Covariance Steering of Discrete-Time Stochastic Linear Systems Based on Wasserstein Distance Terminal Cost

Isin M. Balci and Efstathios Bakolas

***Abstract*— We consider a class of stochastic optimal control problems for discrete-time linear systems whose objective is the characterization of control policies that will steer the probability distribution of the terminal state of the system close to a desired Gaussian distribution. In our problem formulation, the closeness between the terminal state distribution and the desired (goal) distribution is measured in terms of the squared Wasserstein distance which is associated with a corresponding terminal cost term. We recast the stochastic optimal control problem as a finite-dimensional nonlinear program whose performance index can be expressed as the difference of two convex functions. This representation of the performance index allows us to find local minimizers of the original nonlinear program via the so-called convex-concave procedure [1]. Finally, we present non-trivial numerical simulations to demonstrate the efficacy of the proposed technique by comparing it with sequential quadratic programming methods in terms of computation time.**

***Index Terms*— Stochastic Optimal Control, Optimization, Uncertain Systems**

I. Introduction

We consider covariance steering problems for discrete-time stochastic linear systems in which, however, the constraints on the terminal state covariance are enforced indirectly by means of appropriate terminal costs. Specifically we consider the problem of steering the state of a stochastic system, which is originally drawn from a given Gaussian distribution, to a terminal state whose distribution is "close" to a desired (prescribed) Gaussian distribution, where the closeness between the two distributions is measured in terms of the squared Wasserstein distance. We show that the resulting problem can be reduced to a tractable optimization problem which can be solved efficiently if one exploits its structure.

Literature Review: The main focus of the first attempts to study covariance steering problems [2], [3], [4] was on finding stabilizing controllers that drive the state covariance to a desired positive definite matrix asymptotically (infinite-horizon case). The covariance steering problem turns out to be closely related to the so-called Schrödinger's bridge problem, which plays an important role in optimal mass transport and statistical mechanics [5]. Finite-horizon covariance control problems for continuous-time linear systems were recently studied in [6], [7], [8], [9]. Covariance steering problems for discrete-time systems are also receiving significant attention at present. In [10], the constrained covariance steering problem is recast as a finite dimensional convex optimization problem based on a semidefinite relaxation of the constraint on the terminal state covariance. Covariance steering problems with convex chance constraints are studied in [11].

In the previously discussed references, the specifications on the terminal state covariance correspond to hard constraints which often lead to difficult problems (for instance, the analytic solution to the covariance steering problem presented in [6] is only valid for the special case in which the input and noise channels coincide). An alternative problem formulation, which has inspired this paper, is presented in [12] in which a terminal cost is used as a "soft" constraint on the terminal state covariance. The latter cost corresponds to the squared Wasserstein distance between a desired state distribution and the "actual" terminal state distribution. The latter formulation leads to a standard two-point boundary value problem which can be solved by means of indirect shooting methods. It is well known that the success of such methods relies on knowledge of good initial guesses and thus, in general, a systematic process for the computation of the solution to the class of covariance steering problems proposed in [12] with soft terminal constraints is still missing.

Main Contribution: We first formulate the covariance steering problem as a stochastic optimal control problem in which the requirement on the terminal state covariance is encoded in a terminal cost term ("soft constraint"). Similarly with [12], we consider the case in which the terminal cost corresponds to the squared Wasserstein distance between the actual terminal state distribution and the desired Gaussian distribution but in contrast with the latter reference, we consider the discrete-time case. First, we recast this stochastic optimal control problem as a (deterministic) nonlinear program by utilizing an affine state feedback control policy parametrization (the control input at each stage is an affine function of the history of visited states). Then, we show that the performance index of the nonlinear program can be expressed as the difference of two convex functions by using a suitable bilinear transformation of the decision variables. To the best of our knowledge, this is the first paper that shows that covariance steering problems can be formulated as difference of convex functions programs

I. M. Balci (graduate student) and E. Bakolas (Associate Professor) are with the Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin, Austin, Texas 78712-1221, USA, Emails: isinmertbalci@utexas.edu; bakolas@austin.utexas.edu

This research has been supported in part by NSF awards ECCS-1924790 and CMMI-1937957.

(DCPs). By leveraging this fact, one can find local minimizers of the nonlinear program via efficient techniques such as the so-called convex-concave procedure (CCP) [1], [13]. The CCP is an iterative procedure which can compute local minimizers of non-convex optimization problems which correspond to DCP based on successive convexifications. Exploiting this extra structure of the problem reduces its complexity and allows us to use convex optimization solvers which in turn leads to improved scalability and numerical efficiency. Finally, we show the efficacy of our approach in numerical simulations in terms of computation time compared to general purpose NLP solvers.

Outline: The rest of the paper is organized as follows. Section II presents the problem formulation. In Section III, we show that the covariance steering problem with Wasserstein distance terminal cost can be associated with a difference of convex functions program. In Section IV, we present numerical simulations. Finally, Section V concludes the paper with a summary of remarks and directions for future research.

II. Problem Formulation

A. Notation

We denote by $\mathbb{R}^n$ the set of n-dimensional real vectors and by $\mathbb{R}$ and $\mathbb{R}^+$ (resp., $\mathbb{R}^{++}$) the set of real numbers and non-negative (resp., strictly positive) real numbers, respectively. The sets of non-negative and strictly positive integers are denoted by $\mathbb{Z}^+$ and $\mathbb{Z}^{++}$, respectively. We denote by $\mathbb{E}[\cdot]$ the expectation functional. Given a random vector x, we denote its mean vector and covariance matrix by $\mathbb{E}[x]$ and $\mathrm{Cov}[x]$, respectively. The space of $n\times n$ symmetric matrices is denoted by $\mathbb{S}_n$ and the cone of positive semi-definite (positive definite) symmetric matrices by $\mathbb{S}_n^+$ ($\mathbb{S}_n^{++}$). The trace of a square matrix is denoted as $\mathrm{tr}(\cdot)$. The transpose of a matrix $A \in \mathbb{R}^{n\times m}$ is denoted by A^{T}, its nuclear norm by $\|A\|_*$ where $\|A\|_* := \mathrm{tr}((A^{\mathrm{T}}A)^{1/2})$ and its Frobenius norm by $\|A\|_F$ where $\|A\|_F := \sqrt{\mathrm{tr}(A^{\mathrm{T}}A)}$. The block diagonal matrix formed by n matrices $A_1,\dots,A_n$ is denoted by $\mathrm{blkdiag}(A_1,\dots,A_n)$. The zero matrix is denoted as $\mathbf{0}$ whereas the identity matrix as $\boldsymbol{I}$. We write $x\sim\mathcal{N}(\mu,S)$ to denote that x is a Gaussian random vector with mean $\mu\in\mathbb{R}^n$ and covariance $S\in\mathbb{S}_n^{++}$.

B. Wasserstein Distance

In this paper, we shall formulate a stochastic optimal control problem with terminal cost that measures the closeness between the final state probability distribution and a desired distribution in terms of the squared Wasserstein distance. Next, we briefly review the concept of Wasserstein distance between two distributions and we provide its analytic expression when both of them are Gaussian.

Given two random vectors x_1, x_2 over $\mathbb{R}^n$ with probability density functions ρ_1,ρ_2, their squared Wasserstein distance is defined as follows:

$$W^2(\rho_1,\rho_2) := \inf_{\rho\in\mathcal{P}(\rho_1,\rho_2)} \mathbb{E}_y[\|x_1-x_2\|_2^2], \tag{1}$$

where $\rho\in\mathcal{P}(\rho_1,\rho_2)$ is the joint probability density function (pdf) of the random vector $y := [x_1,\ x_2]^{\mathrm{T}}$, and $\mathcal{P}(\rho_1,\rho_2)$ denotes the set of all probability density functions over $\mathbb{R}^{2n}$ with finite second moments and marginals ρ_1 and ρ_2 on x_1 and x_2, respectively. It is worth mentioning that the Wasserstein distance between two probability measures is a valid distance metric (in the strict mathematical sense) because it satisfies all of the properties of a metric (non-negativity, symmetry, triangle inequality, and identity).

If $x_i\sim\mathcal{N}(\mu_i,S_i)$, where $\mu_i\in\mathbb{R}^n$ and $S_i\in\mathbb{S}_n^{++}$ for $i\in\{1,2\}$, then the squared Wasserstein distance can be written in closed-form as follows [14]:

$$W^2(\rho_1,\rho_2) = \|\mu_1-\mu_2\|_2^2 + \mathrm{tr}\Big(S_1+S_2-2(S_2^{1/2}S_1S_2^{1/2})^{1/2}\Big). \tag{2}$$

C. Problem Statement

We consider an uncertain system whose dynamics is described by the following discrete-time stochastic linear state space model:

$$x_{k+1} = A_k x_k + B_k u_k + G_k w_k, \qquad \forall k\in\mathbb{Z}^+, \tag{3}$$

where $\{x_k\}_{k\in\mathbb{Z}^+}$ is the state (random) process over $\mathbb{R}^{n_x}$, $\{u_k\}_{k\in\mathbb{Z}^+}$ is the input process over $\mathbb{R}^{n_u}$ and $\{w_k\}_{k\in\mathbb{Z}^+}$ is the noise (random) process over $\mathbb{R}^{n_w}$. In particular, $\{w_k\}_{k\in\mathbb{Z}^+}$ corresponds to a white Gaussian noise process with $\mathbb{E}[w_k]=0$ and $\mathbb{E}[w_k w_m^{\top}]=\delta(k,m)S_w$, where $S_w\in\mathbb{S}_{n_w}^{++}$ and $\delta(k,m)=1$ when $k=m$ and $\delta(k,m)=0$, otherwise. We also assume that the initial state $x_0\sim\mathcal{N}(\mu_0,S_0)$ and that x_0 and $\{w_k\}$ are mutually independent, which implies that $\mathbb{E}[x_0 w_k^{\top}]=\mathbf{0}$ for all $k\in\mathbb{Z}^+$.

Our objective is to drive the uncertain state of the system (3) from its given initial distribution to a terminal distribution which is close to a desired terminal Gaussian probability distribution $\mathcal{N}(\mu_d,S_d)$, where $\mu_d\in\mathbb{R}^n$ and $S_d\in\mathbb{S}_n^{++}$ are given, at a given finite time while minimizing a relevant performance index. Next, we provide the precise formulation of our problem.

Problem 1. *Let $\mu_0,\mu_f\in\mathbb{R}^{n_x}$, $S_0,S_f\in\mathbb{S}_{n_x}^{++}$, $\lambda>0$ and $N\in\mathbb{Z}^{++}$ be given. In addition, let Π denote the set of all admissible control policies $\pi := \{m_0(\cdot),\dots,m_{N-1}(\cdot)\}$ for system* (3)*, with $u_k=m_k(X^k)$ where X^k denotes the (finite) sequence of states visited up to stage $t=k$, that is, $X^k := \{x_0,x_1,\dots x_k\}$, and $m_k(X^k)$ are measurable functions of the elements of X^k, for $k=0,\dots,N-1$. Then, find a control policy $\pi^*\in\Pi$ that solves the following stochastic optimal control problem:*

$$\underset{\pi\in\Pi}{\text{Minimize}} \quad \mathbb{E}\left[\sum_{k=0}^{N-1} u_k^{\mathrm{T}}u_k\right] + \lambda\varphi(\rho_N,\rho_d) \tag{4a}$$

$$\text{subject to} \quad x_{k+1} = A_k x_k + B_k u_k + G_k w_k \tag{4b}$$

$$x_0\sim\mathcal{N}(\mu_0,S_0) \tag{4c}$$

where ρ_d is the pdf corresponding to the Gaussian probability distribution $\mathcal{N}(\mu_d,S_d)$ (goal terminal state probability distribution), ρ_N is the pdf of the terminal state x_N, and $\varphi(\rho_N,\rho_d)$ denotes the squared Wasserstein distance between ρ_N and ρ_d, that is, $\varphi(\rho_N,\rho_d) := W^2(\rho_N,\rho_d)$.

In order to associate Problem 1 with a tractable, finite-dimensional optimization problem, we only consider admissible control policies that correspond to sequences of control laws $m_k(\cdot)$ which are affine functions of the state history:

$$m_k(X^k) = \sum_{i=0}^{k} K(k,i)\big(x_i - \bar{x}_i\big) + u_{\rm ff}(k), \tag{5}$$

where $\bar{x}_i = \mathbb{E}[x_i]$. Next, we show the main steps for recasting Problem 1, whose decision variable corresponds to the control policy π, as an optimization problem whose decision variables are the controller parameters $u_{\rm ff}(k) \in \mathbb{R}^{n_u}$ and $K(k,j) \in \mathbb{R}^{n_u \times n_x}$, for $k, j \in \{0, \ldots, N-1\}$ with $k \geq j$.

III. Difference of Convex Functions Programming Formulation

In this section, we will show that Problem 1 can be associated with a difference of convex functions program (DCP), that is, a nonlinear program whose performance index is equal to the difference of two convex functions. This will allow us to efficiently compute local minimizers of Problem 1 by means of heuristic and easily implementable algorithms, such as the convex-concave procedure [1]. It is worth mentioning that the set of objective functions which can be expressed as the difference of convex functions is dense in the set of continuous functions; moreover, every twice differentiable function can be represented as the difference of convex functions [15]. However, there is no systematic process that is guaranteed to find such a representation for a given function of interest except for a few special classes of functions.

Next, we recast Problem 1 as a finite-dimensional optimization problem. To this aim, we express the state x_k in terms of a finite-dimensional decision variable. In particular, by propagating forward in time the state of the discrete-time stochastic system (3) and using the control policy parametrization given in (5), we can express x_k as a function of x_0, $\{u_i\}_{i=0}^{k-1}$ and $\{w_i\}_{i=0}^{k-1}$ as follows:

$$x_k = \Phi(k,0)x_0 + \sum_{i=0}^{k-1} \Phi(k,i)B_i u_i + \sum_{i=0}^{k-1} \Phi(k,i) G_i w_i, \tag{6}$$

where $\Phi(k,n) := A_{k-1} \ldots A_n$, $\Phi(n,n) = \boldsymbol{I}$ with $k \geq n$ for $k, n \in \mathbb{Z}^+$ (state transition matrix of (3)). Now, let us define the following quantities:

$$\boldsymbol{x} := [x_0^{\rm T},\ x_1^{\rm T},\ \ldots,\ x_N^{\rm T}]^{\rm T} \in \mathbb{R}^{n_x(N+1)}, \tag{7a}$$
$$\boldsymbol{u} := [u_0^{\rm T},\ u_1^{\rm T},\ \ldots,\ u_{N-1}^{\rm T}]^{\rm T} \in \mathbb{R}^{n_u N}, \tag{7b}$$
$$\boldsymbol{w} := [w_0^{\rm T},\ w_1^{\rm T},\ \ldots,\ w_{N-1}^{\rm T}]^{\rm T} \in \mathbb{R}^{n_w N}. \tag{7c}$$

By using equations (6)-(7), it follows that

$$\boldsymbol{x} = \boldsymbol{\Gamma} x_0 + \boldsymbol{H_u u} + \boldsymbol{H_w w}, \tag{8}$$

where

$$\boldsymbol{\Gamma} := [\boldsymbol{I}\ \Phi(1,0)\ \Phi(2,0)\ \ldots\ \Phi(N,0)], \tag{9}$$

$$\boldsymbol{H_u} := \begin{bmatrix} \boldsymbol{0} & \boldsymbol{0} & \ldots & \boldsymbol{0} \\ B_0 & \boldsymbol{0} & \ldots & \boldsymbol{0} \\ \Phi(2,1)B_0 & B_1 & \ldots & \boldsymbol{0} \\ \vdots & \vdots & \vdots & \vdots \\ \Phi(N,1)B_0 & \Phi(N,2)B_1 & \ldots & B_{N-1} \end{bmatrix}, \tag{10}$$

and $\boldsymbol{H_w}$ is defined similarly, after replacing the matrices B_i in (10) with the matrices G_i. The reader is referred to [10] for the details on the derivation of (8)-(10).

Because the performance index of Problem 1 consists of a terminal cost term, we will use the following equation:

$$x_N = \boldsymbol{F x}, \quad \boldsymbol{F} := [\boldsymbol{0} \cdots\ \boldsymbol{0}\ \boldsymbol{I}], \tag{11}$$

to recover x_N from $\boldsymbol{x}$.

Given the particular affine parametrization of the control policy given in (5) and the fact that the initial state is assumed to be a Gaussian (random) vector, it follows that the states of the system in the subsequent stages will also be Gaussian (random) vectors. In addition, we obtain

$$\boldsymbol{u} = \boldsymbol{K}(\boldsymbol{x} - \bar{\boldsymbol{x}}) + \boldsymbol{u}_{\rm ff}, \tag{12}$$

where $\bar{\boldsymbol{x}} := \mathbb{E}[\boldsymbol{x}]$, $\boldsymbol{u}_{\rm ff} := [u_{\rm ff}^{\rm T}(0), \ldots, u_{\rm ff}^{\rm T}(N-1)]^{\rm T}$ and

$$\boldsymbol{K} := \begin{bmatrix} K(0,0) & \boldsymbol{0} & \ldots & \boldsymbol{0} \\ K(1,0) & K(1,1) & \ldots & \boldsymbol{0} \\ K(2,0) & K(2,1) & \ldots & \boldsymbol{0} \\ \vdots & \vdots & \vdots & \vdots \\ K(N-1,0) & K(N-1,1) & \ldots & \boldsymbol{0} \end{bmatrix}. \tag{13}$$

We proceed with the derivation of the expression of the performance index of Problem 1 in terms of the decision variables $\boldsymbol{u}_{\rm ff}$ and $\boldsymbol{K}$. To this aim, we write $\sum_{k=0}^{N-1} u_k^{\rm T} u_k = \boldsymbol{u}^{\rm T}\boldsymbol{u}$, which in view of basic properties of the trace operator and (12) gives

$$\begin{aligned} \mathbb{E}[\boldsymbol{u}^{\rm T}\boldsymbol{u}] &= \mathbb{E}[\mathrm{tr}(\boldsymbol{u}\boldsymbol{u}^{\rm T})] \\ &= \mathbb{E}[\mathrm{tr}((\boldsymbol{K}(\boldsymbol{x} - \bar{\boldsymbol{x}}) + \boldsymbol{u}_{\rm ff})(\boldsymbol{K}(\boldsymbol{x} - \bar{\boldsymbol{x}}) + \boldsymbol{u}_{\rm ff})^{\rm T})] \\ &= \mathrm{tr}(\boldsymbol{K}\mathbb{E}[\tilde{\boldsymbol{x}}\tilde{\boldsymbol{x}}^{\rm T}]\boldsymbol{K}^{\rm T}) + \|\boldsymbol{u}_{\rm ff}\|_2^2, \end{aligned} \tag{14}$$

where $\tilde{\boldsymbol{x}} := \boldsymbol{x} - \bar{\boldsymbol{x}}$ and in the derivation of the last equality, we have used the fact that $\boldsymbol{u}_{\rm ff}$ is a deterministic quantity.

For the computation of $\mathrm{Cov}[\boldsymbol{x}] = \mathbb{E}[\tilde{\boldsymbol{x}}\tilde{\boldsymbol{x}}^{\rm T}]$, we first have to compute $\bar{\boldsymbol{x}} = \mathbb{E}[\boldsymbol{x}]$. By taking expectation on both sides of (8), we obtain:

$$\begin{aligned} \mathbb{E}[\boldsymbol{x}] &= \mathbb{E}[\boldsymbol{\Gamma} x_0 + \boldsymbol{H_u}(\boldsymbol{K}(\boldsymbol{x} - \bar{\boldsymbol{x}}) + \boldsymbol{u}_{\rm ff}) + \boldsymbol{H_w w}] \\ &= \boldsymbol{\Gamma}\mu_0 + \boldsymbol{H_u}\boldsymbol{u}_{\rm ff}. \end{aligned} \tag{15}$$

After some simple algebraic manipulations, we get:

$$\tilde{\boldsymbol{x}} = (\boldsymbol{I} - \boldsymbol{H_u K})^{-1}(\boldsymbol{\Gamma}(x_0 - \mu_0) + \boldsymbol{H_w w}), \tag{16}$$

from which it follows:

$$\mathbb{E}[\tilde{\boldsymbol{x}}\tilde{\boldsymbol{x}}^{\rm T}] = \bar{\boldsymbol{K}}(\Gamma S_0 \Gamma^{\rm T} + \boldsymbol{H_w S_w H_w}^{\rm T})\bar{\boldsymbol{K}}^{\rm T}, \tag{17}$$

where $\bar{\boldsymbol{K}} := (\boldsymbol{I} - \boldsymbol{H_u K})^{-1}$ and $\boldsymbol{S_w} := \mathbb{E}[\boldsymbol{w}\boldsymbol{w}^{\rm T}] = \mathrm{blkdiag}(S_w, \ldots, S_w)$. From (15) and (17), we can obtain the following expressions for $\mu_N := \mathbb{E}[x_N]$ and $S_N := \mathrm{Cov}[x_N]$:

$$\mu_N = \boldsymbol{F}(\boldsymbol{\Gamma}\mu_0 + \boldsymbol{H_u}\boldsymbol{u}_{\rm ff}), \tag{18a}$$
$$S_N = \boldsymbol{F}(\boldsymbol{I} - \boldsymbol{H_u K})^{-1}\tilde{S}(\boldsymbol{I} - \boldsymbol{H_u K})^{-\rm T}\boldsymbol{F}^{\rm T}, \tag{18b}$$

where $\tilde{S} = (\Gamma S_0 \Gamma^{\rm T} + \boldsymbol{H_w S_w H_w}^{\rm T})$. By plugging (17) into (14), we have:

$$\begin{aligned} \mathbb{E}[\boldsymbol{u}^{\rm T}\boldsymbol{u}] = &\ \mathrm{tr}(\boldsymbol{K}(\boldsymbol{I} - \boldsymbol{H_u K})^{-1}\tilde{S}(\boldsymbol{I} - \boldsymbol{H_u K})^{-\rm T}\boldsymbol{K}^{\rm T}) \\ &+ \|\boldsymbol{u}_{\rm ff}\|^2. \end{aligned} \tag{19}$$

After plugging the expressions of μ_N and S_N in (18a) and (18b) into the expression of $W^2(\rho_N, \rho_d)$ for Gaussian distributions, which is given in (2), we get:

$$\begin{aligned}W^2(\rho_N,\rho_d) =\ & \|\boldsymbol{F}(\boldsymbol{\Gamma}\mu_0 + \boldsymbol{H_u}\boldsymbol{u}_{\mathrm{ff}}) - \mu_d\|_2^2 \\ &+ \operatorname{tr}(\boldsymbol{F}(\boldsymbol{I} - \boldsymbol{H_u}\boldsymbol{K})^{-1}\tilde{S}\boldsymbol{F}(\boldsymbol{I} - \boldsymbol{H_u}\boldsymbol{K})^{-\mathrm{T}}\boldsymbol{F}^{\mathrm{T}} + S_d) \\ &- 2\operatorname{tr}((\sqrt{S_d} \times \\ &\quad (\boldsymbol{F}(\boldsymbol{I} - \boldsymbol{H_u}\boldsymbol{K})^{-1}\tilde{S}\boldsymbol{F}(\boldsymbol{I} - \boldsymbol{H_u}\boldsymbol{K})^{-\mathrm{T}}\boldsymbol{F}^{\mathrm{T}}) \\ &\quad \times \sqrt{S_d})^{1/2}). \end{aligned} \tag{20}$$

At this point, we propose to apply a variable transformation, which was first proposed in [16] and later used for covariance steering problems in [10], to convexify the optimization problem. In particular, we introduce the new variable, $\boldsymbol{\Theta}$, which is defined as follows:

$$\boldsymbol{\Theta} := \boldsymbol{K}(\boldsymbol{I} - \boldsymbol{H_u}\boldsymbol{K})^{-1} =: \varphi(\boldsymbol{K}) \tag{21a}$$

$$\boldsymbol{K} := (\boldsymbol{I} + \boldsymbol{H_u}\boldsymbol{\Theta})^{-1}\boldsymbol{\Theta} =: \phi(\boldsymbol{\Theta}). \tag{21b}$$

Furthermore, by using the identity $(\boldsymbol{I}+\boldsymbol{P})^{-1} = \boldsymbol{I} - \boldsymbol{P}(\boldsymbol{I} + \boldsymbol{P})^{-1}$, we obtain:

$$\begin{aligned}(\boldsymbol{I} - \boldsymbol{H_u}\boldsymbol{K})^{-1} &= \boldsymbol{I} + \boldsymbol{H_u}\boldsymbol{K}(\boldsymbol{I} - \boldsymbol{H_u}\boldsymbol{K})^{-1} \\ &= (\boldsymbol{I} + \boldsymbol{H_u}\boldsymbol{\Theta}).\end{aligned} \tag{21c}$$

As is shown in [16], the functions $\phi(\cdot)$ and $\varphi(\cdot)$ determine a bijective transformation, that is, $\phi(\cdot) = \varphi^{-1}(\cdot)$ and vice versa. Therefore, the right hand sides of equations (19) and (20) can be expressed equivalently in terms of $\boldsymbol{u}_{\mathrm{ff}}$ and the new decision variable $\boldsymbol{\Theta}$, which is defined in (21), as follows:

$$\mathbb{E}[\boldsymbol{u}^{\mathrm{T}}\boldsymbol{u}] = \operatorname{tr}(\boldsymbol{\Theta}\tilde{S}\boldsymbol{\Theta}^{\mathrm{T}}) + \boldsymbol{u}_{\mathrm{ff}}^{\mathrm{T}}\boldsymbol{u}_{\mathrm{ff}} \tag{22}$$

$$\begin{aligned}W^2 &= \|\boldsymbol{F}(\boldsymbol{\Gamma}\mu_0 + \boldsymbol{H_u}\boldsymbol{u}_{\mathrm{ff}}) - \mu_d\|_2^2 \\ &+ \operatorname{tr}(\boldsymbol{F}(\boldsymbol{I} + \boldsymbol{H_u}\boldsymbol{\Theta})\tilde{S}(\boldsymbol{I} + \boldsymbol{H_u}\boldsymbol{\Theta})^{\mathrm{T}}\boldsymbol{F}^{\mathrm{T}}) \\ &- 2\operatorname{tr}((\sqrt{S_d}\boldsymbol{F}(\boldsymbol{I} + \boldsymbol{H_u}\boldsymbol{\Theta})\tilde{S}(\boldsymbol{I} + \boldsymbol{H_u}\boldsymbol{\Theta})^{\mathrm{T}}\boldsymbol{F}^{\mathrm{T}}\sqrt{S_d})^{1/2}) \\ &+ \operatorname{tr}(S_d).\end{aligned} \tag{23}$$

Remark 1. *It should be noted that $\boldsymbol{K}$ is a block lower triangular matrix whose last n_x columns are equal to 0. If we examine equation (21b), we observe that $(\boldsymbol{I} - \boldsymbol{H_u}\boldsymbol{K})^{-1}$ is a block lower triangular matrix given that $\boldsymbol{H_u}$ is also block lower triangular, which in turn implies that $(\boldsymbol{I} - \boldsymbol{H_u}\boldsymbol{K})^{-1}$ is well defined. Finally, left multiplication of $(\boldsymbol{I} - \boldsymbol{H_u}\boldsymbol{K})^{-1}$ with $\boldsymbol{K}$ gives $\boldsymbol{\Theta}$, which is also a block lower triangular matrix with the same dimension as $\boldsymbol{K}$. The reader is referred to [10], [16] for more details. An important observation is that the new decision variable $\boldsymbol{\Theta}$ should have the same structure as $\boldsymbol{K}$ for the control policy to maintain causality.*

Finally, the performance index of Problem 1 can be expressed in terms of the decision variables $\boldsymbol{u}_{\mathrm{ff}}$ and $\boldsymbol{\Theta}$. Let us denote this function as $J(\boldsymbol{u}_{\mathrm{ff}}, \boldsymbol{\Theta})$, where

$$J(\boldsymbol{u}_{\mathrm{ff}}, \boldsymbol{\Theta}) = J_1(\boldsymbol{u}_{\mathrm{ff}}) + J_2(\boldsymbol{\Theta}) + J_3(\boldsymbol{\Theta}) - J_4(\boldsymbol{\Theta}), \tag{24}$$

with

$$J_1(\boldsymbol{u}_{\mathrm{ff}}) := \|\boldsymbol{u}_{\mathrm{ff}}\|_2^2 + \lambda\|\boldsymbol{F}(\boldsymbol{\Gamma}\mu_0 + \boldsymbol{H_u}\boldsymbol{u}_{\mathrm{ff}}) - \mu_d\|_2^2, \tag{25a}$$

$$J_2(\boldsymbol{\Theta}) := \operatorname{tr}(\boldsymbol{\Theta}\tilde{S}\boldsymbol{\Theta}^{\mathrm{T}}), \tag{25b}$$

$$\begin{aligned}J_3(\boldsymbol{\Theta}) &:= \lambda \operatorname{tr}\left(\boldsymbol{F}(\boldsymbol{I} + \boldsymbol{H_u}\boldsymbol{\Theta})\tilde{S}(\boldsymbol{I} + \boldsymbol{H_u}\boldsymbol{\Theta})^{\mathrm{T}}\boldsymbol{F}^{\mathrm{T}}\right) \\ &\quad + \lambda \operatorname{tr}(S_d),\end{aligned} \tag{25c}$$

$$\begin{aligned}J_4(\boldsymbol{\Theta}) &:= 2\lambda \operatorname{tr}\Big(\big(\sqrt{S_d}\boldsymbol{F}(\boldsymbol{I} + \boldsymbol{H_u}\boldsymbol{\Theta}) \\ &\qquad \times \tilde{S}(\boldsymbol{I} + \boldsymbol{H_u}\boldsymbol{\Theta})^{\mathrm{T}}\boldsymbol{F}^{\mathrm{T}}\sqrt{S_d}\big)^{1/2}\Big).\end{aligned} \tag{25d}$$

Thus, Problem 1 can be reduced to the following optimization problem:

Problem 2. *Let $\mu_0, \mu_d \in \mathbb{R}^{n_x}$, $S_0, S_d \in \mathbb{S}_{n_x}^{++}$, $N \in \mathbb{Z}^{++}$ and $\{A_k, B_k, G_k\}_{k=0}^{N-1}$, where $A_k \in \mathbb{R}^{n_x \times n_x}$, $B_k \in \mathbb{R}^{n_x \times n_u}$ and $G_k \in \mathbb{R}^{n_x \times n_w}$, be given. Find a pair $(\boldsymbol{u}_{\mathrm{ff}}^{\star}, \boldsymbol{\Theta}^{\star})$, where $\boldsymbol{\Theta}^{\star} \in \mathbb{R}^{n_u N \times n_x(N+1)}$ is a block lower triangular matrix, that minimizes the objective function $J(\boldsymbol{u}_{\mathrm{ff}}, \boldsymbol{\Theta})$, which is defined in (24)-(25).*

Proposition 1. *Let $\lambda \in \mathbb{R}^{++}$ be given. Then, the functions J_1, J_2, J_3 and J_4, which are defined in (25), are convex and thus Problem 2 corresponds to a difference of convex functions program (DCP).*

Proof. The proof of convexity of the functions $J_1(\cdot)$, $J_2(\cdot)$ and $J_3(\cdot)$ can be found in [10]. For the convexity of $J_4(\cdot)$, we need to define the functions $g(\boldsymbol{\Theta}) := (\sqrt{S_d}\boldsymbol{F}(\boldsymbol{I}+\boldsymbol{H_u}\boldsymbol{\Theta})\tilde{V}\tilde{D}^{1/2})^T$, where $\tilde{V}^{\mathrm{T}}\tilde{D}\tilde{V}$ is the eigenvalue decomposition of $\tilde{S}$, and $f(\boldsymbol{\mathcal{A}}) := \operatorname{tr}((\boldsymbol{\mathcal{A}}^{\mathrm{T}}\boldsymbol{\mathcal{A}})^{1/2}) = \|\boldsymbol{\mathcal{A}}\|_*$. Clearly, $g(\cdot)$ is an affine function. In addition, $f(\cdot)$ corresponds to the nuclear norm, which is a valid matrix norm [17] and thus, $f(\cdot)$ is a convex function. Finally, $J_4(\boldsymbol{\Theta})$ is convex as the composition of the convex function $f(\cdot)$ with the affine function $g(\cdot)$. ☐

Remark 2. *Proposition 1 implies that Problem 1 can be reduced to a DCP, whose (local) minimizers can be found by means of the so-called convex-concave procedure [1], [13] which is known to be efficient and robust in practice.*

Remark 3. *In the formulation of Problem 2, we do not consider state or control constraints because our focus is on studying the role of the Wasserstein terminal cost in covariance steering problems. It is worth mentioning, however, that such constraints can be easily incorporated in our optimization-based approach. For instance, we can impose an explicit upper bound on the expected value of the control effort that can be used similarly to [10] (the latter constraint actually corresponds to a convex constraint).*

IV. Numerical Experiments

In this section, we present numerical experiments to demonstrate the efficacy of our approach. All computations were run on a laptop with 2.8 GHz Intel Core i7-7700HQ CPU and 16 GB RAM. For our simulations, we have used the convex-concave procedure (CCP) with MOSEK [18] to solve Problem 2 and CVXPY [19] to model the convexified sub-problems. We compare our method with general purpose NLP solvers in terms of computation time. In particular, the NLP solvers used are the IPOPT [20] and the L-BFGS-B implementation of the scipy optimization package [21]. First, we consider a discrete-time state space model of a double integrator system which is described by (3) with the following parameters: $A_k = \left[\begin{smallmatrix} 1 & \Delta t \\ 0 & 1 \end{smallmatrix}\right]$, $B_k = [0 \ \Delta t]^{\mathrm{T}}$, $G_k = \boldsymbol{I}$, $w_k \sim \mathcal{N}(\mathbf{0}, \gamma \boldsymbol{I})$, $\forall k \in \mathbb{Z}^{+}$. We also take initial state $x_0 \sim \mathcal{N}(\mu_0, S_0)$, $\mu_0 = [0, 1]^{\mathrm{T}}$, $S_0 = 10\boldsymbol{I}$, desired terminal distribution $\mathcal{N}(\mu_d, S_d)$, with $\mu_d = [10, 12]^{\mathrm{T}}$ and $S_d = \boldsymbol{I}$, and sampling period $\Delta t = 1$. In addition, we consider different experiments for $N \in \{10, 20, 30, 40, 50\}$ and $\gamma \in \{1, 0.5\}$.

Figure 1 illustrates sample trajectories of the closed-loop system along with the corresponding trajectory of the 2-σ confidence ellipsoids which shows the time-evolution of the state distribution. Fig. 2 illustrates how the error of the final covariance matrix, which is measured by $\|S_N - S_d\|_F$, changes for different values of the noise intensity parameter γ and the terminal cost weight parameter λ (all the results shown in this figure were computed for $N = 30$). We observe that as λ increases the covariance error decreases. However, the noise prevents the covariance from "shrinking" to the desired covariance. If the noise intensity is lower than some threshold for a given covariance matrix, increasing λ would eventually make the error vanish as shown in the case where $\gamma = 0.1$. In all the other cases, the final covariance error converged to a non-zero value for sufficiently large λ. Furthermore, in Table I, we compare the computation time of our CCP based approach with off-the shelf NLP solvers.

In our simulations, we have used different values for the problem horizon N and the noise intensity γ. The termination condition for the simulations is $(f_k - f_{k-1})/f_k \leq \epsilon$ with $\epsilon = 10^{-6}$ where f_k is the value of the minimized objective function at the k^{th} iteration of CCP and ϵ is the convergence tolerance. Since this is an unconstrained minimization problem, we select termination parameters adapted to the NLP solvers used to ensure that the returned solutions are close to the solution of the CCP approach. In all our experiments, the objective value returned by CCP was lower than the values returned by the other solvers. The results denoted as '-' correspond to the case in which the solver didn't terminate successfully whereas the asterisk (*) indicates that the returned optimal solution is higher than the ones returned by the other solvers.

Table I reports the computation time of one experiment for different values of certain parameters that appear in the proposed problem formulation. We observe that the proposed CCP-based approach is at least 3 times faster than general purpose NLP solvers in our experiments. One important observation is that the returned solutions are the same for CCP-based approach, however, the NLP solvers returned different locally optimal solutions for some cases (the differences of these solutions were non-negligible and cannot be attributed to numerical fluctuations). This result invites us to further study the existence and uniqueness of local minimizers of the objective function which is defined in (25) in our future work.

Furthermore, we consider a spacecraft rendezvous problem in which an active (controlled) vehicle tries to reach a passive (uncontrolled) vehicle that is moving along a circular orbit with constant angular velocity n. Let $[\delta x, \delta y]^{\mathrm{T}}$ denote the relative position of the active vehicle with respect to the passive vehicle, which is measured in meters (m), and $[\delta v^x, \delta v^y]^{\mathrm{T}}$

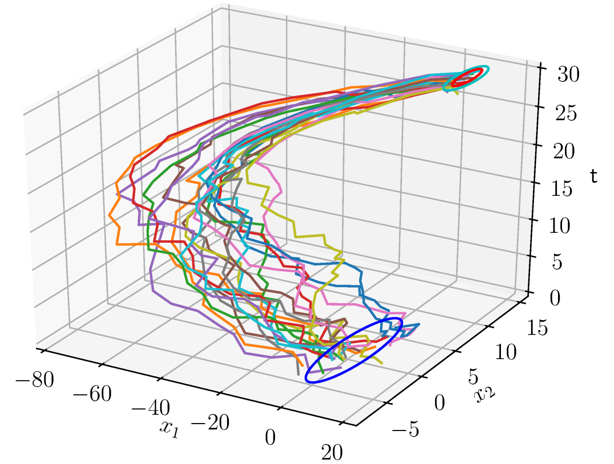

Fig. 1. Sample trajectories drawn from the optimal state process of the double integrator system with $\boldsymbol{N = 30}$, $\boldsymbol{\gamma = 1.0}$, $\boldsymbol{\lambda = 10.0}$.

TABLE I

COMPUTATION TIME (IN SECONDS) FOR DIFFERENT PROBLEM INSTANCES USING DIFFERENT SOLVERS

$\gamma = 1$	$N = 10$	$N = 20$	$N = 30$	$N = 40$	$N = 50$
L-BFGS-B	8.61	95.57	245.05	711.19	1538.85
IPOPT	17.67	119.16	353.37*	-	-
CCP	0.93	5.69	18.20	44.42	92.66
$\gamma = .5$	$N = 10$	$N = 20$	$N = 30$	$N = 40$	$N = 50$
L-BFGS-B	6.77	56.51	246.67	576.57	1097.29
IPOPT	14.77	83.01	97.89*	-	-
CCP	2.14	11.12	38.62	97.03	208.47

its relative velocity, which is measured in meters per second (m/s). The dynamics of the active vehicle is described by the so-called two-dimensional Clohessy–Wiltshire (CW) model which can be approximated by the discrete-time state space model (3) with parameters:

$$A_k = \begin{bmatrix} 1 & \Delta t & 0 & 0 \\ 3n^2 \Delta t & 1 & 0 & 2n\Delta t \\ 0 & 0 & 1 & \Delta t \\ 0 & -2n\Delta t & 0 & 0 \end{bmatrix}, \ B_k = \begin{bmatrix} 0 & 0 \\ \frac{\Delta t}{m} & 0 \\ 0 & 0 \\ 0 & \frac{\Delta t}{m} \end{bmatrix}, G_k = \frac{\sqrt{\Delta t}}{m} \boldsymbol{I},$$

where $x_k := [\delta x_k, \delta v_k^x, \delta y_k, \delta v_k^y]^{\mathrm{T}} \in \mathbb{R}^4$ is the state, $u_k \in \mathbb{R}^2$ is the control input (thrust vector) which is measured in N (Newton), and m is the mass of the active vehicle measured in kg.

For our simulations, we have used $n = 1.113 \times 10^{-3}\mathrm{rad/s}$ and $m = 100\mathrm{kg}$. The parameters that describe the discrete-time CW model (A_k, B_k, G_k) are obtained by discretization of the continuous-time CW model via a forward Euler scheme with sampling period $\Delta t = 1\,\mathrm{s}$. The means and covariances of the initial and goal state (Gaussian) distributions are $\mu_0 = [100, 0, 100, 0]^{\mathrm{T}}$, $S_0 = \mathrm{diag}(10, 1, 10, 1)$, and $\mu_d = [0, 0, 0, 0]^{\mathrm{T}}$, $S_d = \mathrm{diag}(1, 0.1, 1, 0.1)$, respectively. Also, we have taken $N = 30$, $w_k \sim \mathcal{N}(0, \gamma \boldsymbol{I})$ with $\gamma = 1.00$ (the nominal CW model is noise-free; for our simulations, we have added noise to obtain a stochastic system).

Figure 3 illustrates the evolution of the 2-σ confidence ellipsoids of the relative position of the active vehicle. The initial 2-σ ellipsoid is shown in blue, the final one (at stage $t = N$) in cyan whereas the 2-σ confidence ellipsoid of the

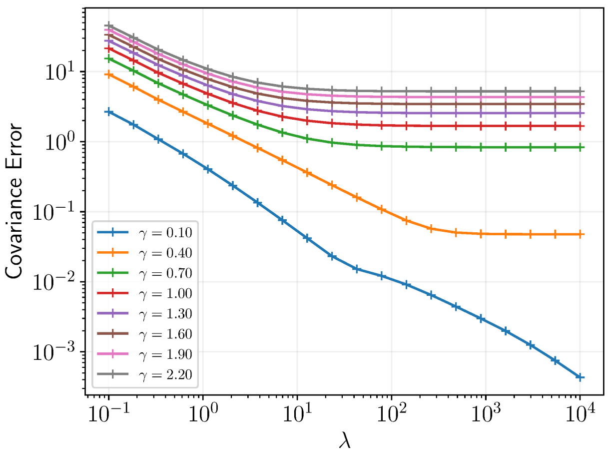

Fig. 2. Error of the final state covariance for the double integrator system v/s the parameter $\boldsymbol{\lambda}$ with $\boldsymbol{N = 30}$ and different values of γ.

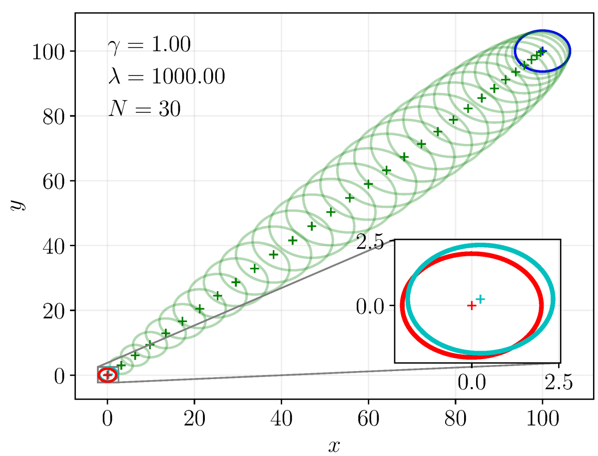

Fig. 3. Evolution of the $\boldsymbol{2\sigma}$-confidence ellipsoids of the relative position of the active vehicle from the rendezvous problem.

goal distribution is shown in red. This problem instance was solved in 282.15 seconds via CCP with the same termination condition that was used in the numerical experiments for the double integrator example with $\epsilon = 10^{-5}$. We wish to highlight here that the general purposes NLP solvers could not solve this problem in a reasonable amount of time for $N > 10$.

V. Conclusion

We have addressed the covariance steering problem with soft terminal constraints in which the terminal cost is defined as the squared Wasserstein distance between the terminal state distribution and a desired distribution. We have shown that, by utilizing an affine control policy parametrization, the proposed covariance steering problem with Wasserstein terminal cost can be reduced to a difference of convex functions program via a bijective variable transformation. The latter problem can be solved efficiently by the so-called convex-concave procedure along with convex optimization solvers. Our numerical experiments have shown that our approach reduces the computation time significantly compared to off-the-shelf nonlinear programming solvers. In our future work, we plan to further analyze the questions of existence and uniqueness of solutions to the finite-dimensional optimization problem (Problem 2) and explore alternative ways (both analytical and numerical) to characterize its minimizers. Furthermore, we plan to extend our approach to covariance steering problems for nonlinear stochastic systems.

References

[1] A. L. Yuille and A. Rangarajan, "The concave-convex procedure," *Neural computation*, vol. 15, no. 4, pp. 915–936, 2003.

[2] A. Hotz and R. E. Skelton, "Covariance control theory," *Int. J. of Control*, vol. 46, no. 1, pp. 13–32, 1987.

[3] C. Hsieh and R. E. Skelton, "All covariance controllers for linear discrete-time systems," *IEEE Trans. on Autom. Control*, vol. 35, no. 8, pp. 908–915, 1990.

[4] R. E. Skelton and T. Iwasaki, "Liapunov and covariance controllers," *Int. J. of Control*, vol. 57, no. 3, pp. 519–536, 1993.

[5] Y. Chen, T. T. Georgiou, and M. Pavon, "On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint," *J. Optim. Theory Appl.*, vol. 169, no. 2, pp. 671–691, 2016.

[6] ——, "Optimal steering of a linear stochastic system to a final probability distribution, part I," *IEEE Trans. on Autom. Control*, vol. 61, no. 5, pp. 1158–1169, 2015.

[7] ——, "Optimal steering of a linear stochastic system to a final probability distribution, part II," *IEEE Trans. on Autom. Control*, vol. 61, no. 5, pp. 1170–1180, 2015.

[8] ——, "Optimal steering of a linear stochastic system to a final probability distribution, part III," *IEEE Trans. on Autom. Control*, vol. 63, no. 9, pp. 3112–3118, 2018.

[9] E. Bakolas, "Optimal covariance control for stochastic linear systems subject to integral quadratic state constraints," in *ACC*, 2016, pp. 7231–7236.

[10] ——, "Finite-horizon covariance control for discrete-time stochastic linear systems subject to input constraints," *Automatica*, vol. 91, pp. 61 – 68, 2018.

[11] K. Okamoto, M. Goldshtein, and P. Tsiotras, "Optimal covariance control for stochastic systems under chance constraints," *IEEE Control Systems Letters*, vol. 2, no. 2, pp. 266–271, 2018.

[12] A. Halder and E. Wendel, "Finite horizon linear quadratic Gaussian density regulator with Wasserstein terminal cost," in *ACC*, 2016, pp. 7249–7254.

[13] X. Shen, S. Diamond, Y. Gu, and S. Boyd, "Disciplined convex-concave programming," in *CDC*, 2016, pp. 1009–1014.

[14] C. R. Givens and R. M. Shortt, "A class of Wasserstein metrics for probability distributions." *Michigan Math. J.*, vol. 31, no. 2, pp. 231–240, 1984.

[15] N. V. T. R. Horst, "DC programming: overview," *J. Optim. Theory Appl.*, vol. 103, no. 1, pp. 1–43, 1999.

[16] J. Skaf and S. P. Boyd, "Design of affine controllers via convex optimization," *IEEE Trans. on Autom. Control*, vol. 55, no. 11, pp. 2476–2487, 2010.

[17] R. A. Horn and C. R. Johnson, *Matrix analysis*. Cambridge university press, 2012.

[18] A. Mosek, "The MOSEK optimization software," *Online at http://www.mosek. com*, vol. 54, no. 2-1, p. 5, 2010.

[19] S. Diamond and S. Boyd, "CVXPY: A Python-embedded modeling language for convex optimization," *J. of Mach. Learn. Res.*, vol. 17, no. 83, pp. 1–5, 2016.

[20] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical programming*, vol. 106, no. 1, pp. 25–57, 2006.

[21] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright *et al.*, "Scipy 1.0: fundamental algorithms for scientific computing in Python," *Nature methods*, vol. 17, no. 3, pp. 261–272, 2020.