Greedy Finite-Horizon Covariance Steering for Discrete-Time Stochastic Nonlinear Systems Based on the Unscented Transform

Efstathios Bakolas Alexandros Tsolovikos

Abstract— **In this work, we consider the problem of steering the first two moments of the uncertain state of a discrete-time nonlinear stochastic system to prescribed goal quantities at a given final time. In principle, the latter problem can be formulated as a density tracking problem, which seeks for a feedback policy that will keep the probability density function of the state of the system close, in terms of an appropriate metric, to the goal density. The solution to the latter infinite-dimensional problem can be, however, a complex and computationally expensive task. Instead, we propose a more tractable and intuitive approach which relies on a greedy control policy. The latter control policy is comprised of the first elements of the control policies that solve a sequence of corresponding linearized covariance steering problems. Each of these covariance steering problems relies only on information available about the state mean and state covariance at the current stage and can be formulated as a tractable (finite-dimensional) convex program. At each stage, the information on the state statistics is updated by computing approximations of the predicted state mean and covariance of the resulting closed-loop nonlinear system at the next stage by utilizing the (scaled) unscented transform. Numerical simulations that illustrate the key ideas of our approach are also presented.**

I. Introduction

This paper deals with the finite-horizon covariance steering problem for discrete-time stochastic nonlinear (DTSN) systems. In particular, we consider the problem of steering the first moment (mean) and the second central moment (covariance) of the uncertain state of a DTSN system to desired quantities at a given (finite) terminal time. We will refer to the latter problem as the nonlinear covariance steering problem to emphasize the fact that it is the steering of the state covariance that constitutes the most challenging and less studied part of this stochastic control problem (steering the state mean essentially corresponds to a standard, but not necessarily trivial, controllability problem). Perhaps, one of the most natural approaches to address nonlinear covariance steering problems would be to place them under the umbrella of PDE tracking problems in which one tries to minimize the distance of the probability density of the state, which evolves in space and time in accordance with the Fokker Planck partial differential equation (PDE), from a desired terminal density function [1]. The solution to the latter infinite-dimensional optimization problem, however, can be a very complex task in general. In this work, we will employ a more practical approach that relies on the solution of a sequence of linearized steering problems which are in turn reduced to tractable convex optimization problems.

E. Bakolas is an Associate Professor and A. Tsolovikos is a PhD student in the Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin, Austin, Texas 78712-1221, USA, Email: bakolas@austin.utexas.edu; tsolovikos@austin.utexas.edu. This work was supported in part by the National Science Foundation (award no. CMMI-1753687).

Literature Review: In the special case of linear Gaussian systems, that is, stochastic linear systems subject to Gaussian white noise, the covariance steering problem corresponds to a distribution steering problem, in the sense that the mean and covariance of the terminal state uniquely determine the (Gaussian) probability distribution of the latter state. Infinite-horizon covariance steering (also known as covariance control) problems for both continuous-time and discrete-time Gaussian systems have been studied extensively by Skelton and his co-authors in a series of papers (see, for instance, [2]–[6]). The finite-horizon problem for the continuous time case was recently revisited and studied in detail in [7], [8] whereas the same problem for the discrete-time case was studied in [9]–[11]. The previous references assume perfect state information (that is, the realization of the state process at each state can be measured perfectly). Covariance control problems in the case of incomplete and imperfect state information have been studied in [12]–[14]. Nonlinear density steering problems for feedback linearizable nonlinear systems were recently studied in [15]. An iterative covariance steering algorithm for nonlinear systems based on a simple linearization of the system dynamics along reference state and input trajectories can be found in [16]. Stochastic nonlinear model predictive control with probabilistic constraints can be found in [17]–[20].

Main Contribution: In this work, we propose a greedy, yet practical and intuitive, solution approach to the nonlinear covariance steering problem. The proposed approach consists of three key steps which are applied iteratively. In the first step, we linearize the system dynamics around the current state of the system (rather than along a reference trajectory). The particular linearization scheme relies on information available at the current stage and in particular, knowledge of approximations of the mean and covariance of the current state of the system. At each new stage, a new linearization will be computed to account for the new information that becomes available at that stage. We refer to the first step as the *recursive linearization* step (RL step). In the second step, we compute a feedback control policy (sequence of feedback control laws) that solves a relevant linear, Gaussian covariance steering problem based on available approximations of the current state mean and covariance and the linear state space model computed at the LN step. The latter policy can be computed in real-time by means of tractable convex optimization techniques by leveraging the results of our previous work in covariance steering problems for Gaussian linear systems [9]–[11]. From the computed policy, only the first control law is executed at each stage. We refer to the latter step as the *linearized Gaussian covariance steering* step (LGCS step).

In the third step, we compute approximations of the one-stage-predictions of the state mean and covariance of the closed-loop system that results by applying the feedback control policy computed at the LGCS step. To compute these approximations, we employ the (scaled) unscented transform [21], [22]. The latter transform relies on the propagation of a small number of points, which are known as "sigma points," in future stages. These points are selected in a deterministic way such that their mean and variance are compatible with prior information [21]. The predicted state mean and covariance of the closed-loop system determine a Gaussian (or normal) approximation of the (predicted) state statistics of the next state. For this reason, we shall refer to the latter step as the *predictive normalization* step (PN step). This three-step process is repeated iteratively until the final stage, when it is expected that the (terminal) state mean and covariance are sufficiently close to the goal quantities.

The previously described iterative process corresponds to an on-line (or real-time) greedy control policy for nonlinear covariance steering. Because predictions of the state statistics in this approach do not go beyond the next stage, there cannot be explicit performance considerations as in a typical model predictive control approach [23]. Instead, the emphasis of the proposed greedy approach is placed on satisfying as closely as possible the boundary conditions (by steering the state mean and state covariance to desired prescribed quantities).

Structure of the paper: The rest of the paper is organized as follows. In Section II, we formulate the nonlinear covariance steering problem. A greedy algorithm for the solution to the latter problem is presented in Section III. Furthermore, we present numerical simulations in Section IV and we conclude the paper with a number of remarks and directions for future research in Section V.

II. Problem Formulation

A. Notation

We denote by $\mathbb{R}^n$ the set of n-dimensional real vectors. Given integers α, β with $\alpha \leq \beta$, we denote by $[\alpha, \beta]_d$, the discrete interval from α to β. We denote by $\mathbb{E}[\cdot]$ the expectation operator. Given a random vector x, we denote by $\mathbb{E}[x]$ its mean and by $\mathrm{Cov}[x]$ it covariance, where $\mathrm{Cov}[x] := \mathbb{E}[(x - \mathbb{E}[x])(x - \mathbb{E}[x])^{\mathrm{T}}]$. 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. Finally, we write $\mathrm{bdiag}(A_1, \ldots, A_\ell)$ to denote the block diagonal matrix formed by the matrices A_i, $i \in \{1, \ldots, \ell\}$.

B. Problem setup

We consider the following discrete-time nonlinear stochastic system

$$x(t+1) = f(x(t), u(t)) + w(t), \tag{1}$$

for $t \in [0, N-1]_d$, where N is a positive integer, and $x(0) = x_0$, where x_0 is a random vector with $\mathbb{E}[x_0] = \mu_0$ and $\mathrm{Cov}[x_0] = \Sigma_0$, with $\mu_0 \in \mathbb{R}^n$ and $\Sigma_0 \in \mathbb{S}_n^{++}$ be given quantities. Furthermore, $f(\cdot)$ is a C^1 function. In addition, $x_{0:N} := \{x(t) \in \mathbb{R}^n : t \in [0, N]_d\}$ and $u_{0:N-1} := \{u(t) \in \mathbb{R}^m : t \in [0, N-1]_d\}$ correspond to the state and input (random) processes, respectively. In addition, $w_{0:N-1} := \{w(t) \in \mathbb{R}^n : t \in [0, N-1]_d\}$ corresponds to the noise process which is assumed to be a sequence of independent and identically distributed random variables with

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

for all $t, \tau \in [0, N-1]_d$, where $W_t \in \mathbb{S}_n^+$, and $\delta(t, \tau) := 1$, when $t = \tau$, and $\delta(t, \tau) := 0$, otherwise. Furthermore, x_0 is independent of $w_{0:N-1}$, that is,

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

for all $t \in [0, N-1]_d$. Finally, throughout this paper we assume that we have perfect state information, that is, at each stage t, the realization $x(t)$ of the state process is perfectly known (measured).

Because the system given in (1) is nonlinear, even if the initial state is drawn from a normal distribution and the noise is white Gaussian, it is not guaranteed that the state at future stages will remain Gaussian. For this reason, it is not meaningful to require that the terminal state of the system should be steered to a prescribed normal distribution as in the standard formulation of a finite-horizon covariance steering problem for Gaussian linear systems. A more practical approach would be to require that the state mean and covariance of the nonlinear system attain (exactly or approximately) prescribed quantities. In particular, let us denote by $\mu_x(t)$ and $\Sigma_x(t)$ the state mean and covariance at stage t, that is,

$$\mu_x(t) := \mathbb{E}[x(t)], \quad \Sigma_x(t) := \mathrm{Cov}[x(t)]. \tag{4}$$

The class of admissible control policies is taken to be the set of sequences of control laws that are measurable functions of the realization of the current state of the system. Then, the nonlinear covariance steering problem can be formulated as follows:

Problem 1: Let $\mu_0, \mu_{\mathrm{f}} \in \mathbb{R}^n$ and $\Sigma_0, \Sigma_{\mathrm{f}} \in \mathbb{S}_n^{++}$ be given. Find a control policy $\pi := \{\kappa(t, \cdot) : t \in [0, N-1]_d)\}$ that will steer the system (1) from $x(0) = x_0$ with $\mathbb{E}[x_0] = \mu_0$ and $\mathrm{Cov}[x_0] = \Sigma_0$ to a terminal state $x(N)$ with

$$\mu_x(N) = \mu_{\mathrm{f}}, \quad (\Sigma_{\mathrm{f}} - \Sigma_x(N)) \in \mathbb{S}_n^+. \tag{5}$$

C. Collection of Finite-Horizon Linearized Covariance Steering Problems

Next, we associate the DTSN system (1) at stage $t = k \in [0, N-1]_d$ with a discrete-time stochastic linear system. The latter system corresponds to a linearization of the DTSN system around a given point $(\mu_k, \nu_k) \in \mathbb{R}^n \times \mathbb{R}^m$ which is given by

$$\begin{aligned} z(t+1) = A_k(z(t) - \mu_k) + B_k(u(t) - \nu_k) \\ + r_k + w(t), \end{aligned} \tag{6}$$

for $t \in [k, N-1]_d$ and $z(k) = \mathsf{z}_k$, with $\mathbb{E}[\mathsf{z}_k] = \mu_k$ and $\mathrm{Cov}[\mathsf{z}_k] = \Sigma_k$, where $\mu_k \in \mathbb{R}^n$, $\Sigma_k \in \mathbb{S}_n^{++}$, and $\nu_k \in \mathbb{R}^m$. In addition, it is assumed that

$$\mathbb{E}\left[\mathsf{z}_k w(t)^{\mathrm{T}}\right] = 0, \quad \mathbb{E}\left[w(t)\mathsf{z}_k^{\mathrm{T}}\right] = 0, \tag{7a}$$

for all $t \in [k, N-1]_d$. Furthermore, A_k and B_k are constant (time-invariant) matrices whereas r_k is a constant vector, and in particular,

$$A_k := \frac{\partial}{\partial x} f(x,u)\Big|_{\substack{x=\mu_k \\ u=\nu_k}}, \quad B_k := \frac{\partial}{\partial u} f(x,u)\Big|_{\substack{x=\mu_k \\ u=\nu_k}}, \tag{8a}$$

$$r_k := f(\mu_k, \nu_k). \tag{8b}$$

We can equivalently write (6) in a slightly more compact form as follows:

$$z(t+1) = A_k z(t) + B_k u(t) + d_k + w(t), \tag{9}$$

where $d_k := -A_k\mu_k - B_k\nu_k + r_k$.

We will refer to the latter linear model as the k-th linearized state space model ($t = k$ corresponds to the initial stage). It is worth noting that the triple (A_k, B_k, r_k) remains constant throughout the whole horizon $[k, N-1]_d$. However, for a different k, one obtains a different linearized system with a different but time-invariant triplet (A_k, B_k, r_k). Therefore, (6) describes essentially a collection of $N - k$ different (one for each k) time-invariant systems. An implicit assumption here is that the pair (A_k, B_k) is controllable.

An alternative linear model can be derived if one linearizes the DTSN system (1) around $(\mu_{\mathrm{f}}, \nu_{\mathrm{f}}) \in \mathbb{R}^n \times \mathbb{R}^m$, where μ_{f} is the goal mean of the terminal state and ν_{f} is such that $\mu_{\mathrm{f}} = f(\mu_{\mathrm{f}}, \nu_{\mathrm{f}})$. Then, $A_k = A_{\mathrm{f}}$, $B_k = B_{\mathrm{f}}$, and $r_k = r_{\mathrm{f}}$, where

$$A := \frac{\partial}{\partial x} f(x,u)\Big|_{\substack{x=\mu_{\mathrm{f}} \\ u=\nu_{\mathrm{f}}}}, \quad B := \frac{\partial}{\partial u} f(x,u)\Big|_{\substack{x=\mu_{\mathrm{f}} \\ u=\nu_{\mathrm{f}}}} \tag{10a}$$

$$r_k = f(\mu_{\mathrm{f}}, \nu_{\mathrm{f}}), \tag{10b}$$

for all $k \in [0, N-1]$. The previous linearization assumes that the DTSN system operates "near" the terminal target point $(\mu_{\mathrm{f}}, \nu_{\mathrm{f}}) \in \mathbb{R}^n \times \mathbb{R}^m$.

Note that both of the previously described linearized models are different from the one obtained after linearizing a nonlinear system around a given pair of reference state and input sequences $\bar{z}_{0:N} := \{\bar{z}(t) : t \in [0, N]_d\}$ and $\bar{u}_{0:N-1} := \{\bar{u}(t) : t \in [0, N-1]_d\}$, respectively, as is proposed, for instance, in [16]. In the latter case, one would consider a single time-varying linearized system described by the following equation:

$$z(t+1) = A(t)z(t) + B(t)u(t) + r(t) + w(t), \tag{11}$$

for $t \in [0, N-1]_d$, where $A(t)$, $B(t)$, and $r(t)$ are time-varying matrices which are defined as follows:

$$A(t) := \frac{\partial}{\partial x} f(x,u)\Big|_{\substack{x=\bar{z}(t) \\ u=\bar{u}(t)}}, \quad B(t) := \frac{\partial}{\partial u} f(x,u)\Big|_{\substack{x=\bar{z}(t) \\ u=\bar{u}(t)}},$$

$$r(t) := f(\bar{z}(t), \bar{u}(t)) - A(t)\bar{z}(t) - B(t)\bar{u}(t),$$

for all $t \in [0, N-1]_d$.

However, finding a reference state sequence $\bar{z}_{0:N}$ and a corresponding (compatible) reference input sequence $\bar{u}_{0:N-1}$ may be a non-trivial task. In particular, the reference state sequence should satisfy the desired boundary conditions whereas the reference input sequence should generate the corresponding reference state sequence.

Next, we will formulate a linearized covariance steering problem for the system described in (6) for a given $k \in [0, N-1]_d$. The class $\mathcal{U}$ of admissible control policies for the latter problem will consist of sequence of control laws $\{\phi_k(t, \cdot) : \ t \in [k, N-1]_d\}$, where

$$\phi_k(t, z) = \upsilon_k(t) + K_k(t)z, \quad t \in [k, N-1]_d. \tag{12}$$

We next formulate the linearized covariance steering based on information available at stage $t = k$.

Problem 2 (k-th linearized covariance steering problem): Let $\mu_k, \mu_{\mathrm{f}} \in \mathbb{R}^n$ and $\Sigma_k, \Sigma_{\mathrm{f}} \in \mathbb{S}_n^{++}$ be given. Among all admissible control policies $\varpi_k := \{\phi_k(k, \cdot), \ldots, \phi_k(N-1, \cdot)\} \in \mathcal{U}$, where $\phi_k(t, \cdot)$ satisfies (12) for $t \in [k, N-1]_d$, find a control policy $\varpi_k^\star$ that minimizes the following performance index

$$J_k(\varpi_k) := \mathbb{E}\Big[\sum_{t=k}^{N-1} \phi_k(t, z(t))^{\mathrm{T}} \phi_k(t, z(t))\Big] \tag{13}$$

subject to the recursive dynamic constraints (6) and the following boundary conditions:

$$\mathbb{E}[\mathsf{z}_k] = \mu_k, \qquad \mathrm{Cov}[\mathsf{z}_k] = \Sigma_k, \tag{14a}$$

$$\mathbb{E}[z(N)] = \mu_{\mathrm{f}}, \qquad (\Sigma_{\mathrm{f}} - \mathrm{Cov}[z(N)]) \in \mathbb{S}_n^{+}. \tag{14b}$$

Remark 1 The choice of the performance index is to ensure that the control input will have finite energy and thus avoid excessive actuation (as we have already mentioned, performance considerations are not of primary interest in this work). Problem 2 does not correspond to a standard finite-horizon linear quadratic Gaussian (LQG) problem due to the presence of the (non-standard) terminal positive semi-definite constraint $(\Sigma_{\mathrm{f}} - \mathrm{Cov}[\Sigma_z(N)]) \in \mathbb{S}_n^{+}$. Although we do not explicitly consider state or input constraints in the formulation of Problem 2, we will present an optimization-based solution which is also applicable to more general problem formulations including those with input and/or state constraints as in [11].

Remark 2 Note that finding a policy ϖ_k that solves Problem 2 is equivalent to finding a sequence $\{(\upsilon_k(t), K_k(t)) : t \in [k, N-1]_d\}$. The main idea of the proposed solution approach is that the first control law of the control policy that solves the k-th linearized covariance steering problem (Problem 2) can be used as the control law corresponding to the stage $t = k$ from the control policy that is a candidate solution to the nonlinear covariance steering problem (Problem 1). Later on, we will see that this idea will have to be applied iteratively in the sense that the data of the linear covariance problem will change at each stage, and consequently, the corresponding feedback policy has to be updated accordingly to reflect the new information available.

D. Solution to the k-th Linearized Covariance Steering

Next, we will present the main steps of the solution to the k-th linearized covariance steering problem (Problem 2). To this aim, Eq. (6) can be written in compact form as follows:

$$\boldsymbol{z} = \mathbf{G}_z^k \mathsf{z}_k + \mathbf{G}_u^k \boldsymbol{u} + \mathbf{G}_w^k (\boldsymbol{w} + \boldsymbol{d}_k), \tag{15}$$

where

$$\boldsymbol{z} := [z(k)^{\mathrm{T}}, \ldots, z(N)^{\mathrm{T}}]^{\mathrm{T}}, \quad \boldsymbol{u} := [u(k)^{\mathrm{T}}, \ldots, u(N-1)^{\mathrm{T}}]^{\mathrm{T}}$$
$$\boldsymbol{w} := [w(k)^{\mathrm{T}}, \ldots, w(N-1)^{\mathrm{T}}]^{\mathrm{T}} \quad \boldsymbol{d}_k := [d_k^{\mathrm{T}}, \ldots, d_k^{\mathrm{T}}]^{\mathrm{T}}.$$

In addition, $\mathbf{G}_u^k$, $\mathbf{G}_w^k$, and $\mathbf{G}_z^k$ are defined as follows:

$$\mathbf{G}_u^k := \begin{bmatrix} 0 & 0 & \ldots & 0 \\ B_k & 0 & \ldots & 0 \\ A_k B_k & B_k & \ldots & 0 \\ \vdots & \vdots & \ldots & \vdots \\ A_k^{N-1-k} B_k & A_k^{N-2-k} B_k & \ldots & B_k \end{bmatrix},$$
$$\mathbf{G}_w^k := \begin{bmatrix} 0 & 0 & \ldots & 0 \\ I & 0 & \ldots & 0 \\ A_k & I & \ldots & 0 \\ \vdots & \vdots & \ldots & \vdots \\ A_k^{N-1-k} & A_k^{N-2-k} & \ldots & I \end{bmatrix},$$
$$\mathbf{G}_z^k := \begin{bmatrix} I & A_k^{\mathrm{T}} & \ldots & (A_k^{N-k})^{\mathrm{T}} \end{bmatrix}^{\mathrm{T}}.$$

In view of (12), an admissible control sequence can be written compactly a follows:

$$\boldsymbol{u} = \boldsymbol{\mathcal{K}}_k \boldsymbol{z} + \boldsymbol{\upsilon}_k, \tag{16}$$

where

$$\boldsymbol{\mathcal{K}}_k := [\mathrm{bdiag}(K_k(k), \ldots, K_k(N-1)),\ 0],$$
$$\boldsymbol{\upsilon}_k := [\upsilon_k(k)^{\mathrm{T}}, \ldots, \upsilon_k(N-1)^{\mathrm{T}}]^{\mathrm{T}}.$$

Consequently, after plugging (16) into (15), we can express the closed-loop dynamics in compact form as follows:

$$\boldsymbol{z} = \mathbf{T}_z^k \mathsf{z}_k + \mathbf{T}_\upsilon^k \boldsymbol{\upsilon}_k + \mathbf{T}_w^k (\boldsymbol{w} + \boldsymbol{d}_k) \tag{17}$$

where

$$\mathbf{T}_z^k := (I - \mathbf{G}_u^k \boldsymbol{\mathcal{K}}_k)^{-1} \mathbf{G}_z^k \tag{18a}$$
$$\mathbf{T}_\upsilon^k := (I - \mathbf{G}_u^k \boldsymbol{\mathcal{K}}_k)^{-1} \mathbf{G}_u^k \tag{18b}$$
$$\mathbf{T}_w^k := (I - \mathbf{G}_u^k \boldsymbol{\mathcal{K}}_k)^{-1} \mathbf{G}_w^k. \tag{18c}$$

Note that the matrix $(I - \mathbf{G}_u^k \boldsymbol{\mathcal{K}}_k)$ corresponds to a block lower triangular matrix whose diagonal blocks are equal to the identity matrix (for more details, the reader may refer to [11], [24]). Thus, $(I - \mathbf{G}_u^k \boldsymbol{\mathcal{K}}_k)^{-1}$ is well-defined.

In view of equation (17), (16) becomes

$$\boldsymbol{u} = \mathbf{H}_z^k \mathsf{z}_k + \mathbf{H}_\upsilon^k \boldsymbol{\upsilon}_k + \mathbf{H}_w^k (\boldsymbol{w} + \boldsymbol{d}_k), \tag{19}$$

where

$$\mathbf{H}_z^k := \boldsymbol{\mathcal{K}}_k (I - \mathbf{G}_u^k \boldsymbol{\mathcal{K}}_k)^{-1} \mathbf{G}_z^k \tag{20a}$$
$$\mathbf{H}_\upsilon^k := I + \boldsymbol{\mathcal{K}}_k (I - \mathbf{G}_u^k \boldsymbol{\mathcal{K}}_k)^{-1} \mathbf{G}_u^k \tag{20b}$$
$$\mathbf{H}_w^k := \boldsymbol{\mathcal{K}}_k (I - \mathbf{G}_u^k \boldsymbol{\mathcal{K}}_k)^{-1} \mathbf{G}_w^k. \tag{20c}$$

After plugging (19) in (13), one can obtain an expression for the (predicted) cost as a function of the decision variables $\boldsymbol{\mathcal{K}}_k$ and $\boldsymbol{\upsilon}_k$. In particular,

$$\begin{aligned} J_k(\varpi_k) &= \mathbb{E}\big[\boldsymbol{u}^{\mathrm{T}} \boldsymbol{u}\big] = \mathrm{trace}\big(\mathbb{E}\big[\boldsymbol{u}\boldsymbol{u}^{\mathrm{T}}\big]\big) \\ &= \mathrm{trace}\big(\mathbb{E}\big[\big(\mathbf{H}_z^k \mathsf{z}_k + \mathbf{H}_\upsilon^k \boldsymbol{\upsilon}_k + \mathbf{H}_w^k (\boldsymbol{w} + \boldsymbol{d}_k)\big) \\ &\qquad \times \big(\mathbf{H}_z^k \mathsf{z}_k + \mathbf{H}_\upsilon^k \boldsymbol{\upsilon}_k + \mathbf{H}_w^k (\boldsymbol{w} + \boldsymbol{d}_k)\big)^{\mathrm{T}}\big]\big) \\ &=: \tilde{J}_k(\boldsymbol{\mathcal{K}}_k, \boldsymbol{\upsilon}_k). \end{aligned} \tag{21}$$

It follows readily that

$$\begin{aligned} \tilde{J}_k(\boldsymbol{\mathcal{K}}_k, \boldsymbol{\upsilon}_k) = \mathrm{trace}\big(&\mathbf{H}_z^k (\Sigma_k + \mu_k \mu_k^{\mathrm{T}}) (\mathbf{H}_z^k)^{\mathrm{T}} \\ &+ 2\mathbf{H}_z^k \mu_k \boldsymbol{\upsilon}_k^{\mathrm{T}} (\mathbf{H}_\upsilon^k)^{\mathrm{T}} + 2\mathbf{H}_z^k \mu_k \boldsymbol{d}_k^{\mathrm{T}} (\mathbf{H}_w^k)^{\mathrm{T}} \\ &+ \mathbf{H}_\upsilon^k \boldsymbol{\upsilon}_k \boldsymbol{\upsilon}_k^{\mathrm{T}} (\mathbf{H}_\upsilon^k)^{\mathrm{T}} + 2\mathbf{H}_\upsilon^k \boldsymbol{\upsilon}_k \boldsymbol{d}_k^{\mathrm{T}} (\mathbf{H}_w^k)^{\mathrm{T}} \\ &+ \mathbf{H}_w^k (\mathbf{W}_{k:N-1} + \boldsymbol{d}_k \boldsymbol{d}_k^{\mathrm{T}}) (\mathbf{H}_w^k)^{\mathrm{T}}\big), \end{aligned} \tag{22}$$

where $\mathbf{W}_{k:N-1} := \mathrm{bdiag}(W_k, \ldots, W_{N-1})$. In the previous derivation, we have used the available information about the statistics of z_k and in particular, that $\mathbb{E}[\mathsf{z}_k] = \mu_k$, $\mathbb{E}[\mathsf{z}_k \mathsf{z}_k^{\mathrm{T}}] = \Sigma_k + \mu_k \mu_k^{\mathrm{T}}$.

Next, we express the terminal constraints in terms of the decision variables $(\boldsymbol{\mathcal{K}}_k, \boldsymbol{\upsilon}_k)$. In particular, we have

$$\begin{aligned} \mathbb{E}[z(N)] &= \mathbb{E}[\mathbf{P}_N \boldsymbol{z}] = \mathbf{P}_N \mathbb{E}[\boldsymbol{z}] \\ &= \mathbf{P}_N \big(\mathbf{T}_z^k \mu_k + \mathbf{T}_\upsilon^k \boldsymbol{\upsilon}_k + \mathbf{T}_w^k \boldsymbol{d}_k\big) \\ &=: \mathfrak{f}(\boldsymbol{\mathcal{K}}_k, \boldsymbol{\upsilon}_k), \end{aligned} \tag{23}$$

where $\mathbf{P}_N := [0, \ldots, 0,\ I]$. Therefore, the constraint $\mathbb{E}[z(N)] = \mu_{\mathrm{f}}$ can be written as follows:

$$C_1(\boldsymbol{\mathcal{K}}_k, \boldsymbol{\upsilon}_k) = 0, \quad C_1(\boldsymbol{\mathcal{K}}_k, \boldsymbol{\upsilon}_k) := \mathfrak{f}(\boldsymbol{\mathcal{K}}_k, \boldsymbol{\upsilon}_k) - \mu_{\mathrm{f}}, \tag{24}$$

where $\mathfrak{f}(\boldsymbol{\mathcal{K}}_k, \boldsymbol{\upsilon}_k)$ is given in (23). Furthermore, we have that

$$\mathrm{Cov}[z(N)] = \mathbb{E}[z(N) z(N)^{\mathrm{T}}] - \mu_{\mathrm{f}} \mu_{\mathrm{f}}^{\mathrm{T}}, \tag{25}$$

where

$$\mathbb{E}[z(N)z(N)^{\mathrm{T}}] = \mathbf{P}_N\mathbb{E}\big[\big(\mathbf{T}_z^k \mathrm{z}_k + \mathbf{T}_v^k \boldsymbol{v}_k + \mathbf{T}_w^k(\boldsymbol{w}+\boldsymbol{d}_k)\big) \times \big(\mathbf{T}_z^k \mathrm{z}_k + \mathbf{T}_v^k \boldsymbol{v}_k + \mathbf{T}_w^k(\boldsymbol{w}+\boldsymbol{d}_k)\big)^{\mathrm{T}}\big]\mathbf{P}_N^{\mathrm{T}} =: \mathfrak{g}(\boldsymbol{\mathcal{K}}_k, \boldsymbol{v}_k). \quad (26)$$

Therefore, the terminal state covariance constraint: $(\Sigma_{\mathrm{f}} - \Sigma_z(N)) \in \mathbb{S}_n^+$, can be written as the following positive semi-definite constraint:

$$C_2(\boldsymbol{\mathcal{K}}_k, \boldsymbol{v}_k) \in \mathbb{S}_n^+, \quad (27a)$$

$$C_2(\boldsymbol{\mathcal{K}}_k, \boldsymbol{v}_k) := \Sigma_{\mathrm{f}} - \mathfrak{g}(\boldsymbol{\mathcal{K}}_k, \boldsymbol{v}_k) + \mu_{\mathrm{f}}\mu_{\mathrm{f}}^{\mathrm{T}}, \quad (27b)$$

where $\mathfrak{g}(\boldsymbol{\mathcal{K}}_k, \boldsymbol{v}_k)$ is defined in (26).

Problem 3: Find a pair $(\boldsymbol{\mathcal{K}}_k^\star, \boldsymbol{v}_k^\star)$ that minimizes the predicted cost $\tilde{J}_k(\boldsymbol{\mathcal{K}}_k, \boldsymbol{v}_k)$ subject to the constraints:

$$C_1(\boldsymbol{\mathcal{K}}_k, \boldsymbol{v}_k) = 0, \qquad C_2(\boldsymbol{\mathcal{K}}_k, \boldsymbol{v}_k) \in \mathbb{S}_n^+, \quad (28)$$

where $C_1(\boldsymbol{\mathcal{K}}_k, \boldsymbol{v}_k)$ and $C_2(\boldsymbol{\mathcal{K}}_k, \boldsymbol{v}_k)$ are defined in (24) and (27b), respectively.

Problem 3 is not convex as is explained in [11]. One can associate it, however, with a convex program by applying suitable transformations to the pair of decision variables $(\boldsymbol{\mathcal{K}}_k, \boldsymbol{v}_k)$ in order to obtain a new pair of decision variables, $(\boldsymbol{\mathcal{L}}_k, \boldsymbol{\nu}_k)$, which are defined as follows [24]:

$$\boldsymbol{\mathcal{L}}_k := \boldsymbol{\mathcal{K}}_k(I - \mathbf{G}_u^k \boldsymbol{\mathcal{K}}_k)^{-1}, \quad (29a)$$

$$\boldsymbol{\nu}_k := (I + \boldsymbol{\mathcal{L}}_k \mathbf{G}_u^k)\boldsymbol{v}_k. \quad (29b)$$

As is shown in [9], [11], the predicted cost can be expressed as a convex function of the new decision variables $(\boldsymbol{\mathcal{L}}_k, \boldsymbol{\nu}_k)$; this new expression is denoted as $\mathcal{J}(\boldsymbol{\mathcal{L}}_k, \boldsymbol{\nu}_k)$. In addition, the constraint functions $C_1(\boldsymbol{\mathcal{K}}_k, \boldsymbol{v}_k)$ and $C_2(\boldsymbol{\mathcal{K}}_k, \boldsymbol{v}_k)$ become $\mathcal{C}_1(\boldsymbol{\mathcal{L}}_k, \boldsymbol{\nu}_k)$ and $\mathcal{C}_2(\boldsymbol{\mathcal{L}}_k, \boldsymbol{\nu}_k)$, respectively. In particular, $\mathcal{C}_1(\boldsymbol{\mathcal{L}}_k, \boldsymbol{\nu}_k)$ corresponds to an affine function in $(\boldsymbol{\mathcal{L}}_k, \boldsymbol{\nu}_k)$ whereas the constraint $\mathcal{C}_2(\boldsymbol{\mathcal{L}}_k, \boldsymbol{\nu}_k) \in \mathbb{S}_n^+$ can be expressed as an LMI constraint in terms of $(\boldsymbol{\mathcal{L}}_k, \boldsymbol{\nu}_k)$ as is shown in [9], [11], [13]. The reader may refer to the latter references for the technical details on the conversion of the latter problems into tractable convex programs.

E. Closed-Loop Nonlinear Dynamics and Propagation of Uncertainty

Now let $\pi = \{\kappa(t,\cdot) : t \in [0, N-1]\}$ be an admissible control policy for Problem 1. Then, the state space model of the closed loop system is given by

$$x(t+1) = f_{\mathrm{cl}}(t, x(t)) + w(t), \quad (30)$$

where

$$f_{\mathrm{cl}}(t,x) := f(x, \kappa(t,x)). \quad (31)$$

Next, we describe the main steps for the propagation of the mean and the covariance of the uncertain state of the nonlinear system described by (30) based on the (scaled) unscented transform [21], [22]. To this aim, let us assume that the mean $\mu_k := \mathbb{E}[x(k)]$ and the covariance $\Sigma_k := \mathrm{Cov}[x(k)]$ of the state of (30) are known at stage k (in practice only estimates / approximations of the latter quantities will be known). Then, we will compute $2n+1$ (deterministic) points, known as *sigma points*, by using the following equation:

$$\sigma_k^{(i)} = \begin{cases} \mu_k, & \text{if } i = 0, \\ \mu_k + \sqrt{n+\lambda}\Sigma_k^{1/2}\boldsymbol{e}_i, & \text{if } i \in [1,n]_d, \\ \mu_k - \sqrt{n+\lambda}\Sigma_k^{1/2}\boldsymbol{e}_{i-n}, & \text{if } i \in [n+1, 2n]_d, \end{cases} \quad (32)$$

where $\{\boldsymbol{e}_i : \ i \in [1,n]_d\}$ denotes the standard orthonormal basis of $\mathbb{R}^n$. To each sigma point, we associate a pair of gains $(\gamma_k^{(i)}, \delta_k^{(i)})$ where

$$\gamma_k^{(i)} = \begin{cases} \lambda/(\lambda+n), & \text{if } i = 0 \\ 1/(2(\lambda+n)), & \text{if } i \in [1,2n]_d, \end{cases} \quad (33)$$

and

$$\delta_k^{(i)} = \begin{cases} 1 - \alpha^2 + \beta + \lambda/(\lambda+n), & \text{if } i = 0, \\ 1/(2(\lambda+n)), & \text{if } i \in [1,2n]_d. \end{cases} \quad (34)$$

The parameter α determines the spread around μ_k whereas β is a positive number and $\lambda := \alpha^2 n - n$. Typically, $0 < \alpha \ll 1$ and $\beta = 2$ for Gaussian approximations as suggested in [22], [25].

Subsequently, we propagate the set of sigma points $\{\sigma_k^{(i)} : \ i \in [1, 2n+1]_d\}$ at the next stage $t = k+1$ to obtain a new set of points $\{\hat{\sigma}_{k+1}^{(i)} : \ i \in [1, 2n+1]_d\}$, where

$$\hat{\sigma}_{k+1}^{(i)} = f_{\mathrm{cl}}(k, \sigma_k^{(i)}), \quad i \in [0, 2n]_d. \quad (35)$$

Using the point-set $\{\hat{\sigma}_{k+1}^{(i)} : \ i \in [0,2n]_d\}$, one can approximate the (predicted) state mean and state covariance at stage $t = k+1$ as follows:

$$\hat{\mu}_x(k+1) = \sum_{i=0}^{2L} \gamma_k^{(i)} \hat{\sigma}_{k+1}^{(i)}, \quad (36a)$$

$$\hat{\Sigma}_x(k+1) = \sum_{i=0}^{2L} \delta_k^{(i)} (\hat{\sigma}_{k+1}^{(i)} - \hat{\mu}_x(k+1)) \times (\hat{\sigma}_{k+1}^{(i)} - \hat{\mu}_x(k+1))^{\mathrm{T}} + W_k, \quad (36b)$$

where $W_k \in \mathbb{S}_n^+$ corresponds to the noise covariance at stage $t = k$.

III. A GREEDY ALGORITHM FOR NONLINEAR COVARIANCE STEERING

The proposed algorithm consists of three main steps. We will describe these steps starting at stage $t = k$, where $k \in [0, N-1]_d$, and we will assume that approximations of the state mean $\hat{\mu}_k$, the state covariance $\hat{\Sigma}_k$, and the input mean $\hat{\nu}_k$ are known (if $k = 0$, then we set $\hat{\mu}_k = \mu_0$, $\hat{\Sigma}_k = \Sigma_0$, and $\hat{\nu}_k = 0$).

We refer to the first step as the *recursive linearization* step (RL step). In the RL step, we construct a linearization (A_k, B_k, r_k) of (1) around the point $(\hat{\mu}_k, \hat{\nu}_k)$ by using (8a)–(8b). Note that the approximations $\hat{\mu}_k$ and $\hat{\nu}_k$ will be updated at the end of each stage and consequently, the linearized model will also have to be updated at each new stage to reflect the new information and hence the "recursive" qualifier in the name of this step. We will write

$$(A_k, B_k, r_k) = \Lambda\big[\hat{\mu}_k, \hat{\nu}_k; f(\cdot)\big]. \quad (37)$$

In the second step, which we refer to as the *linearized Gaussian covariance steering* step (LGCS step), we compute a feedback control policy (sequence of feedback control laws) that solves the k-th linearized covariance steering problem (Problem 2). To solve the latter problem, we need to know the linearized model (A_k, B_k, r_k), the approximations of the predicted mean and covariance $(\hat{\mu}_k, \hat{\Sigma}_k)$ at stage k assuming that the goal state mean and state covariance $(\mu_\mathrm{f}, \Sigma_\mathrm{f})$ are known a priori. The triplet (A_k, B_k, r_k) is computed in the RL step, whereas the pair $(\hat{\mu}_k, \hat{\Sigma}_k)$ is computed at the previous stage (by executing the third step of the algorithm that will be discussed shortly next). The policy $\varpi_k^\star$ that solves the k-th linearized covariance steering problem with boundary conditions

$$\mathbb{E}[z(k)] = \hat{\mu}_k, \qquad \mathrm{Cov}[z(k)] = \hat{\Sigma}_k, \quad (38a)$$

$$\mathbb{E}[z(N)] = \mu_\mathrm{f}, \qquad (\Sigma_\mathrm{f} - \mathrm{Cov}[z(N)]) \in \mathbb{S}_n^+, \quad (38b)$$

where z corresponds to the state of the linearized system. We write

$$\varpi_k^\star := \mathcal{S}_k\big[A_k, B_k, r_k, \hat{\mu}_k, \hat{\nu}_k, \hat{\Sigma}_k\big], \quad (39)$$

where $\varpi_k^\star := \{\phi_k^\star(k,\cdot), \dots, \phi_k^\star(N-1,\cdot)\}$. The computation of the control policy $\varpi_k^\star$ can be done in real-time by means of robust and efficient convex optimization techniques (for details the reader should refer to [9], [11]). We refer to the latter step as the *linearized Gaussian covariance steering* step (LGCS step). After the computation of $\varpi_k^\star$, we extract from it its first control law, $\phi_k^\star(k,\cdot)$, that is, the control law that corresponds to stage k. We write

$$\phi_k^\star(k,z) := \mathcal{P}_1\left[\varpi_k^\star\right] = \upsilon_k^\star(k) + K_k^\star(k) z,$$

where $\mathcal{P}_1(\cdot)$ denotes the truncation operator that truncates all the elements of a sequences except from the first one. Then, we set the control law $\kappa^\star(k,\cdot)$ corresponding to the k-th element of the feedback control policy $\pi^\star$ for the original nonlinear covariance steering problem (Problem 1) to be equal $\phi_k^\star(k,\cdot)$, that is,

$$\kappa^\star(k,x) := \phi_k^\star(k,x) = \upsilon_k^\star(k) + K_k^\star(k) x, \quad (40)$$

where x is the state of the original nonlinear system. Consequently, the one-stage transition map for the closed-loop dynamics based on information available at stage k is described by the following equation:

$$x(k+1) = f_\mathrm{cl}^k(k,x) + w(k), \quad (41)$$

where

$$f_\mathrm{cl}^k(k,x) := f(x, \kappa^\star(k,x)) = f(x, \upsilon_k^\star(k) + K_k^\star(k)x). \quad (42)$$

In the third step, we compute approximations $\hat{\mu}_x(k+1)$ and $\hat{\Sigma}_x(k+1)$ of the (predicted) mean and covariance of the state of the closed-loop system at stage $t = k+1$. To this aim, we first compute a set of sigma points $\{\hat{\sigma}_k : \; k \in [0, 2n]_d\}$ and their corresponding weights $(\gamma_k^{(i)}, \delta_k^{(i)})$ based on equations (32) and (33)-(33), respectively. Next, we compute the point-set $\{\hat{\sigma}_{k+1} : \; k \in [0, 2n]_d\}$ by using equation (35) and the closed loop one-stage transition map $f_\mathrm{cl}^k(k,x)$ which is defined in (42). Subsequently, we compute $\hat{\mu}_x(k+1)$ and $\hat{\Sigma}_x(k+1)$ by using (36a)-(36b). The pair $(\hat{\mu}_x(k+1), \hat{\Sigma}_x(k+1))$ determines a Gaussian approximation of the statistics of the state of the closed-loop system at stage $t = k+1$. We set $\hat{\mu}_{k+1} := \hat{\mu}_x(k+1)$ and $\hat{\Sigma}_{k+1} := \hat{\Sigma}_x(k+1)$. Finally, we set $\hat{\nu}_{k+1} := \phi_k^\star(k+1, \hat{\mu}_{k+1})$. We refer to the third step as the *predictive normalization* step (PN step). We write

$$(\hat{\mu}_{k+1}, \hat{\nu}_{k+1}, \hat{\Sigma}_{k+1}) := \mathcal{F}_k\big[\hat{\mu}_k, \hat{\Sigma}_k; f_\mathrm{cl}^k(k,\cdot)\big]. \quad (43)$$

These three steps of the previously described iterative process are repeated for all stages $t \in [k, N-1]_d$ for a given $k \in [0, N-1]_d$. At the end of the process, the predicted approximations of the state mean and covariance are sufficiently close to their corresponding goal quantities. The output of this iterative process will be a control policy $\pi_{k:N-1}^\star := \{\kappa^\star(t,x) : \; t \in [k, N-1]_\mathbb{Z}\}$. If $k \in [1, N-1]_d$, then the policy $\pi_{k:N-1}^\star$ corresponds to the truncation of the control policy $\pi^\star$ that solves Problem 1, which is comprised of the "last" $N-k$ elements of the latter policy. If we start the iterative process at $k=0$, then the output of the process is the control policy $\pi^\star$ that solves Problem 1. The pseudocode of the previous process is given in Algorithm 1.

Algorithm 1 Computation of feeback policy $\pi_{k:N-1}^\star := \{\kappa^\star(t,\cdot) : t \in [k, N-1]_\mathbb{Z}\}$ that solves Problem 1

```
procedure GREEDY NONLINEAR COVARIANCE STEERING
Input data: N, μ_f, Σ_f, f(·)
Input variables: k, μ̂_k, ν̂_k, Σ̂_k
Output variables: π_{k:N-1}, {μ̂_x(t)}_{t=k}^N, {Σ̂_x(t)}_{t=k}^N
    for t = k : N − 1 do
        (A_t, B_t, r_t) := Λ[μ̂_t, ν̂_t, Σ̂_t]
        ϖ*_t := S_t[A_t, B_t, r_t, μ̂_t, Σ̂_t]
        κ*(t,·) := P_1[ϖ*_t]
        f^t_cl(t,·) := f(x, κ*(t,·))
        (μ̂_{t+1}, ν̂_{t+1}, Σ̂_{t+1}) := F_t[μ̂_t, Σ̂_t; f^t_cl(t,·)]
    π*_{k:N-1} := {κ*(t,x) : t ∈ [k, N−1]_Z}
```

IV. Numerical Simulations

In this section, we present numerical simulations to illustrate the basic ideas of this paper. In particular, we consider the following DTSN system:

$$x_1(t+1) = x_1(t) + \tau x_2(t), \quad (44a)$$

$$x_2(t+1) = x_2(t) - \tau(\delta x_1(t) + \zeta x_1(t)^3 + \gamma x_2(t)) + \tau u(t) + \sqrt{\tau} w(t), \quad (44b)$$

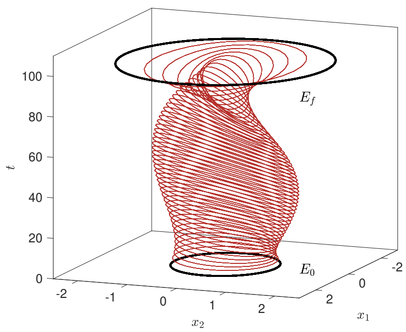

Fig. 1. Time evolution of the sequence $\{E_t\}_{t=0}^{N}$. The vertical axis in this 3D graph corresponds to the (discrete) time-axis.

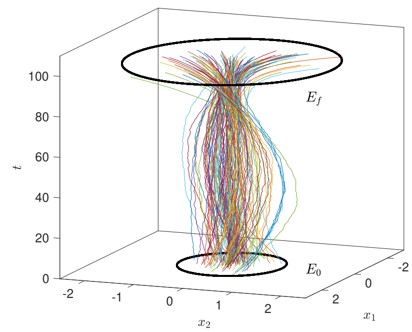

Fig. 2. Sample trajectories of the closed loop system.

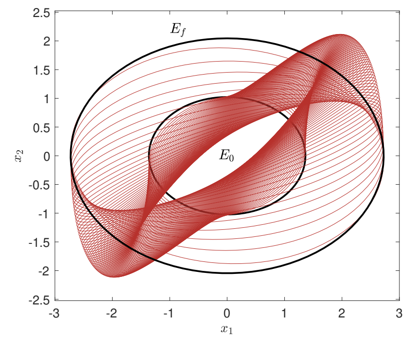

Fig. 3. Time evolution of $\{E_t\}_{t=0}^{N}$. The black ellipses correspond to E_0 and $E_{\rm f}$.

where $[x_1(0),\ \ x_2(0)]^{\rm T} \sim \mathcal{N}(\mu_0, \Sigma_0)$ with $\mu_0 = [0,\ \ 0]^{\rm T}$ and $\Sigma_0 = \mathrm{diag}(\sigma_1^2, \sigma_2^2)$, where $\sigma_1 = 2.5$ and $\sigma_2 = 2.0$. In addition, the desired terminal state mean and covariance are taken to be, respectively, $\mu_{\rm f} = [0,\ 0]^{\rm T}$ and $\Sigma_{\rm f} = \mathrm{diag}(s_1^2, s_2^2)$, where $s_1 = 1.25$ and $s_2 = 1.0$. For our simulations, we consider the following parameter values: $\tau = 0.01$, $N = 100$ and $\zeta = 0.05$, $\gamma = 0.05$, $\delta = -1$, $\alpha = 0.05$, and $\beta = 2$.

Figure 1 illustrates the time evolution of the predicted state covariance $\hat{\Sigma}_x(t)$ in terms of the evolution of the sequence of ellipsoids $\{E_t\}_{t=0}^{N}$, where

$$E_t := \{\mathsf{x} \in \mathbb{R}^2 : (\mathsf{x} - \hat{\mu}_x(t))^{\rm T} \hat{\Sigma}_x(t)^{-1} (\mathsf{x} - \hat{\mu}_x(t)) = 1\},$$

for $t \in [0, N]_d$ (the ellipsoid E_t is in an one-to-one correspondence with $\hat{\Sigma}_x(t)$). To the desired terminal state covariance $\Sigma_{\rm f}$, we associate the ellipsoid $E_{\rm f}$, where

$$E_{\rm f} := \{\mathsf{x} \in \mathbb{R}^2 : (\mathsf{x} - \mu_{\rm f})^{\rm T} \Sigma_{\rm f}^{-1} (\mathsf{x} - \mu_{\rm f}) = 1\}.$$

In particular, Fig. 1 illustrates the evolution of the sequence $\{E_t\}_{t=0}^{N}$ in a 3D graph whose vertical axis corresponds to the time-axis. Sample trajectories of the closed loop system are illustrated in Fig. 2. The projection on the $x_1 - x_2$ plane of the 3D graph given in Fig. 1 is illustrated in Fig. 3. In these three figures, the black ellipses correspond to E_0 and $E_{\rm f}$. We observe that E_N is very close to $E_{\rm f}$ and thus, the predicted covariance of the terminal state $\hat{\Sigma}_{\rm f}$ is very close to the goal covariance $\Sigma_{\rm f}$.

The evolution of the sigma points used in the unscented transform for the prediction of the state mean and covariance are illustrated in Fig. 4. In particular, the red diamonds correspond to the sigma points associated with the initial state mean and covariance, whereas the magenta circles correspond to the predicted sigma points generated for all subsequent stages $t \in [1, N]_d$. The sigma points corresponding to the original pair (μ_0, Σ_0) and the terminal pair $(\mu_{\rm f}, \Sigma_{\rm f})$ belong to the ellipses $\mathcal{E}_0$ and $\mathcal{E}_{\rm f}$, respectively, where

$$\mathcal{E}_0 := \{\mathsf{x} \in \mathbb{R}^2 : (\mathsf{x} - \mu_0)^{\rm T} \Sigma_0^{-1} (\mathsf{x} - \mu_0) = 2 + \lambda\},$$
$$\mathcal{E}_{\rm f} := \{\mathsf{x} \in \mathbb{R}^2 : (\mathsf{x} - \mu_{\rm f})^{\rm T} \Sigma_{\rm f}^{-1} (\mathsf{x} - \mu_{\rm f}) = 2 + \lambda\}.$$

V. Conclusion

In this work, we have proposed a greedy covariance steering algorithm for discrete-time stochastic nonlinear systems. The proposed approach relies on the solution of a sequence of linearized covariance steering problems combined with the (scaled) unscented transform that provides the one-stage predictions of the mean and covariance of the state of the closed loop system.

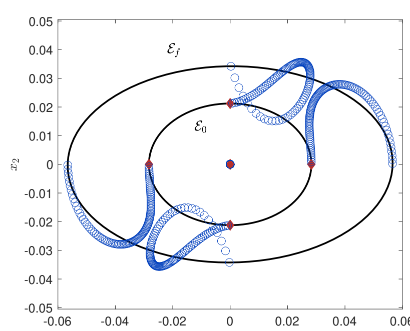

Fig. 4. Time evolution of the sigma points for $t \in [0, N]_d$. The black ellipses correspond to the ellipses $\mathcal{E}_0$ and $\mathcal{E}_{\rm f}$.

To put the presented work under the umbrella of stochastic model predictive control, it is necessary that performance and stability considerations as well as notions of invariance based on reachability analysis are integrated in the proposed algorithm. It is worth noting that the reachability analysis for nonlinear covariance steering problems requires the characterization of "admissible" sets of positive-definite matrices from which the system can be steered to the desired state terminal covariance in the given time horizon. To the best of our knowledge, the latter reachability problem constitutes, at least for the case of stochastic nonlinear systems, an open problem. In our future work, we plan to study the latter problem and we will also explore possible connections of this work with modern techniques of stochastic model predictive control.

Another important problem in the context of nonlinear covariance steering is the problem of verification of the results obtained with the proposed greedy algorithm. At present, one can expect that the predicted state mean and state covariance of the SNDT system, which are computed by means of the unscented transform, will end up sufficiently close to their goal quantities but this is not automatically the case for the true state mean and state covariance of the SNDT system. Finally, we plan to consider the case of incomplete state information and also explore connections with recent results on PDE tracking for distribution steering problems.

References

[1] A. Fleig and L. Grune, "On dissipativity of the Fokker–Planck equation for the Ornstein Uhlenbeck process," pp. 13–18, 2019.

[2] A. Hotz and R. E. Skelton, "Covariance control theory," *Int. J. Control*, vol. 16, pp. 13–32, Oct 1987.

[3] J.-H. Xu and R. E. Skelton, "An improved covariance assignment theory for discrete systems," *IEEE Trans. Autom. Control*, vol. 37, pp. 1588–1591, Oct 1992.

[4] K. Yasuda, R. E. Skelton, and K. M. Grigoriadis, "Covariance controllers: A new parametrization of the class of all stabilizing controllers," *Automatica*, vol. 29, no. 3, pp. 785–788, 1993.

[5] K. M. Grigoriadis and R. E. Skelton, "Minimum-energy covariance controllers," *Automatica*, vol. 33, no. 4, pp. 569–578, 1997.

[6] B. C. Levy and A. Beghi, "Discrete-time Gauss-Markov processes with fixed reciprocal dynamics," *J. Math. Syst. Est. Control*, vol. 7, pp. 55–80, 1997.

[7] Y. Chen, T. Georgiou, and M. Pavon, "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, 2016.

[8] Y. Chen, T. Georgiou, and M. Pavon, "Optimal steering of a linear stochastic system to a final probability distribution, Part II," *IEEE Trans. Autom. Control*, vol. 61, no. 5, pp. 1170–1180, 2016.

[9] E. Bakolas, "Optimal covariance control for discrete-time stochastic linear systems subject to constraints," in *IEEE CDC (2016)*, pp. 1153–1158, Dec 2016.

[10] M. Goldshtein and P. Tsiotras, "Finite-horizon covariance control of linear time-varying systems," in *IEEE CDC (2017)*, pp. 3606–3611, Dec. 2017.

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

[12] E. Bakolas, "Covariance control for discrete-time stochastic linear systems with incomplete state information," in *2017 American Control Conference (ACC)*, pp. 432–437, May 2017.

[13] E. Bakolas, "Finite-horizon separation-based covariance control for discrete-time stochastic linear systems," in *2018 IEEE Conference on Decision and Control (CDC)*, pp. 3299–3304, Dec 2018.

[14] E. Bakolas, "Dynamic output feedback control of the liouville equation for discrete-time siso linear systems," *IEEE Transactions on Automatic Control*, vol. 64, no. 10, pp. 4268–4275, 2019.

[15] K. F. Caluya and A. Halder, "Finite horizon density control for static state feedback linearizable systems," *arXiv preprint arXiv:1904.02272*, 2019.

[16] J. Ridderhof, K. Okamoto, and P. Tsiotras, "Nonlinear uncertainty control with iterative covariance steering," *arXiv preprint arXiv:1903.10919*, 2019.

[17] A. Mesbah, S. Streif, R. Findeisen, and R. D. Braatz, "Stochastic nonlinear model predictive control with probabilistic constraints," in *ACC (2014)*, pp. 2413–2419, June 2014.

[18] C. Liu, A. Gray, C. Lee, J. K. Hedrick, and J. Pan, "Nonlinear stochastic predictive control with unscented transformation for semi-autonomous vehicles," in *2014 American Control Conference*, pp. 5574–5579, June 2014.

[19] M. A. Sehr and R. R. Bitmead, "Particle model predictive control: Tractable stochastic nonlinear output-feedback MPC," in *20th IFAC World Congress*, pp. 15361 – 15366, 2017.

[20] J. A. Paulson and A. Mesbah, "Nonlinear model predictive control with explicit backoffs for stochastic systems under arbitrary uncertainty," in *IFAC NMPC (2018)*, pp. 523–534, 2018.

[21] S. J. Julier, "The scaled unscented transformation," in *ACC (2002)*, vol. 6, pp. 4555–4559, IEEE, 2002.

[22] S. J. Julier and J. K. Uhlmann, "Unscented filtering and nonlinear estimation," *Proceedings of the IEEE*, vol. 92, no. 3, pp. 401–422, 2004.

[23] L. Grune and J. Pannek, "Nonlinear model predictive control," in *Nonlinear Model Predictive Control*, pp. 45–69, Springer, 2017.

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

[25] E. A. Wan and R. Van Der Merwe, "The unscented Kalman filter for nonlinear estimation," in *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium*, pp. 153–158, 2000.