Systems & Control Letters 113 (2018) 109–116

Contents lists available at ScienceDirect

Systems & Control Letters

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

Constrained minimum variance control for discrete-time stochastic linear systems

E. Bakolas

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

ARTICLE INFO

Article history:
Received 19 February 2017
Received in revised form 5 January 2018
Accepted 1 February 2018
Available online 17 February 2018

Keywords:
Minimum-variance control
Stochastic optimal control
Discrete-time stochastic systems
Convex optimization

ABSTRACT

We propose a computational scheme for the solution of the so-called minimum variance control problem for discrete-time stochastic linear systems subject to an explicit constraint on the 2-norm of the input (random) sequence. In our approach, we utilize a state space framework in which the minimum variance control problem is interpreted as a finite-horizon stochastic optimal control problem with incomplete state information. We show that if the set of admissible control policies for the stochastic optimal control problem consists exclusively of sequences of causal (non-anticipative) control laws that can be expressed as linear combinations of all the past and present outputs of the system together with its past inputs, then the stochastic optimal control problem can be reduced to a deterministic, finite-dimensional optimization problem. Subsequently, we show that the latter optimization problem can be associated with an equivalent convex program and in particular, a quadratically constrained quadratic program (QCQP), by means of a bilinear transformation. Finally, we present numerical simulations that illustrate the key ideas of this work.

1. Introduction

We propose a computational framework for the characterization of control policies for a special class of stochastic optimal control problems for discrete-time stochastic linear systems with incomplete state information. Specifically, our objective is to compute a control policy that will minimize the expected value of a finite sum of cost-per-stage functions, which are (convex) quadratic functions of the system's output, subject to an explicit constraint on (the expected value) of the ℓ_2-norm of the input (random) sequence. The CMVC problem can find many real world applications in, for instance, the so-called web-forming processes including thickness control of paper sheets, cold or hot rolled sheets and coils, and plastic film extrusion by means of compressive forces [1–3]. Another example is trajectory optimization problems for uncertain dynamical systems in which the objective is to minimize the dispersion of the endpoints of a representative sample of their state trajectories around the terminal goal (mean) state. The latter problem is also related to the problem of steering the distribution of the uncertain state of a stochastic dynamical system to a goal state distribution, which has recently received some notable attention [4–6].

Literature review: The CMVC problem in the absence of constraints reduces to the standard *Minimum Variance Control* (MVC) problem, which is a well studied problem in the literature [7–9]. Typically, the scope of the MVC problem is limited to SISO systems and its solution is based on transfer function design techniques given that in its state-space formulation, the MVC problem corresponds to a *singular* linear quadratic stochastic optimal control problem whose performance index does not reflect any penalty on the control effort. For this reason, one cannot use the standard Riccati-based techniques used for similar, but non-singular, problems and may have to resort instead to more sophisticated geometric techniques [10,11]. It is well-known that the optimal control policy that solves the MVC problem can be characterized by passing the system's output through a certain stable linear filter [12]. The previous interpretation of the solution to the MVC problem implies that the control input that should be applied to the system at each stage can be expressed as a linear combination of the past and present output measurements of the system together with its past inputs. This observation will play an instrumental role in the proposed solution approach for the CMVC problem.

One of the main limitations of the most popular transfer function design techniques for the MVC problem is that their applicability requires the solution of the so-called Diophantine (polynomial) equation, which can be a complex task, especially for high-dimensional and/or time-varying systems [13]. Solution techniques for the MVC problem based on state-space methods have also appeared in the literature [13–15]. A comprehensive presentation and analysis of several formulations of the MVC problem for stochastic linear systems with an emphasis placed on the so-called ARMAX (Auto-Regressive, Moving Average, with eXogenous input) model can be found in [12, pp. 236–251].

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

https://doi.org/10.1016/j.sysconle.2018.02.001
0167-6911/

Main contribution: This work proposes a computational solution approach for the CMVC problem, which is based on convex optimization techniques [16]. The main idea of the proposed solution approach is centered around the interpretation of the CMVC problem as a stochastic optimal control problem with incomplete state information. This particular formulation of the CMVC problem will allow us to leverage certain convex optimization tools and techniques, which are used in control design problems for discrete-time stochastic linear systems (see, for instance, [17–22]), for the development of an algorithmic procedure for the efficient computation of its solution. Motivated by the structure of the optimal policy of the standard MVC problem, we will restrict our search for the optimal control policy of the CMVC problem to the set of sequences of causal (non-anticipative) control laws that can be expressed as linear combinations of the past and present output measurements of the system together with its past inputs. Under this assumption, it turns out that the CMVC problem can be reduced to a tractable deterministic convex program, which can be addressed by means of efficient and robust computational tools. It should be highlighted that this particular parametrization of the admissible control policies has its roots in the so-called Youla–Kucera parametrization of all stabilizing controllers for a given discrete-time linear system as well as the affine/linear disturbance feedback parametrization for discrete-time stochastic linear systems, which was proposed in [23]. For the reduction of the stochastic optimal control problem to a convex program, we will make use of some of the key ideas presented in [22]. Finally, we wish to highlight that despite the fact that in the formulation of the CMVC problem we only consider a single input constraint, the proposed approach can be extended in a natural way to the case of multiple similar state/input constraints. One can use the solution to the problem with such constraints as a high-level roadmap to the control design problem and subsequently employ more specialized techniques from, for instance, the literature of stochastic MPC problems [17,20,24,25], to enforce either hard constraints or tight chance constraints on the applied control inputs point-wisely in time.

Structure of the paper: The remainder of the paper is organized as follows. In Section 2, we formulate the CMVC problem, which we subsequently reduce to a deterministic, finite-dimensional optimization problem, which may not be convex in general, in Section 3. In Section 4, we show that by employing a certain bilinear transformation, the previous optimization problem reduces to a tractable convex program. Numerical simulations that illustrate some of the key ideas of the proposed solution techniques are presented in Section 5. Finally, Section 6 concludes the paper with a summary of remarks.

2. Problem formulation

2.1. Notation

We denote by $\mathbb{R}$ and $\mathbb{R}_{\geq 0}$ the set of real numbers and the set of non-negative real numbers, respectively, and by $\mathbb{R}^n$ and $\mathbb{R}^{m\times n}$ the set of n-dimensional real vectors and $m \times n$ real matrices, respectively. We write $|\boldsymbol{\alpha}|$ to denote the 2-norm of a vector $\boldsymbol{\alpha} \in \mathbb{R}^n$. We write $\mathbb{Z}_{\geq 0}$ and $\mathbb{Z}_{>0}$ to denote the set of non-negative integers and strictly positive integers, respectively. For a given $N \in \mathbb{Z}_{\geq 0}$, we denote by $\mathbb{T}_N$ the discrete set $\{0, \dots, N\} \subset \mathbb{Z}_{\geq 0}$. Given a probability space $(\Omega, \mathfrak{F}, P)$ and $N \in \mathbb{Z}_{>0}$, we denote by $\ell_2^n(\mathbb{T}_N; \Omega, \mathfrak{F}, P)$ the Hilbert space of mean square summable random sequences $\{x(t) : t \in \mathbb{T}_N\}$ on $(\Omega, \mathfrak{F}, P)$, where $x(t)$ is an n-dimensional (random) vector for each $t \in \mathbb{T}_N$. Given $\{x(t) : t \in \mathbb{T}_N\} \in \ell_2^n(\mathbb{T}_N; \Omega, \mathfrak{F}, P)$, we write $\|x(\cdot)\|_{\ell_2}$ to denote its norm in $\ell_2^n(\mathbb{T}_N; \Omega, \mathfrak{F}, P)$, that is, $\|x(\cdot)\|_{\ell_2} := (\mathbb{E}\big[\sum_{t=0}^N |x(t)|^2\big])^{1/2} = (\sum_{t=0}^N \mathbb{E}\big[|x(t)|^2\big])^{1/2}$, where $\mathbb{E}[\cdot]$ denotes the expectation operator. Given a square matrix $\mathbf{A}$, we denote its trace by $\mathrm{trace}(\mathbf{A})$. The induced matrix 2-norm of $\mathbf{A}$ is denoted by $\|\mathbf{A}\|_2$, where $\|\mathbf{A}\|_2 = (\lambda_{\max}(\mathbf{A}^{\mathrm{T}}\mathbf{A}))^{1/2}$ and $\lambda_{\max}(\mathbf{M})$ denotes the maximum eigenvalue of a real symmetric matrix $\mathbf{M}$. We write $\mathbf{0}_{m\times p}$ (or simply, $\mathbf{0}$) and $\mathbf{I}_m$ (or simply, $\mathbf{I}$) to denote the $m \times p$ zero matrix and the $m \times m$ identity matrix, respectively. Furthermore, we denote by $\mathrm{bdiag}(\mathbf{A}_1, \dots, \mathbf{A}_\ell)$ the block diagonal matrix whose diagonal blocks are matrices $\mathbf{A}_i$, $i \in \{1, \dots, \ell\}$, of compatible dimensions. The set of $N \times N$ block square and lower triangular (real) matrices whose blocks have dimension $m \times n$ will be denoted by $\mathfrak{BL}_N(m, n)$; note that $\mathfrak{BL}_N(m, n) \subset \mathbb{R}^{Nm\times Nn}$. We will denote the convex cone of $n \times n$ symmetric positive definite and positive semi-definite matrices by $\mathbb{P}_n$ and $\overline{\mathbb{P}}_n$, respectively. Finally, for a given a matrix $\mathbf{A} \in \overline{\mathbb{P}}_n$, we will denote by $\mathbf{A}^{1/2}$ its (unique) square root in $\overline{\mathbb{P}}_n$.

2.2. Formulation of the constrained minimum variance control problem

For a given $N \in \mathbb{Z}_{>0}$, let $\{\mathbf{A}(t) \in \mathbb{R}^{n\times n} : t \in \mathbb{T}_{N-1}\}$, $\{\mathbf{B}(t) \in \mathbb{R}^{n\times m} : t \in \mathbb{T}_{N-1}\}$, $\{\mathbf{C}(t) \in \mathbb{R}^{n\times p} : t \in \mathbb{T}_{N-1}\}$, $\{\mathbf{G}(t) \in \mathbb{R}^{n\times q} : t \in \mathbb{T}_{N-1}\}$, and $\{\mathbf{N}(t) \in \mathbb{R}^{n\times r} : t \in \mathbb{T}_{N-1}\}$ denote known sequences of matrices of appropriate dimensions. Let us also consider a discrete-time stochastic linear system that satisfies the following stochastic difference equation and output equation, respectively:

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

$$y(t) = \mathbf{C}(t)x(t) + \mathbf{N}(t)v(t), \tag{1b}$$

for $t \in \mathbb{T}_{N-1}$, where $x(0) = x_0$ is a random vector drawn from the Gaussian distribution $\mathcal{N}(\mu_0, \boldsymbol{\Sigma}_0)$ with μ_0 and $\boldsymbol{\Sigma}_0$ be, respectively, a given vector in $\mathbb{R}^n$ and a given matrix in $\mathbb{P}_n$. In addition, $\{x(t) : t \in \mathbb{T}_N\}$, $\{u(t) : t \in \mathbb{T}_{N-1}\}$, and $\{y(t) : t \in \mathbb{T}_{N-1}\}$ denote, respectively, the state, the control input, and the output (random) sequences on a complete probability space $(\Omega, \mathfrak{F}, P)$. In addition, the control input sequence $\{u(t) : t \in \mathbb{T}_{N-1}\}$ is assumed to belong to $\ell_2^m(\mathbb{T}_{N-1}; \Omega, \mathfrak{F}, P)$ and to have finite k-moments for all $k > 0$. We will henceforth refer to a control input sequence that satisfies these properties as *admissible*. In addition, $\{w(t) : t \in \mathbb{T}_{N-1}\}$ and $\{v(t) : t \in \mathbb{T}_{N-1}\}$ are sequences of independent normal random variables with zero mean and unit covariance, that is,

$$\mathbb{E}[w(t)] = \mathbf{0}, \qquad \mathbb{E}\left[w(t)w(\tau)^{\mathrm{T}}\right] = \delta(t, \tau)\mathbf{I}, \tag{2a}$$

$$\mathbb{E}[v(t)] = \mathbf{0}, \qquad \mathbb{E}\left[v(t)v(\tau)^{\mathrm{T}}\right] = \delta(t, \tau)\mathbf{I}, \tag{2b}$$

for all $t, \tau \in \mathbb{T}_{N-1}$, with $\delta(t, \tau) := 1$, if $t = \tau$, and $\delta(t, \tau) := 0$, otherwise. It is assumed that x_0 and $\{w(t) : t \in \mathbb{T}_{N-1}\}$ as well as $\{w(t) : t \in \mathbb{T}_{N-1}\}$ and $\{v(t) : t \in \mathbb{T}_{N-1}\}$ are mutually independent, which implies that

$$\mathbb{E}\left[w(t)v(\tau)^{\mathrm{T}}\right] = \mathbf{0}, \tag{3a}$$

$$\mathbb{E}\left[v(t)x_0^{\mathrm{T}}\right] = \mathbf{0}, \quad \mathbb{E}\left[w(t)x_0^{\mathrm{T}}\right] = \mathbf{0}, \tag{3b}$$

for all $t, \tau \in \mathbb{T}_{N-1}$.

Our objective is to find a control policy that minimizes the expected value of a finite sum of cost-per-stage functions, which are convex quadratic functions of the output measurement $y(t)$ of the stochastic linear system (1a)–(1b) as t runs through $\mathbb{T}_{N-1}$, subject to an explicit inequality constraint on the ℓ_2-norm of the input sequence (realization of the control policy). We will assume that the set of admissible control policies, which is denoted by Π, consists of all control policies π which are sequences of control laws $\kappa(\cdot; t)$ that are causal (non-anticipative), measurable functions of the elements of the so-called *information set* up to time t. For a given $t \in \mathbb{T}_{N-1}$, the information set, which is a random discrete set, is denoted as $\mathcal{I}_t$ and is defined as follows: $\mathcal{I}_t := \mathcal{I}_t^y \times \mathcal{I}_{t-1}^u$, where $\mathcal{I}_t^y := \{y(\tau) \in \mathbb{R}^p : \tau \in \mathbb{T}_t\}$ and $\mathcal{I}_{t-1}^u := \{u(\sigma) \in \mathbb{R}^m : \sigma \in \mathbb{T}_{t-1}\}$.