Automatica 91 (2018) 61–68

Contents lists available at ScienceDirect

Automatica

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

Brief paper

Finite-horizon covariance control for discrete-time stochastic linear systems subject to input constraints ☆

Efstathios Bakolas

Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin, Austin, TX 78712-1221, USA

ARTICLE INFO

Article history:
Received 6 September 2016
Received in revised form 7 September 2017
Accepted 3 January 2018

Keywords:
Covariance control
Stochastic optimal control
Discrete-time linear systems
Convex optimization

ABSTRACT

This work deals with a finite-horizon covariance control problem for discrete-time, stochastic linear systems with complete state information subject to input constraints. First, we present the main steps for the transcription of the covariance control problem, which is originally formulated as a stochastic optimal control problem, into a deterministic nonlinear program (NLP) with a convex performance index and with both convex and non-convex constraints. In particular, the convex constraints in this nonlinear program are induced by the input constraints of the stochastic optimal control problem, whereas the non-convex constraints are induced by the requirement that the terminal state covariance be equal to a prescribed positive definite matrix. Subsequently, we associate this nonlinear program, via a simple convex relaxation technique, with a (convex) semi-definite program, which can be solved numerically by means of modern computational tools of convex optimization. Although, in general, the endpoints of a representative sample of closed-loop trajectories generated by the control policy that corresponds to the solution of the relaxed convex program are not expected to follow exactly the goal terminal Gaussian distribution, they are more likely to be concentrated near the mean of this distribution than if they were drawn from the latter, which is a desirable feature in practice. Numerical simulations that illustrate the key ideas of this work are also presented.

1. Introduction

Given a stochastic discrete-time linear system subject to a white noise process, we seek to find a feedback control policy that will steer the uncertain state of this system from a given Gaussian distribution to another prescribed Gaussian distribution after a fixed (finite) number of stages under the assumption of complete state information. In our problem formulation, we consider explicit constraints on the (weighted) ℓ_2-norm of the (random) input sequence/process. (We will see that the latter constraints will allow us to also enforce, in principle, point-wise in time constraints on the expected value of the norm of the input vector). Without loss of generality (or perhaps, with minimal loss), we will assume that the mean of both the initial and terminal Gaussian distributions are equal to zero, which means that the latter distributions are described completely in terms of their covariance matrices. For this reason, we will broadly refer to the special class of *distribution steering* problems we consider herein as the finite-horizon *covariance control* problem with perfect state information.

Literature review: The covariance control problem was first introduced to the controls community by Hotz and Skelton (Hotz & Skelton, 1985, 1987). This class of problems for both continuous-time and discrete-time stochastic linear systems has been studied extensively in the literature (the reader may refer, for instance, to Grigoriadis & Skelton, 1997; Xu & Skelton, 1992; Yasuda, Skelton, & Grigoriadis, 1993). All these references, however, focus on the infinite-horizon problem in which the objective is to steer the state covariance of a stochastic linear system to a *steady state* covariance matrix, which is a positive definite matrix that satisfies a relevant algebraic Lyapunov matrix equation. The finite-horizon covariance control problem for continuous-time stochastic linear systems has been recently addressed in Chen, Georgiou, and Pavon (2016a, b). It turns out that the continuous-time covariance control problem becomes amenable to analysis and computation, when the input and noise channels of the stochastic linear system are identical (Chen et al., 2016a). On the other hand, the more general case in which the input and the noise channels do not necessarily match turns out to be a much harder problem, whose solvability is in general difficult to be concluded a priori (Chen et al., 2016b). The finite-horizon covariance control problem for continuous-time stochastic linear systems in the presence of "soft" state constraints was addressed in our previous work (Bakolas, 2016b). A finite-horizon covariance control problem in which a soft constraint on the terminal state

☆ The material in this paper was partially presented at the 55th IEEE Conference on Decision and Control, December 12–14, 2016, Las Vegas, NV, USA. This paper was recommended for publication in revised form by Editor Ian R. Petersen.

E-mail address: bakolas@austin.utexas.edu.

https://doi.org/10.1016/j.automatica.2018.01.029

covariance is enforced via an appropriate terminal cost term is addressed in Halder and Wendel (2016).

Problems related to the discrete-time version of the problem considered in Chen et al. (2016a, b) have appeared in Beghi (1996), Levy and Beghi (1997) and Vladimirov and Petersen (2015). In particular, (Beghi, 1996; Levy & Beghi, 1997) deal with the problem of constructing a Markov process with fixed reciprocal dynamics (Jamison, 1970) that connects two prescribed (marginal) probability densities at the endpoints of a given time-interval. Ref. (Vladimirov & Petersen, 2015) deals with the problem of characterizing the noise process that will steer the state of a (control-free) discrete-time stochastic linear system, emanating from a known initial Gaussian distribution to a prescribed terminal Gaussian distribution at a given terminal stage and explores connections between dissipativity theory and robust performance analysis for discrete-time stochastic linear systems. It should be mentioned at this point that despite the fact that Beghi (1996), Levy and Beghi (1997) and Vladimirov and Petersen (2015) present some very important and insightful results, it is not clear how one can directly use these results for the design of feedback control policies that will realize the proposed transitions between the prescribed (marginal) distributions at the endpoints of a given time-interval. The design of such control policies becomes even more challenging when practical input constraints come into play. Problems of control synthesis for discrete-time stochastic linear systems, including stochastic MPC problems (see Kouvaritakis & Cannon, 2015 and references therein), have received a lot of attention in the literature (Agarwal, Cinquemani, Chatterjee, & Lygeros, 2009; Chatterjee, Hokayem, & Lygeros, 2011; Hokayem, Cinquemani, Chatterjee, Ramponi, & Lygeros, 2012; Primbs & Sung, 2009; Skaf & Boyd, 2010). Many of these references rely on convex optimization techniques. It is in a way surprising that, to the best of our knowledge, the idea of applying these powerful techniques to covariance control problems have never been explored in depth before.

Main contribution: This work is purported to fill the gap in the literature regarding the synthesis of feedback control policies for covariance control problems in the presence of input constraints by leveraging some of the powerful techniques of convex optimization (Bertsekas, 2015; Boyd & Vandenberghe, 2004) for control synthesis problems (Agarwal et al., 2009; Chatterjee et al., 2011; Goulart, Kerrigan, & Maciejowski, 2006; Hokayem et al., 2012; Primbs & Sung, 2009; Skaf & Boyd, 2010). Specifically, we present a solution approach to the finite-horizon covariance control problem for discrete-time stochastic linear systems, which is based on the transcription of the stochastic optimal control problem into a deterministic nonlinear program (NLP) with a convex performance index and both convex and non-convex constraints. In particular, the convex constraints of the NLP are induced by the input constraints, whereas the non-convex constraints are induced by the requirement that the terminal state covariance be equal to a prescribed positive definite matrix. We show that the latter matrix equality constraint can be associated with a positive semi-definite (convex) constraint by means of a convex relaxation technique.

It should be mentioned that the endpoints of a representative sample of closed-loop trajectories generated by the control policy induced by the solution to the relaxed convex program are not expected to follow exactly the goal terminal Gaussian distribution. However, they are actually more likely to concentrate near the mean of the goal distribution than if they were drawn from the latter. The previous observation along with the fact that the original covariance control problem can be associated with a convex optimization problem, for the solution of which efficient, scalable and robust algorithms exist (Bertsekas, 2015; Calafiore & El Ghaoui, 2014), outweigh the fact that the latter problem is not equivalent to the original problem in the strict mathematical sense.

Finally, we wish to mention that a preliminary version of this paper has appeared in Bakolas (2016a). The latter reference, however, does not present a complete and detailed description of a systematic approach for the computation of the feedback control policy that solves the covariance control problem subject to input constraints.

Structure of the paper: The rest of the paper is organized as follows. In Section 2, we formulate the covariance control problem as a stochastic optimal control problem, which we transcribe into a finite-dimensional nonlinear program in Section 3. The latter problem is subsequently associated with a convex program, via a convex relaxation technique. Illustrative numerical simulations are presented in Section 4, and finally, Section 5 concludes the paper with a summary of remarks.

2. Problem formulation

2.1. Notation

We denote by $\mathbb{R}^n$ and $\mathbb{R}^{m\times n}$ the set of real n-dimensional (column) vectors and real $m \times n$ matrices, respectively. We write $\mathbb{Z}^+$ and $\mathbb{Z}^{++}$ to denote the set of non-negative integers and strictly positive integers, respectively. Given $z_\alpha, z_\beta \in \mathbb{Z}^+$ with $z_\alpha \le z_\beta$, we denote the *discrete interval* from z_α to z_β as $[z_\alpha, z_\beta]_d$; note that $[z_\alpha, z_\beta]_d = [z_\alpha, z_\beta] \cap \mathbb{Z}^+$. Given a complete probability space $(\Omega, \mathfrak{F}, \mathbb{P})$ and $N \in \mathbb{Z}^{++}$, we denote by $\ell_2^n([0, N]_d; \Omega, \mathfrak{F}, \mathbb{P})$ the Hilbert space of mean square summable and $\mathbb{R}^n$-valued random sequences or processes $X_N := \{x(t) : t \in [0, N]_d\}$ on $(\Omega, \mathfrak{F}, P)$. Given a process X_N in $\ell_2^n([0, N]_d; \Omega, \mathfrak{F}, P)$, we denote its norm by $\|X_N\|_{\ell_2}$, with $\|X_N\|_{\ell_2} := \left(\mathbb{E}\left[\sum_{t=0}^{N} x(t)^{\mathrm{T}} x(t)\right]\right)^{1/2}$, where $\mathbb{E}[\cdot]$ denotes the expectation operator. Given a matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$, we will denote by $\mathrm{vec}(\mathbf{A})$ the mn-dimensional column vector formed by stacking the n columns of $\mathbf{A}$ one below the other. If $\mathbf{A} \in \mathbb{R}^{n\times n}$, then we denote its trace by $\mathrm{trace}(\mathbf{A})$ and by $\mathbf{A}^{-1}$ its inverse (provided that the latter is well defined). We write $\mathbf{0}$ and $\mathbf{I}$ to denote the zero matrix and the identity matrix, respectively.

We will denote by $\mathrm{bdiag}(\mathbf{A}_1, \dots, \mathbf{A}_k)$ the block diagonal matrix formed by matrices $\mathbf{A}_1, \dots, \mathbf{A}_k$ of compatible dimensions. We will denote by $\mathbb{BL}_{P\times Q}(m, n)$ the set of $P \times Q$ block lower triangular matrices whose blocks are $m \times n$ (real) matrices; in the special case when $Q = P$ and $m = n$ we will write $\mathbb{BSL}_P(m)$. Recall that a block matrix $\mathbf{A} = [\mathbf{A}_{ij}]$ is block lower triangular when $\mathbf{A}_{ij} = \mathbf{0}$ for all $j > i$. Note also that $\mathbb{BL}_{P\times Q}(m, n)$ and $\mathbb{BSL}_P(m)$ are convex subsets of $\mathbb{R}^{Pm\times Qn}$ and $\mathbb{R}^{Pm\times Pm}$, respectively. We will write $\mathbf{A} = [\mathbf{A}_{ij}]$, if we want $\mathbf{A}$ to be viewed as an element of $\mathbb{BL}_{P\times Q}(m, n)$, in which case $\mathbf{A}_{ij} \in \mathbb{R}^{m\times n}$, whereas the notation $\mathbf{A} = [\mathbf{A}^{(i,j)}]$ implies that $\mathbf{A}$ should be viewed as an element of $\mathbb{R}^{Pm\times Qn}$, in which case $\mathbf{A}^{(i,j)} \in \mathbb{R}$. The space of real symmetric $n \times n$ matrices will be denoted by $\mathbb{S}_n$. Furthermore, we will denote the convex cone of $n \times n$ (symmetric) positive semi-definite and (symmetric) positive definite matrices by $\mathbb{S}_n^+$ and $\mathbb{S}_n^{++}$, respectively. Given a matrix $\mathbf{A} \in \mathbb{S}_n^{++}$ (resp. $\mathbf{A} \in \mathbb{S}_n^+$), we will also write $\mathbf{A} \succ \mathbf{0}$ (resp., $\mathbf{A} \succeq \mathbf{0}$). In addition, if $\mathbf{A} \succeq \mathbf{0}$, we will denote by $\mathbf{A}^{1/2}$ its (unique) square root in $\mathbb{S}_n^+$. Finally, given two functions $f : \mathcal{Y} \to \mathcal{Z}$ and $g : \mathcal{X} \to \mathcal{Y}$, we denote by $f \circ g : \mathcal{X} \to \mathcal{Z}$, where $(f \circ g)(x) = f(g(x))$, the composition of f with g.

2.2. Formulation of the optimal covariance control problem

For a given $N \in \mathbb{Z}^{++}$, let $\{\mathbf{A}(t) \in \mathbb{R}^{n\times n} : t \in [0, N-1]_d\}$, $\{\mathbf{B}(t) \in \mathbb{R}^{n\times m} : t \in [0, N-1]_d\}$, and $\{\mathbf{C}(t) \in \mathbb{R}^{n\times p} : t \in [0, N-1]_d\}$ be known sequences of real matrices. We consider a stochastic discrete-time linear system that is described by the following stochastic difference equation:

$$x(t+1) = \mathbf{A}(t)x(t) + \mathbf{B}(t)u(t) + \mathbf{C}(t)w(t), \tag{1}$$