2021 American Control Conference (ACC)
New Orleans, USA, May 25-28, 2021

Steering the State of Linear Stochastic Systems: A Constrained Minimum Principle Formulation

Ali Pakniyat and Panagiotis Tsiotras

Abstract— **In order to optimally steer the state of a stochastic system to a desired value over a finite time horizon, a novel approach based on the Stochastic Minimum Principle is presented, which enforces a constraint on the expectation of the terminal state at all instances of time. In order to solve the associated optimal control problem, we invoke a version of the Stochastic Minimum Principle which we call the Terminally Constrained Stochastic Minimum Principle (TC-SMP). For linear stochastic systems with quadratic costs, analytical solutions to the adjoint equation of the TC-SMP are derived and are explicitly represented in terms of controllability Gramians and solutions of Riccati equations. Numerical examples are provided to illustrate the results, and the performance of the TC-SMP approach is compared to both penalty-based and covariance-steering alternative approaches.**

I. Introduction

In several engineering applications, it is desired to bring a system from an initial configuration to a specific terminal configuration. A classical example is balancing the upright configuration of an inverted pendulum. A more complex example is the vertical landing of a reusable rocket, e.g., the booster rocket of SpaceX Falcon 9, which is required to come to a full stop at an exact location on the landing platform. If dynamic uncertainties are negligible, powerful theoretical tools are available in the control theory literature, the most notable being the Pontryagin Minimum Principle (MP) [1], which determines the optimal input, among all controllers that steering the state to the desired terminal value.

In the presence of a stochastic diffusion, these state steering problems are more challenging and have been the subject of a limited number of studies. More precisely, the majority of studies assume linearity of the dynamics and a quadratic form for the cost, so that the associated probability distributions are Gaussian. In this case, and in the absence of any additional state constraints, the dynamics of the mean state process and the covariance state process are decoupled. Within an infinite time horizon setting, the problem has been formulated as the association of a steady-state distribution with its mean being at the desired terminal location, and a comprehensive study over the assignable covariances for the infinite horizon problem is presented in [2]–[5]. For linear stochastic systems over finite time horizons, a similar philosophy is taken in both continuous time and discrete time settings [6]–[14]. The accommodation of input constraints is considered in [10], and convex relaxations for linear systems subject to chance constraints, which are probabilistic constraints that impose a maximum probability of constraint violation, are studied in [12], [13]. Extensions of the probability distribution assignment to nonlinear systems has been presented for feedback-linearizable systems [15], and implementation through iterative linearization is proposed in [16].

A fundamental limit of the current methodologies based on the assignment of terminal probability distributions is that the studied probabilities are conditioned on the filtration at the initial time. For linear systems, the information obtained as time progresses is accommodated in a model predictive control (MPC) based approach in [17]–[20]. However, as proposed in this paper, the employment of the Stochastic Minimum Principle (SMP) yields a natural accommodation of filtration-adapted updates because the same adaptation requirement must be provided for the adjoint process. In other words, the optimal input expressed in terms of the adjoint process is adapted to the current time filtration, since the solution of the backward stochastic differential equation (BSDE) for the adjoint process must remain adapted to the same forward filtration. This important characteristic provides an opportunity to impose terminal state constraints at all times, as apposed to the current literature where constraints are imposed on probability distributions as viewed at the initial time. In order to solve the associated problem, we invoke the Stochastic Maximum Principle (SMP) presented in [21] and, in particular, the version with terminal state constraints [21, Theorem 5], henceforth called the Terminally Constrained Stochastic Minimum Principle (TC-SMP). While, in general, obtaining numerical solutions to the BSDEs of the adjoint process are computationally expensive, for a class of linear stochastic systems with quadratic costs, we derive analytical solutions to the adjoint equation in terms of the system's state transition matrix, its controllability Gramian and the solution of a differential matrix Riccati equation.

The organization of the paper is as follows. In Section II, the steering problem for the general nonlinear case is formulated, a representation form for the terminal state constraints is proposed, and the associated necessary optimality conditions are presented in the form of the TC-SMP. The specialization of the results to stochastic systems with linear dynamics and quadratic cost is presented in Section III, and analytical solutions to the TC-SMP are provided in terms of the controllability Gramians and the solutions of Riccati

A. Pakniyat is a post-doctoral fellow with the Institute for Robotics and Intelligent Machines, Georgia Institute of Technology, Atlanta, GA 30332, USA. Email: pakniyat@gatech.edu

P. Tsiotras is the David and Andrew Lewis Professor at the School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA. Email: tsiotras@gatech.edu

2021 American Control Conference (ACC) | 978-1-6654-4197-1/21/$31.00 ©2021 AACC | DOI: 10.23919/ACC50511.2021.9483276

equations. To illustrate the results, numerical examples are studied in Section IV and the performance of the TC-SMP approach is compared to both penalty-based approaches and the covariance control methodology. Concluding remarks are provided in Section V.

II. The General Nonlinear Case

Let $\big(\Omega, \mathcal{F}, \{\mathcal{F}^t\}_{t=t_0}^{t_f}, \mathsf{P}\big)$ be a filtered probability space with $\mathcal{F}^t$ being an increasing family of sub σ-algebras of $\mathcal{F}$ such that $\mathcal{F}^{t_0}$ contains all the P-null sets, and $\mathcal{F}^{t_f} = \mathcal{F}$ for a fixed terminal time $t_f < \infty$. Consider the nonlinear stochastic system governed by a controlled Itô differential equation

$$\mathsf{d}x_s = f(s, x_s, u_s)\,\mathsf{d}s + g(s, x_s)\,\mathsf{d}w_s, \tag{1}$$

where $x_s \in \mathbb{R}^n$, $u_s \in U \subset \mathbb{R}^m$ are, respectively, the state and the input values at instant $s \in [t_0, t_f]$, and w is a standard k-dimensional Wiener process, such that $w_s \in \mathbb{R}^k$, $\mathbb{E}[w_s] = 0$, and $\mathbb{E}[\mathsf{d}w_s\,\mathsf{d}w_s^\mathsf{T}] = \sqrt{\mathsf{d}t}\,I_{k\times k}$. In order to ensure the existence and uniqueness of solutions, we assume that f and g are smooth and bounded functions of their arguments.

Let us denote by $[u] \equiv [u]_t^{t_f} := \big\{u_s, s \in [t, t_f],\ u_s \in U : \mathcal{F}^s\!-\text{adapted}\big\}$ a general control process, and denote by $\mathcal{U}$ the set of all such inputs. In other words, $[u] \in \mathcal{U}$ whenever it is a progressively measurable process over the interval $[t, t_f]$, taking values from the set $U \subset \mathbb{R}^m$. We remark that the underlying policy for the determination of the input process $[u]$ can take any form, as long as the policy remains causal, that is, u_s does not depend on future values of the noise or the state.

In this paper, we consider only the case with complete and accurate observations of the state. Thus, for time instances t and s within the interval $[t_0, t_f]$, and under the filtration $\mathcal{F}^t$, the variable x_s is treated as a *deterministic* variable whenever $s \le t$, and is treated as a *random* variable whenever $s > t$. We define the notation

$$\mathbb{E}_{\mathcal{F}^t}^{[u]}[x_s] := \mathbb{E}\big[x_s\big|\mathcal{F}^t; [u]_t^{t_f}\big] \equiv \mathbb{E}\big[x_s\big|\mathcal{F}^t; [u]_t^{s}\big], \tag{2}$$

for the expected value of x_s at $s \in [t, t_f]$, under the filtration $\mathcal{F}^t$ and given the input process $[u]_t^{t_f}$, where the last equality (conditioning on $[u]_t^s$ instead of $[u]_t^{t_f}$) is a consequence of the causality of the controlled process in (1).

In order for the state to be steered to a desired value $\mu_f \in \mathbb{R}^n$, we enforce the constraint

$$\mathbb{E}_{\mathcal{F}^t}^{[u]}[x_{t_f}] = \mu_f, \tag{3}$$

at all time instances $t \in [t_0, t_f]$.

Remark 2.1: A significant distinction of the proposed problem formulation, in comparison to the covariance steering literature [2]–[14] for example, is that here the constraint (3) is enforced under all filtrations $\mathcal{F}^t$, $t \in [t_0, t_f]$, whereas in [2]–[14], it is enforced only under the filtration at $t = t_0$, that is, $\mathbb{E}_{\mathcal{F}^{t_0}}^{[u]}[x_{t_f}] = \mu_f$, and this terminal constraint is accompanied by a constraint on the covariance at the initial time, namely, $\mathrm{cov}_{\mathcal{F}^{t_0}}^{[u]}[x_{t_f}] = \Sigma_f$, with Σ_f a desired positive definite covariance matrix for the distribution of the terminal state conditioned on the information at t_0.

Whenever the class of controllers satisfying (3) is nonempty, the performance of such $[u]$ is evaluated by

$$J\big(t, x_t, [u]_t^{t_f}\big) := \mathbb{E}_{\mathcal{F}^t}^{[u]}\left[\int_t^{t_f} \ell(x_s, u_s)\,\mathsf{d}s + L\big(x_{t_f}\big)\right]. \tag{4}$$

The objective of the optimal control problem is to find $[u^*]$ satisfying (3) such that the cost (4) is minimized.

In order to solve this optimal control problem, we invoke the following result from [21].

Theorem 2.2: (Terminally-Constrained Stochastic Minimum Principle (TC-SMP)) For the system (1), the optimal input for the cost (4) subject to the constraint (3) is determined from

$$u_s^* = \underset{u\in\mathbb{R}^m}{\operatorname{argmin}}\Big\{\ell(x_s, u) + \lambda_s^\mathsf{T} f(x_s, u)\Big\}, \tag{5}$$

where the adjoint pair (λ_s, Λ_s), $s \in [t, t_f]$ are governed by the backward stochastic differential equation

$$\mathsf{d}\lambda_s = -\left(\frac{\partial f(x_s^*, u_s^*)}{\partial x}\lambda_s + \frac{\partial \ell(x_s^*, u_s^*)}{\partial x}\right)\mathsf{d}s + \Lambda_s\,\mathsf{d}w_s, \tag{6}$$

subject to the terminal condition

$$\lambda_{t_f} = \alpha\frac{\partial L(x_{t_f}^*)}{\partial x} + \beta, \tag{7}$$

where $\alpha \in \mathbb{R}$ and $\beta \in \mathbb{R}^n$ are constants which are not simultaneously zero. □

Proof: Please see [21, Theorem 5]. ■

III. TC-SMP for Linear Quadratic Problems

In this section, we specialize the results of Theorem 2.2 to linear stochastic systems with quadratic cost, and provide an analytical solution to the TC-SMP. To this end, let the dynamics (1) be of the form

$$\mathsf{d}x_s = \big(A_s x_s + B_s u_s\big)\,\mathsf{d}s + D_s\,\mathsf{d}w_s, \tag{8}$$

with $A \in L^\infty([t_0, t_f]; \mathbb{R}^{n\times n})$, $B \in L^\infty([t_0, t_f]; \mathbb{R}^{n\times m})$, $D \in L^\infty([t_0, t_f]; \mathbb{R}^{n\times k})$, essentially bounded measurable matrix functions of time.

For simplicity, we assume that the cost (4) is a quadratic function of the input and the terminal state, that is,

$$J\big(t, x_t, [u]_t^{t_f}\big) := \frac{1}{2}\mathbb{E}_{\mathcal{F}^t}^{[u]}\bigg[\int_t^{t_f} u_s^\mathsf{T} R_s u_s\,\mathsf{d}s + (x_{t_f} - \mu_f)^\mathsf{T} H_f\,(x_{t_f} - \mu_f)\bigg], \tag{9}$$

with $R \in L^\infty([t_0, t_f]; \mathcal{S}^{m\times m})$, $R_s > 0$, for all $s \in [t_0, t_f]$, and $H_f \in \mathcal{S}^{n\times n}$, $H_f \ge 0$, where $\mathcal{S}^{m\times m}$ denotes the space of $m \times m$-dimensional symmetric matrices.

We assume that the system (A_s, B_s) is controllable[1], and that the system is noise controllable[2], equivalently, $\mathrm{Im}(D_s) \subset \mathrm{Im}(B_s)$, for all $s \in [t_0, t_f]$, that is,

$$\forall w \in \mathbb{R}^k, \exists u \in \mathbb{R}^m \text{ s.t. } B_s u = D_s w. \tag{10}$$

The following theorem summarizes the main result of the paper.

[1]Hence, the Gramian (13) is full rank.

[2]As a requirement for solvability of the Riccati equations (14) and (15).

Theorem 3.1: For the system (8) and the cost (9) subject to the constraint (3), the optimal input is determined by

$$
\begin{aligned}
u_s^* &= -R_s^{-1}B_s^\mathsf{T}\Phi(t_f,s)^\mathsf{T}\left[\mathcal{G}(t,t_f)\right]^{-1}\left(\Phi(t_f,t)x_t-\mu_f\right)\\
&\quad - R_s^{-1}B_s^\mathsf{T}\Pi(s;t_f)\bigg(x_s-\Phi(s,t)x_t\\
&\quad + \mathcal{G}(t,s)\Phi(t_f,s)^\mathsf{T}\left[\mathcal{G}(t,t_f)\right]^{-1}\Big(\Phi(t_f,t)x_t-\mu_f\Big)\bigg),
\end{aligned}
\tag{11}
$$

where $\Phi(s,t)\in\mathbb{R}^{n\times n}$ is the state transition matrix from t to s for the system (8), which is the solution of

$$
\dot{\Phi}\equiv\frac{\partial\Phi(s,t)}{\partial s}=A_s\Phi,\qquad \Phi(t,t)=I_{n\times n},\tag{12}
$$

and where

$$
\mathcal{G}(\tau,t):=\int_t^\tau\Phi(\tau,s)B_sR_s^{-1}B_s^\mathsf{T}\Phi(\tau,s)^\mathsf{T}\,\mathsf{d}s,\tag{13}
$$

is the controllability Gramian (see e.g., [22, Theorem 6.1]) over the horizon $[t,\tau]\subset[t_0,t_f]$, and $\Pi(s;t_f)$ is the solution of the following Riccati equation

$$
\dot{\Pi}_s\equiv\frac{\mathsf{d}}{\mathsf{d}s}\Pi(s;t_f)=\Pi_sB_sR_s^{-1}B_s^\mathsf{T}\Pi_s-\Pi_sA_s-A_s^\mathsf{T}\Pi_s,\tag{14}
$$

subject to the terminal condition

$$
\Pi(t_f;t_f)=H_f.\tag{15}
$$

□

Proof: For the brevity of notation, the time index s is dropped for the matrices A_s, B_s, etc.

We first invoke (5) to obtain

$$
u_s^*=\underset{u_s\in\mathbb{R}^m}{\operatorname{argmin}}\Big\{\frac{1}{2}u_s^\mathsf{T}Ru_s+\lambda_s^\mathsf{T}\big(Ax_s+Bu_s\big)\Big\}=-R^{-1}B^\mathsf{T}\lambda_s.\tag{16}
$$

Therefore, the joint dynamics of the (optimal) state process and the adjoint process become

$$
\mathsf{d}x_s=(Ax_s-BR^{-1}B^\mathsf{T}\lambda_s)\mathsf{d}s+D\,\mathsf{d}w_s,\tag{17}
$$

$$
\mathsf{d}\lambda_s=-A^\mathsf{T}\lambda_s\mathsf{d}s+\Lambda_s\mathsf{d}w_s,\tag{18}
$$

where the state dynamics (17) and the adjoint dynamics (18) are, respectively, subject to the initial condition $x_t=x(t)$ and the terminal conditions obtained from (7) as

$$
\lambda_{t_f}=\alpha H_f\left(x_{t_f}-\mu_f\right)+\beta.\tag{19}
$$

Defining now

$$
\bar{x}_s=\mathbb{E}_{\mathcal{F}^t}^{[u^*]}(x_s),\qquad \tilde{x}_s=x_s-\bar{x}_s\tag{20}
$$

$$
\bar{\lambda}_s:=\mathbb{E}_{\mathcal{F}^t}^{[u^*]}(\lambda_s),\qquad \tilde{\lambda}_s:=\lambda_s-\bar{\lambda}_s,\tag{21}
$$

the dynamics (17) and (18) and the associated initial and terminal conditions are expressed in terms of the mean processes

$$
\mathsf{d}\bar{x}_s=(A\bar{x}_s-BR^{-1}B^\mathsf{T}\bar{\lambda}_s)\mathsf{d}s,\tag{22}
$$

$$
\mathsf{d}\bar{\lambda}_s=-A^\mathsf{T}\bar{\lambda}_s\mathsf{d}s,\tag{23}
$$

subject to the initial and terminal conditions

$$
\bar{x}_t=x_t,\qquad \bar{x}_{t_f}\overset{(3)}{=}\mu_f,\tag{24}
$$

$$
\bar{\lambda}_t=\text{free},\qquad \bar{\lambda}_{t_f}\underset{(3)}{\overset{(19)}{=}}\beta\equiv\text{free},\tag{25}
$$

together with the error processes

$$
\mathsf{d}\tilde{x}_s=(A\tilde{x}_s-BR^{-1}B^\mathsf{T}\tilde{\lambda}_s)\mathsf{d}s+D\,\mathsf{d}w_s,\tag{26}
$$

$$
\mathsf{d}\tilde{\lambda}_s=-A^\mathsf{T}\tilde{\lambda}_s\mathsf{d}s+\Lambda_s\mathsf{d}w_s,\tag{27}
$$

with the initial and terminal conditions

$$
\tilde{x}_t=0,\qquad \tilde{x}_{t_f}\overset{(3)}{=}\text{free},\tag{28}
$$

$$
\tilde{\lambda}_t=\text{free},\qquad \tilde{\lambda}_{t_f}\overset{(19)}{=}\alpha H_f\tilde{x}_s.\tag{29}
$$

Since the dynamics (23) of $\bar{\lambda}$ is decoupled from the dynamics (22) of $\bar{x}$, it yields that

$$
\bar{\lambda}_s=\Phi(t,s)^\mathsf{T}\bar{\lambda}_t,\tag{30}
$$

whose substitution in (22) yields

$$
\begin{aligned}
\bar{x}_{t_f}&=\Phi(t_f,t)x_t-\int_t^{t_f}\Phi(t_f,\theta)B\,R^{-1}B^\mathsf{T}\Phi(\theta,t)^\mathsf{T}\bar{\lambda}_t\,\mathsf{d}\theta\\
&=\Phi(t_f,t)x_t\\
&\quad-\left[\int_t^{t_f}\Phi(t_f,\theta)B\,R^{-1}B^\mathsf{T}\Phi(t_f,\theta)^\mathsf{T}\,\mathsf{d}\theta\right]\Phi(t,t_f)^\mathsf{T}\bar{\lambda}_t\\
&=\Phi(t_f,t)x_t-\mathcal{G}(t_f,t)\Phi(t,t_f)^\mathsf{T}\bar{\lambda}_t.
\end{aligned}
\tag{31}
$$

Therefore,

$$
\bar{\lambda}_t=\Phi(t_f,t)^\mathsf{T}\left[\mathcal{G}(t_f,t)\right]^{-1}\left(\Phi(t_f,t)x_t-\mu_f\right),\tag{32}
$$

which, together with (30), yields

$$
\bar{\lambda}_s=\Phi(t_f,s)^\mathsf{T}\left[\mathcal{G}(t_f,t)\right]^{-1}\left(\Phi(t_f,t)x_t-\mu_f\right).\tag{33}
$$

Furthermore, it can be easily verified that (26), (27), (28) and (29) correspond to a classical LQG system for the cost

$$
\begin{aligned}
\tilde{J}\big(t,\tilde{x}_t,[\tilde{u}]\big)&:=\mathbb{E}_{\mathcal{F}^t}^{[\tilde{u}]}\bigg[\int_t^{t_f}\frac{1}{2}\tilde{u}_s^\mathsf{T}R\tilde{u}_s\,\mathsf{d}s\\
&\quad+\frac{1}{2}\alpha(\tilde{x}_{t_f}-\mu_f)^\mathsf{T}H_f\left(\tilde{x}_{t_f}-\mu_f\right)\bigg].
\end{aligned}
\tag{34}
$$

Thus, $\alpha=1$ needs to hold for compliance with (9), and by invoking the stochastic Riccati formalism (see, e.g., [23, Chapter 6, Section 6]), we obtain

$$
\tilde{\lambda}_s=\Pi(s;t_f)\tilde{x}_s,\tag{35}
$$

$$
\Lambda_s=\Pi(s;t_f)D,\tag{36}
$$

resulting in the dynamics (14) and the terminal condition (15). Moreover, substitution of $\tilde{x}$ from the definition (20), yields

$$
\begin{aligned}
\tilde{\lambda}_s&=\Pi(s;t_f)(x_s-\bar{x}_s)=\Pi(s;t_f)\bigg(x_s-\Phi(s,t)x_t\\
&\quad+\mathcal{G}(s,t)\Phi(t_f,s)^\mathsf{T}\Phi(t_f,s)^\mathsf{T}\mathcal{G}(t_f,t)^{-1}\Big(\Phi(t_f,t)x_t-\mu_f\Big)\bigg).
\end{aligned}
\tag{37}
$$

Hence, (11) is obtained with the substitution of $\bar{\lambda}_s$ from (33) and $\tilde{\lambda}_s$ from (37) into $\lambda_s=\bar{\lambda}_s+\tilde{\lambda}_s$, and subsequently into (16). ■

IV. Numerical Examples

In this section we provide a series of examples to demonstrate the proposed approach and compare it with other standard approaches in the literature.

Example 4.1: Consider the scalar case of a linear stochastic system with the dynamics

$$\mathsf{d}x_s := (ax_s - bu_s)\,\mathsf{d}s + d\,\mathsf{d}w_s, \tag{38}$$

with a, b, d scalar constants, and consider the problem of steering the state to the desired value $\mu_f \in \mathbb{R}$ by enforcing

$$\mathbb{E}^{[u]}_{\mathcal{F}^t}[x_{t_f}] = \mu_f, \tag{3}$$

at all $t \in [t_0, t_f]$, with the cost

$$J\big(t, x_t, [u]_t^{t_f}\big) := \frac{1}{2}\mathbb{E}^{[u]}_{\mathcal{F}^t}\left[\int_t^{t_f} r\,u_s^2\,\mathsf{d}s + h(x_{t_f} - \mu_f)^2\right]. \tag{39}$$

where $r > 0$ and $h \in \mathbb{R}_{\geq 0} \setminus \{2ar/b^2\}$[3].

For this problem, we can analytically represent Φ, $\mathcal{G}$ and Π as

$$\Phi(s,t) = e^{a(s-t)}, \tag{40}$$

$$\mathcal{G}(t,\tau) = \frac{b^2}{2a}e^{2a\,t_f}\left(e^{-2a\,t} - e^{-2a\,\tau}\right), \tag{41}$$

$$\Pi(s;t_f) = \frac{2ar}{b^2\left(1 - \frac{h}{h-\frac{2ar}{b^2}}e^{\frac{b^2}{r}(t_f-s)}\right)}. \tag{42}$$

and, therefore, the optimal input (11) becomes

$$\begin{aligned} u_s^* = {} & \frac{-2a}{br\left(e^{2a(t_f-t)} - 1\right)}e^{a(t_f-s)}\left(e^{a(t_f-t)}x_t - \mu_f\right) \\ & - \frac{2a}{b\left(1 - \frac{h}{h-\frac{2ar}{b^2}}e^{\frac{b^2}{r}(t_f-s)}\right)}\Bigg(x_s - e^{a(s-t)}x_t \\ & + e^{a(t_f-s)}\frac{e^{-2at} - e^{-2as}}{e^{-2at} - e^{-2at_f}}\left(e^{a(t_f-t)}x_t - \mu_f\right)\Bigg). \end{aligned} \tag{43}$$

Let $a = b = d = r = h = 1$, and the time horizon be $[t_0, t_f] = [0, 1]$. For the steering towards the desired stated $\mu_f = 2$, from the initial condition $x_0 = -4$, the optimal input satisfying (3) for all $t \in [0, 1]$, and the associated trajectories for 50 sample paths are illustrated in Figure 1. As can be seen from the figure, all trajectories are almost surely driven to the required terminal state at $t = t_f$.

In order to illustrate the role of the constraint (3), the associated trajectories and inputs of the unconstrained (without the constraint (3)) linear quadratic system with the dynamics (38) and the cost (39) with the terminal cost $\frac{1}{2}(x_{t_f} - \mu_f)^2$ are displayed in Figure 2. It can be observed from these sample paths that neither the expected value of x_{t_f}, nor almost any of its realizations achieve the requirement $\mu_f = 2$. While it is possible to push trajectories closer to μ_f by increasing the terminal cost (i.e., by selecting larger values for h), the inherent inability of such penalty-based approaches to meet the constraint *exactly* is quite obvious, even from this simple example.

[3]The exclusion of this value, which occurs only if $a > 0$, is due to the appearance of $h - 2ar/b^2$ as a denominator in (43).

It is also worth comparing the results of TC-SMP with the covariance steering methodology (e.g., [6], [7]). By defining $\mu_s := \mathbb{E}^{[u]}_{\mathcal{F}^{t_0}}[x_s]$ and $\sigma_s := \mathbb{E}^{[u]}_{\mathcal{F}^{t_0}}[(x_s - \mu_s)^2]$ and by considering the linear feedback law $u_s = k_s x_s + v_s$, we obtain

$$\dot{\mu} = (a + b\,k_s)\,\mu_s + b\,v_s, \tag{44}$$

$$\dot{\sigma} = 2\,(a + b\,k_s)\,\sigma_s + d^2. \tag{45}$$

For the cost we have, accordingly, that

$$\begin{aligned} J'\big(t_0, x_0, [u]_{t_0}^{t_f}\big) &:= \frac{1}{2}\mathbb{E}^{[u]}_{\mathcal{F}^{t_0}}\left[\int_{t_0}^{t_f} r\,u_s^2\,\mathsf{d}s\right] \\ &= \frac{r}{2}\mathbb{E}^{[u]}_{\mathcal{F}^{t_0}}\left[\int_{t_0}^{t_f}(k_s x_s + v_s)^2\,\mathsf{d}s\right] \\ &= \frac{r}{2}\int_{t_0}^{t_f}(k_s\mu_s + v_s)^2\,\mathsf{d}s + \frac{r}{2}\int_{t_0}^{t_f}K_s^2\sigma_s^2\,\mathsf{d}s., \end{aligned} \tag{46}$$

The associated covariance steering problem is to find k_s and v_s such that $\mu_0 = x_0$, $\mu_{t_f} = \mu_f$, $\sigma_0 = 0$, $\sigma_{t_f} = \sigma_f$, where σ_f is a desired terminal state provided in the problem statement. For $x_0 = -4$, $\mu_f = 2$ and $\sigma_f = 1$, which is equivalent to assigning the normal distribution $\mathcal{N}(2,1)$ as the terminal state distribution under the filtration $\mathcal{F}^{t_0}$, the associated optimal trajectories and inputs for 50 sample paths are displayed in Figure 3. As seen from this figure, the trajectories meet the covariance constraint, but have a larger spread around the final mean value, as expected from the problem formulation.

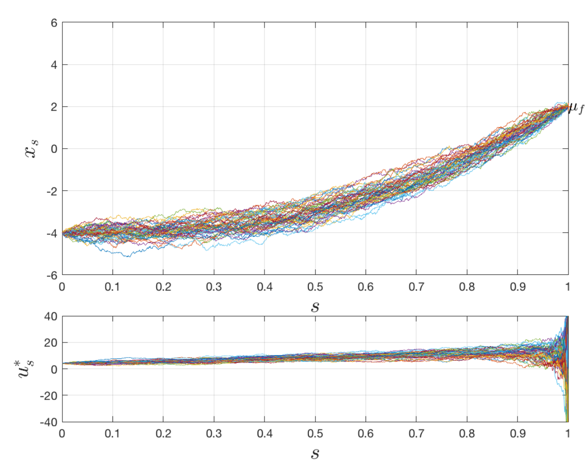

Fig. 1: The implementation of the Terminally Constrained Stochastic Minimum Principle (TC-SMP) on the system in Example 4.1.

Example 4.2: Consider the system governed by

$$\mathsf{d}x_s = \left(\begin{bmatrix} 0 & 1 \\ 1 & 2 \end{bmatrix}\begin{bmatrix} x_s^{(1)} \\ x_s^{(2)} \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix}u_s\right)\mathsf{d}s + \begin{bmatrix} 0 \\ 1 \end{bmatrix}\mathsf{d}w_s, \tag{47}$$

over the time horizon $[t_0, t_f] = [0, 1]$, starting from the initial condition $x_0 = [1, 1]^{\mathsf{T}}$, and steered towards the desired

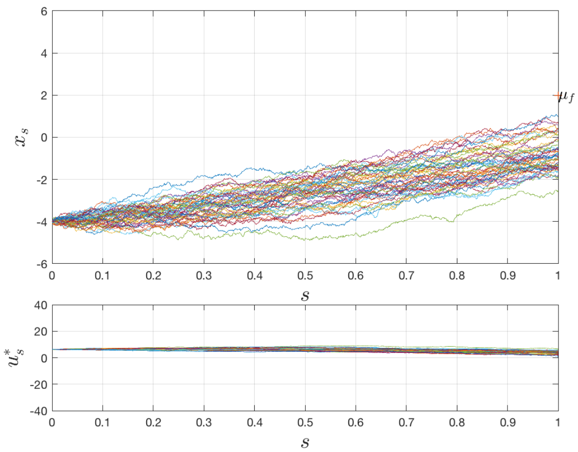

Fig. 2: The implementation of a naïve approach in attempting to reach the target via the terminal cost $\frac{1}{2}(x_{t_f} - \mu_f)^2$.

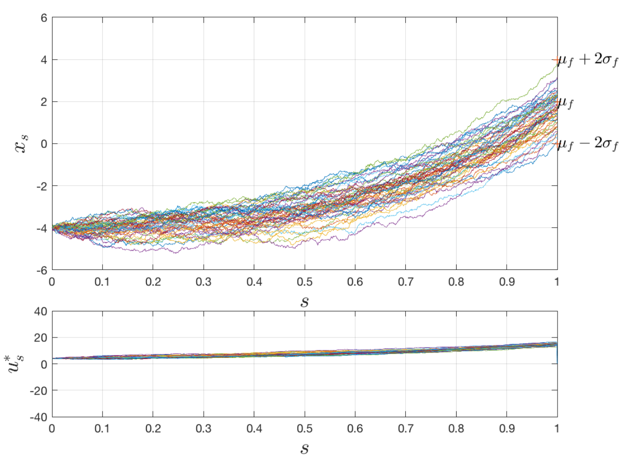

Fig. 3: The implementation of the Covariance Control methodology, assigning the normal distribution $\mathcal{N}(2, 1)$ to the terminal state.

terminal state by enforcing

$$\mathbb{E}^{[u]}_{\mathcal{F}^t}[x_1] = \begin{bmatrix} -1 \\ -1 \end{bmatrix}, \tag{48}$$

at all $t \in [0, 1]$, and consider the associated optimal control problem with the cost

$$J\big(t, x_t, [u]\big) := \mathbb{E}^{[u]}_{\mathcal{F}^t}\left[\int_t^{t_f} \frac{1}{2} u_s^2 \,\mathsf{d}s + \frac{1}{2}\left\|x_{t_f} - \mu_f\right\|^2\right]. \tag{49}$$

The implementation of the TC-SMP for 50 Sample paths are illustrated in Figure 4.

Example 4.3: Consider a similar system as in Example 4.2 but with two dimensional input and noise processes, i.e.,

$$\mathsf{d}x_s = \left(\begin{bmatrix} 0 & 1 \\ 1 & 2 \end{bmatrix}\begin{bmatrix} x_s^{(1)} \\ x_s^{(2)} \end{bmatrix} + \begin{bmatrix} 1 & 0.2 \\ 0.5 & 1 \end{bmatrix}\begin{bmatrix} u_s^{(1)} \\ u_s^{(2)} \end{bmatrix}\right)\mathsf{d}s + \begin{bmatrix} 0.2 \\ 1 \end{bmatrix}\mathsf{d}w_s, \tag{50}$$

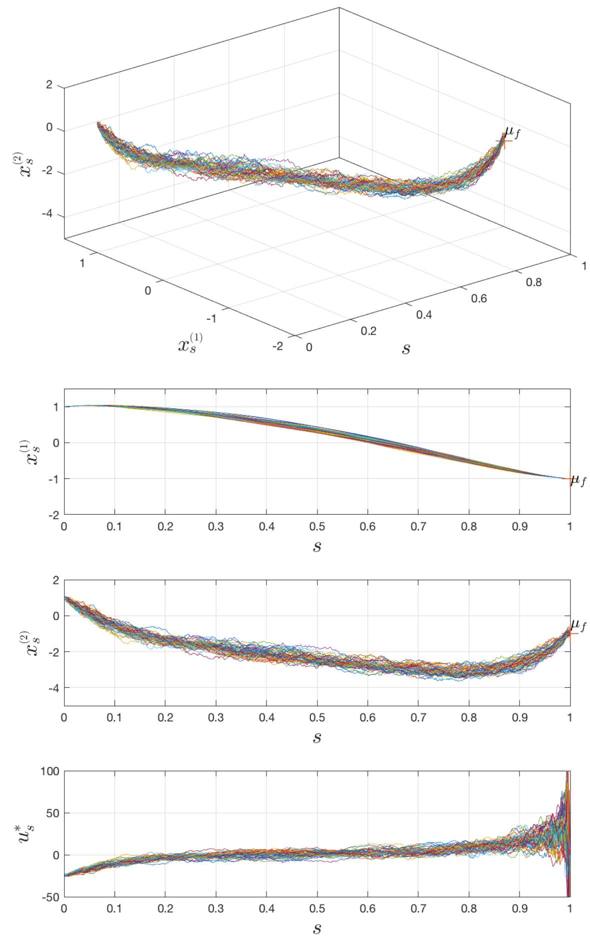

Fig. 4: The implementation of the Terminally Constrained Stochastic Minimum Principle (TC-SMP) on the system in Example 4.2.

over the time horizon $[t_0, t_f] = [0, 1]$, starting from the initial condition $x_0 = [1, 1]^{\mathsf{T}}$, and steered towards the desired terminal state by enforcing

$$\mathbb{E}^{[u]}_{\mathcal{F}^t}[x_1] = \begin{bmatrix} -1 \\ -1 \end{bmatrix}, \tag{51}$$

at all $t \in [0, 1]$, and consider the associated optimal control problem with the cost

$$J\big(t, x_t, [u]\big) := \mathbb{E}^{[u]}_{\mathcal{F}^t}\left[\int_t^{t_f} \frac{1}{2}\left\|u_s\right\|^2 \mathsf{d}s + \frac{1}{2}\left\|x_{t_f} - \mu_f\right\|^2\right]. \tag{52}$$

The resulting trajectories from the implementation of the TC-SMP on this system for 50 sample paths are illustrated in Figure 5. Again, it is shown that the TC-SMP drives all trajectory realizations to the final desired state.

V. Concluding Remarks

A new formulation of the state steering problem for stochastic systems is proposed. The formulation hinges on enforcing the terminal state constraint under the problem filtrations at every instance of time, and the associated

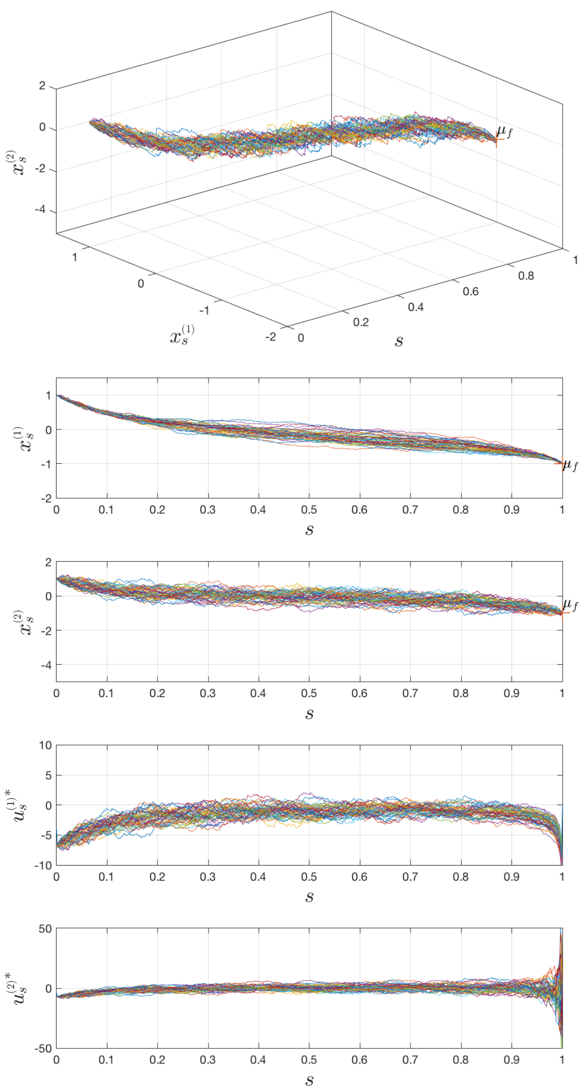

Fig. 5: The implementation of the Terminally Constrained Stochastic Minimum Principle (TC-SMP) on the system in Example 4.3.

solution methodology is based on the Terminally Constrained Stochastic Minimum Principle (TC-SMP). The solution consists of a coupled set of forward/backward stochastic differential equations. For the case of linear systems with quadratic costs, analytical expressions for the optimal inputs are provided in terms of the system's state transition matrix, its controllability Gramian and the solution of a Riccati equation. Numerical examples demonstrate that the TC-SMP performs successfully in terms of steering of the state towards the desired location at the expense of increased control effort at the final time. Future work includes the development of numerical algorithms for the solution of the general nonlinear TC-SMP, and the extension of the theory to stochastic hybrid systems.

References

[1] L. Pontryagin, V. Boltyanskii, R. Gamkrelidze, and E. Mishchenko, *The Mathematical Theory of Optimal Processes*. Wiley Interscience, New York, 1962, vol. 4.

[2] A. Hotz and R. E. Skelton, "Covariance Control Theory," *International Journal of Control*, vol. 46, no. 1, pp. 13–32, 1987.

[3] J.-H. Xu and R. Skelton, "An Improved Covariance Assignment Theory for Discrete Systems," *IEEE transactions on Automatic Control*, vol. 37, no. 10, pp. 1588–1591, 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] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal Steering of a Linear Stochastic System to a Final Probability Distribution, Part I," *IEEE Transactions on Automatic 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 Transactions on Automatic Control*, vol. 61, no. 5, pp. 1170–1180, 2015.

[8] A. Beghi, "On the Relative Entropy of Discrete-Time Markov Processes with Given End-Point Densities," *IEEE Transactions on Information Theory*, vol. 42, no. 5, pp. 1529–1535, 1996.

[9] E. Bakolas, "Optimal covariance control for discrete-time stochastic linear systems subject to constraints," in *Proceedings of the 55th IEEE Conference on Decision and Control*, 2016, pp. 1153–1158.

[10] ——, "Finite-Horizon Covariance Control for Discrete-Time Stochastic Linear Systems subject to Input Constraints," *Automatica*, vol. 91, pp. 61–68, 2018.

[11] M. Goldshtein and P. Tsiotras, "Finite-Horizon Covariance Control of Linear Time-Varying Systems," in *Proceedings of the 56th IEEE Conference on Decision and Control*, 2017, pp. 3606–3611.

[12] 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.

[13] K. Okamoto and P. Tsiotras, "Optimal Stochastic Vehicle Path Planning using Covariance Steering," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2276–2281, 2019.

[14] A. Halder and E. D. Wendel, "Finite Horizon Linear Quadratic Gaussian Density Regulator with Wasserstein Terminal Cost," in *Proceedings of the American Control Conference*, 2016, pp. 7249–7254.

[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," in *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 3484–3490.

[17] A. Mesbah, S. Streif, R. Findeisen, and R. D. Braatz, "Stochastic Nonlinear Model Predictive Control with Probabilistic Constraints," in *Proceedings of the American Control Conference*, 2014, pp. 2413–2419.

[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*. IEEE, 2014, pp. 5574–5579.

[19] M. A. Sehr and R. R. Bitmead, "Particle Model Predictive Control: Tractable Stochastic Nonlinear Output-Feedback MPC," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 15 361–15 366, 2017.

[20] K. Okamoto and P. Tsiotras, "Stochastic model predictive control for constrained linear systems using optimal covariance steering," 2019, http://arxiv.org/abs/1905.13296.

[21] S. Peng, "A General Stochastic Maximum Principle for Optimal Control Problems," *SIAM Journal on Control and Optimization*, vol. 28, no. 4, pp. 966–979, Jun. 1990.

[22] C.-T. Chen, "Linear System Theory and Design," *Rinhart and Winston, Orlando*, 1984.

[23] J. Yong and X. Y. Zhou, *Stochastic Controls: Hamiltonian Systems and HJB Equations*. Springer-Verlag New York, 1999.