements to the network manager over a shared communication channel. Because this communication channel has a finite capacity, we seek to optimize the allocation of channel capacity to each sensor in order to minimize the *total estimation* error. We work under the assumption that the samples sent by the subsystems all take the *same*, fixed amount of bandwidth.

To proceed, we first describe the model adopted in precise terms. We consider N weakly-coupled stochastic linear systems with sampled outputs

$$S_i := \begin{cases} dx_i = \left(A_i x_i + \epsilon \displaystyle\sum_{j\neq i} A_{ij} x_j\right) dt + G_i dw_i \\ y_i(k\tau_0) = \bar{c}_i^\top x_i(k\tau_0) + v_i(k\tau_0), \end{cases} \tag{1}$$

where $\tau_0 > 0$ is the sampling period ($1/\tau_0 > 0$ is the sampling rate) of the sensors and k is a positive integer. We have that $A_i, A_{ij} \in \mathbb{R}^{n\times n}$, $\bar{c}_i \in \mathbb{R}^{n\times p}$, and $G_i \in \mathbb{R}^{n\times n}$, and that $|\epsilon|$ is small. The assumptions that the subsystems have the same state-dimension n and the outputs y_i have the same dimension p for all i, and the assumption that the coupling parameter ϵ is the same for all pair (i, j), for $i \neq j$, are made to simplify the presentation of the results, but are not necessary for them to hold. The w_i's are pairwise independent standard Wiener processes and the $v_i(k\tau_0)$ are pairwise independent normal random variables. The w_i and v_i are also assumed to be independent. We refer to the system described in (1) as subsystem S_i.

The samples $y_i(k\tau_0)$, $k \in \mathbb{N}$, are sent over a common channel to a network manager whose objective is to estimate the states x_i of the subsystems S_i, for all $i = 1, \ldots, N$, from these samples. The network manager needs to decide the *schedule* with which it receives the samples in order to minimize the estimation error. It is important to note that since the systems are *coupled*, the knowledge of y_i can help with the estimation of x_j, for $i \neq j$.

Problem description. We now describe in detail the scheduling problem that the network manager has to solve. See Fig. 1 for an illustration. The network manager has at his disposal N linear sensors from which he can request samples in order to estimate the states of the subsystems. We only consider *periodic* schedules: we assume that over a fixed scheduling period $\tau > 0$, the network manager can request up to $r_{\text{tot}} = r_1 + r_2 + \cdots + r_N$ samples from the sensors, where r_i is the number of samples from the ith sensor, and over the following scheduling period τ, the same requests are made. We furthermore assume that the r_i's are bounded below by a

☆ Research of T. Başar was supported in part by the U.S. Air Force Office of Scientific Research (AFOSR) MURI grant FA9550-10-1-0573 and in part by National Science Foundation (NSF) grant CCF 11-11342; Research of M.-A. Belabbas was supported in part by National Science Foundation (NSF) ECCS 13-07791 and in part by National Science Foundation (NSF) ECCS CAREER 13-51586. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Tongwen Chen under the direction of Editor Ian R. Petersen.

E-mail addresses: xudong.chen@colorado.edu (X. Chen), belabbas@illinois.edu (M.-A. Belabbas), basar1@illinois.edu (T. Başar).

http://dx.doi.org/10.1016/j.automatica.2017.07.039
0005-1098/

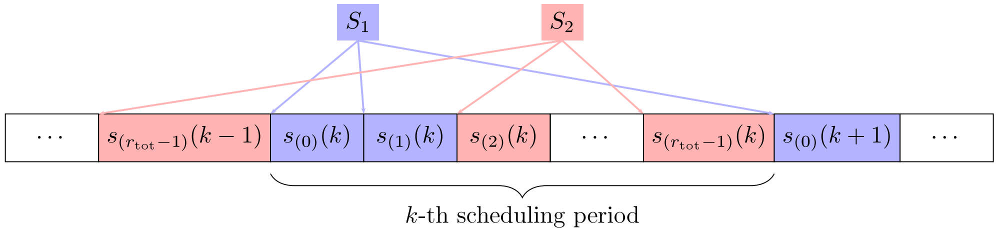

Fig. 1. Over a scheduling period $\tau > 0$, there are r_{tot} slots available to S_1 and S_2 to send their samples to a network manager in charge of the estimation.

positive number $r_{\min}$. We thus have $\tau = r_{\text{tot}}\tau_0$ and we can assume that the time-period τ is then divided evenly into r_{tot} time slots. In each time slot, the network manager can only have one sample sent over the common channel from one of the N sensors. Thus, the problem faced by the network manager is to decide how to assign these r_{tot} slots to the sensors to send their samples to minimize the estimation error.

We note here that the problem has *two* natural scales, τ and τ_0 which are proportional to each other, with ratio r_{tot}. We use the following notation to refer to the time slots. The sub-index l refers to the current position within a time period, and the main index k refers to the current scheduling period. More specifically, for an arbitrary time signal $s(t)$, we let

$$s_{(l)}(k) := s((kr_{\text{tot}} + l)\tau_0) \tag{2}$$

where l is only allowed to take values in the set $\{0, \ldots, r_{\text{tot}} - 1\}$. We illustrate the convention in Fig. 1 for the case $N = 2$. With this convention, we can write the output of the ith sensor as

$$y_{i,(l)}(k) = \bar{c}_i^{\top} x_{(l)}(k) + v_{i,(l)}(k), \tag{3}$$

where the $v_{i,(l)}(k)$'s are pairwise independent normal variables.

We call an **allocation strategy** an assignment of the time slots to the sensors over a scheduling period τ, and denote by $\mathcal{R}$ the set of all possible allocation strategies. We call $\mathcal{R}$ the **strategy set**. Our objective is thus to find an allocation strategy that minimizes the time-averaged (infinite horizon) estimation error. We refer to this problem as the **optimal allocation problem**. A precise formulation of the problem is presented in Section 2. As can be seen from its description, the main difficulty of the problem is a computational one, and we will see below that most of the work deals with finding methods to reduce the complexity.

State of the art. The optimal allocation problem, which is known as the optimal scheduling problem when the dynamics of the state x is in discrete-time, has been a subject of study for the past few decades. Among its numerous applications in networked control and estimation, we mention localization of mobile robot formations (Mourikis & Roumeliotis, 2006), navigation of underwater vehicles using sonar sensors (Meduna, Rock, & McEwen, 2008), target tracking (He & Chong, 2006), and trajectory planning (Singh, Kantas, Vo, Doucet, & Evans, 2007), to name just a few. Because of its widespread relevance, there have been continuing efforts to design efficient algorithms to optimize the allocation of the sensing bandwidth.

We first mention the seminal work (Meier, Peschon, & Dressler, 1967) by Meier, Peschon, and Dressler: the authors there considered a *discrete-time* linear control system with multiple sensors, where only *one* sensor can be used at each time step. Their objective was to determine the schedule of the sensors in order to minimize the total estimation error over a finite horizon. They proposed in that work a method based on dynamic programming to obtain an optimal schedule. However, such a method is often computationally intractable when the number of sensors is large and schedule horizon is long (here, N and r_{tot} are large). Following (Meier et al., 1967), various methods have been established to reduce the computational complexity. Among the deterministic methods, greedy algorithms have been used to find suboptimal solutions (see, for example, Chhetri, Morrell, & Papandreou-Suppappola, 2007; Kagami & Ishikawa, 2006; Oshman, 1994). A different strand of algorithms is based on the observation that the scheduling problem is a tree-search problem. Algorithms in this category rely on pruning of the tree-search (Vitus, Zhang, Abate, Hu, & Tomlin, 2012), which yields trade-offs between the quality of the solution and the complexity of the algorithm through a tuning parameter. We further refer to Alriksson and Rantzer, 2005 for a suboptimal algorithm using relaxed dynamic programming.

In addition to the deterministic algorithms mentioned above, stochastic methods to handle the computational complexity of the optimal scheduling problem have also been developed. For example, the authors in Gupta, Chung, Hassibi, and Murray, 2006 select a sensor randomly at each time step according to a carefully defined probability distribution: an upper-bound for the expected value of the steady-state estimation error is then established, and the probability distribution is chosen so as to minimize the upper-bound. In He and Chong, 2006, the authors proposed a Monte-Carlo method, and in Singh et al., 2007, a more empirical, simulation-based approach. Finally, we point out that the optimal scheduling problem has been investigated for specific nonlinear processes as well. For example, Baras and Bensoussan, 1989 established the existence of an optimal solution for nonlinear diffusion processes.

Closer to this paper is the work on periodic scheduling of actuators and/or sensors (Brockett, 1995; Han, Wu, Zhang, & Shi, 2017; Jiang, Zou, & Zhang, 2008; Shi & Chen, 2013a, 2013b; Zhang & Hristu-Varsakelis, 2005, 2006). We first mention the seminal work (Brockett, 1995), where the author introduced the notion of a (periodic) communication sequence (which is what we called the *allocation strategy* in this paper) and investigated the problem of how to design the sequence so as to stabilize a networked control system. The idea has then been further explored in Jiang et al., 2008; Zhang and Hristu-Varsakelis, 2005, 2006 where issues of controllability, observability and feedback stabilizability have all been addressed to some extent. The most closely related works are Han et al., 2017; Shi and Chen, 2013a, 2013b; the authors in Shi & Chen, 2013a investigated the problem of optimizing over all periodic allocation strategies so as to minimize the time-average infinite-horizon estimation error using a discrete-time linear stochastic system. Unlike what is done here, the authors assumed that each pair $(A, \bar{c}_i)$ is detectable and imposed a minimum dwell-time condition on each sensor, i.e. each sensor is selected for at least a few consecutive samples. In particular, the dwell-time has to be large enough so that the error covariance matrix converges to a positive semi-definite matrix that is ϵ-close to the steady-state of the associated algebraic Riccati equation. The authors then approximated the original optimization problem by replacing every transient error covariance matrix with the corresponding steady-state. In contrast, none of the pairs $(A, \bar{c}_i)$ needs to be detectable in our work and no dwell-time constraint