

Partially Observed Steering the State of Linear Stochastic Systems

Ali Pakniyat and Panagiotis Tsiotras

Abstract—The presence of terminal state constraints in terms of expectations is studied for steering the state of partially observed linear stochastic systems. Three scenarios for the observation process are considered, namely, (i) continuous-time exact observations of the state, (ii) discrete-time exact observations of the state, and (iii) discrete-time exact observations of the state accompanied by continuous-time noisy observations of the state. Closed form expressions are presented for the optimal inputs enforcing the terminal state constraint under these information structures, which are expressed in terms of controllability Gramians and solutions of Riccati and Lyapunov equations. Numerical examples are provided to illustrate the results.

I. INTRODUCTION

In several engineering applications, it is desired to bring a system to a specific terminal configuration. A classical example is bringing an inverted pendulum to the upright position with zero velocity. 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 control inputs for the steering of the state to the desired terminal value.

In the presence of a stochastic diffusion, these state steering problems are more challenging to solve and have been the subject of only a limited number of studies [2]–[14]. For stochastic systems with linear dynamics and quadratic cost, and in the absence of any additional state constraints, the associated probability distributions are Gaussian, and the dynamics for the mean state process and the covariance state process are decoupled. This has led to the formulation of such problems as the association of a desired Gaussian distribution to the total probability distribution of the state in both infinite time horizon [2]–[5], and finite time horizon [6]–[14] settings. The accommodation of input constraints is considered in [10], and convex relaxations for linear systems subject to chance constraints (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. The accommodation of the information obtained from exact observations of the state are studied in [17]–[20] within the context of stochastic model predictive control (MPC). However, these MPC-based methodologies require recomputation of the entire procedure, and the effect of incomplete and noisy observations of the state cannot be easily accommodated into these studies.

In past work of the authors [21], an alternative approach based on the Stochastic Minimum Principle (SMP) [22] is presented for both linear and nonlinear stochastic systems under full observation of the state. The proposed methodology in [21] provides a natural accommodation of filtration, permitting input policies to be adjusted to the obtained information in order to enforce the desired terminal state constraint. In this paper, we expand upon those results by considering incomplete and noisy observations of the state within the context of linear stochastic systems. In particular, by considering three scenarios for the observation process, we establish closed form expressions for the optimal inputs for the enforcement of the terminal state constraint under the information structure corresponding to these scenarios.

The organization of the paper is as follows. In Section II a class of continuous time linear stochastic systems and the associated filtration-adapted inputs are presented. The optimal state steering problem is defined as the determination of an input to satisfy the filtration-adapted constraint on the conditional expectation of the terminal state, while minimizing a performance measure under all filtrations. The underlying assumptions on the system along with the necessary notation and fundamental notions appearing in the results are presented in Section III. The considered scenarios for the observation process are presented in Section IV. Namely, Section IV-A presents the case of continuous in time exact observations of the state, Section IV-B considers the case of discrete-time exact observations of the state, and Section IV-C studies the case with discrete-time exact observations accompanied by continuous-time noisy observations of the state, and where the terminal state constraint is enforced under the exact (discrete) subset of the information. To illustrate the results, numerical examples are provided in Section V and the performance of the controllers under each information structure are compared. Discussions about the results and further concluding remarks are provided in Section VI.

II. PROBLEM STATEMENT

Consider a linear stochastic system governed by the controlled Itô differential equation

$$dx_s = (A_s x_s + B_s u_s) ds + G_s dw_s, \quad (1)$$

A. Pakniyat is an Assistant Professor in the department of Mechanical Engineering at the University of Alabama, Tuscaloosa, AL, USA, and P. Tsiotras is a Professor in the School of Aerospace Engineering at Georgia Institute of Technology, Atlanta, GA, USA. Email: apakniyat@ua.edu, tsiotras@gatech.edu

where, $x_s \in \mathbb{R}^{n_x}$, $u_s \in \mathbb{R}^{n_u}$ are, respectively, the values of the state and the input processes at time $s \in [t_0, t_f]$, and $w_s \in \mathbb{R}^{n_w}$ is a standard n_w -dimensional Wiener process, and where $A_s \in \mathbb{R}^{n_x \times n_x}$, $B_s \in \mathbb{R}^{n_x \times n_u}$ and $G_s \in \mathbb{R}^{n_x \times n_w}$ are deterministic matrix-valued bounded functions of time.

Let $\mathcal{F}_t^o$ denote the information available from observations up until time $t \in [t_0, t_f]$, which will be specified later, and let an $\mathcal{F}_s^o$ -adapted input be denoted by $\llbracket u \rrbracket \equiv \llbracket u \rrbracket_{t_0}^{t_f} := \{u_s, s \in [t_0, t_f]; u_s \in \mathbb{R}^{n_u}, u_s : \mathcal{F}_s^o\text{-adapted}\}$. We introduce the notation

$$\mathbb{E}_{\mathcal{F}_t^o}^{\llbracket u \rrbracket}[\cdot] := \mathbb{E}[\cdot | \mathcal{F}_t^o; \llbracket u \rrbracket] \quad (2)$$

for the conditional expectation of any random variable given the available information at time t and the input process $\llbracket u \rrbracket$.

The objective of this paper is the characterization of inputs for the enforcement of the terminal state constraint

$$\mathbb{E}_{\mathcal{F}_t^{st}}^{\llbracket u \rrbracket}[x_{t_f}] = \mu_f, \quad (3)$$

for all $t \in [t_0, t_f]$, where $\mathcal{F}_t^{st} \subseteq \mathcal{F}_t^o$, to be specified in the next section, is a subset of the available filtration over which the state steering constraint (3) is enforced.

The cost-to-go associated with an input $\llbracket u \rrbracket$ is defined by

$$J(t, x_t; \llbracket u \rrbracket) := \frac{1}{2} \mathbb{E}_{\mathcal{F}_t^o}^{\llbracket u \rrbracket} \left[\int_t^{t_f} u_s^T R_s u_s ds + (x_{t_f} - \mu_f)^T H_f (x_{t_f} - \mu_f) \right], \quad (4)$$

where $\mu_f \in \mathbb{R}^{n_x}$ is a fixed desired terminal state, and $R_s = R_s^T \in \mathbb{R}^{n_u \times n_u}$, $R_s > 0$ is a deterministic matrix-valued bounded function of time and $H_f = H_f^T \in \mathbb{R}^{n_x \times n_x}$, $H_f \geq 0$.

III. DEFINITIONS AND UNDERLYING ASSUMPTIONS

Let $\Phi(s, t) \in \mathbb{R}^{n_x \times n_x}$ denote the state transition matrix from t to s for the system (1), which is the solution of

$$\dot{\Phi} \equiv \frac{\partial \Phi(s, t)}{\partial s} = A_s \Phi, \quad \Phi(t, t) = I. \quad (5)$$

For each $t, \tau \in [t_0, t_f]$ such that $t_0 \leq \tau < t \leq t_f$, let us define the Gramian as

$$\mathcal{G}(\tau, t) := \int_\tau^t \Phi(t_f, s) B_s R_s^{-1} B_s^T \Phi(t_f, s) ds. \quad (6)$$

The above definition of the Gramian is a variant of the conventional controllability Gramian (see e.g., [23, Theorem 6.1]) which is related to (6) by taking $R_s \equiv I_{n_u \times n_u}$ and selecting $\tau = t_0$ and $t = t_f$.

We also define $\Pi(s; t_f)$ as the solution of the following Riccati equation

$$\begin{aligned} \dot{\Pi}_s &\equiv \frac{d}{ds} \Pi(s; t_f) = \Pi_s B_s R_s^{-1} B_s^T \Pi_s - \Pi_s A_s - A_s^T \Pi_s, \\ \Pi(t_f; t_f) &= H_f. \end{aligned} \quad (7)$$

Assumption 3.1: In this paper, we assume that

- (i) The pair (A, B) is controllable, hence, the Gramian (6) is non-singular and, therefore, it is invertible.
- (ii) The nullity of the diffusion coefficient G is zero, i.e., G_s is full-rank for all $s \in [t_0, t_f]$.
- (iii) The system is noise controllable, i.e., for all $s \in [t_0, t_f]$, $\text{Im}(G_s) \subset \text{Im}(B_s)$, that is, $\forall w \in \mathbb{R}^{n_w}, \exists u \in \mathbb{R}^{n_u}$ s.t. $B_s u = G_s w$.

IV. OBSERVATION-BASED STEERING OF THE STATE

While it is possible to consider various cases for the observation filtration $\mathcal{F}_t^o$ and the filtration $\mathcal{F}_t^{st}$ for the steering constraint enforcement, we restrict our attention to three cases. Namely, continuously in time exact measurements of the state; sampled data (discrete in time) exact measurements of the state; and sampled discrete measurements of the state combined with continuous noisy observations.

A. CFO: Continuous-Time Full Observation of the State

In this case, the exact value of the state x_t is measured at each time $t \in [t_0, t_f]$ and it is made available to the controller. In other words, the information available at time t is $\mathcal{F}_t^o = \mathcal{F}_t^{x,c}$, where

$$\mathcal{F}_t^{x,c} := \sigma\{x_s; s \in [t_0, t]\}. \quad (8)$$

In this case, the steering constraint (3) is enforced under the same filtration.

Theorem 4.1: For the enforcement of (3) under the filtration $\mathcal{F}_t^{st} = \mathcal{F}_t^{x,c}$, the minimum value of the cost (4) under the filtration $\mathcal{F}_t^o = \mathcal{F}_t^{x,c}$ is achieved by

$$\begin{aligned} u_s^* &= -R_s^{-1} B_s^T \Phi(t_f, s)^T [\mathcal{G}(t, t_f)]^{-1} (\Phi(t_f, t) x_t - \mu_f) \\ &\quad - R_s^{-1} B_s^T \Pi(s; t_f) \left(x_s - \Phi(s, t) x_t \right. \\ &\quad \left. + \mathcal{G}(t, s) \Phi(t_f, s)^T [\mathcal{G}(t, t_f)]^{-1} (\Phi(t_f, t) x_t - \mu_f) \right), \end{aligned} \quad (9)$$

where the state transition matrix Φ is defined in (5), the Gramian $\mathcal{G}$ is defined in (6), and the Riccati matrix Π is defined in (7). $\square$

Proof: Since G is bounded and under the Assumption 3.1(ii), its nullity is zero, the sigma algebras $\mathcal{F}_t^{x,c}$ and $\mathcal{F}_t^w := \sigma\{w_s; s \in [t_0, t]\}$ are equivalent [24, Lemma 1.1] and the set of $\mathcal{F}_t^{x,c}$ -adapted inputs are dense in the set of $\mathcal{F}_t^w$ -adapted inputs [24, Lemma 1.2]. Hence, the requirements (3) and (4) are equivalent to the conditioning on $\mathcal{F}_t^w$ since $x_{t_0} = x_0$ is deterministic and, therefore, [21, Theorem 3.1] can be invoked to obtain (9). $\blacksquare$

B. DFO: Discrete-Time Full Observation of the State

In this case, instead of perpetual measurements of the state as in Section IV-A, the observations are made only over a finite set of strictly increasing sampling times $\{\tau_i\}_{i=0}^N$ within $[t_0, t_f]$, such that $\tau_0 := t_0 < \tau_1 < \dots < \tau_N < \tau_{N+1} := t_f$. Thus $\mathcal{F}_t^o = \mathcal{F}_t^{x,d}$, where

$$\mathcal{F}_t^{x,d} := \sigma\{x_{\tau_i}; \tau_i \in \{\tau_j\}_{\tau_j \leq t}\}. \quad (10)$$

This case is very common in engineering applications where a continuous time physical system is observed by digital measurement devices. Due to its nature, this class of observations is sometimes referred to as sampled-data observations. We consider the case where the steering requirement (3) is enforced based on the same filtration.

Theorem 4.2: For the enforcement of (3) under the filtration $\mathcal{F}_t^{st} = \mathcal{F}_t^{x,d}$, the minimum value of the cost (4) under filtration $\mathcal{F}_t^o = \mathcal{F}_t^{x,d}$ is achieved by the input policy

$$u_s^* = -R_s^{-1} B_s^T \Phi(t_f, s)^T [\mathcal{G}(\tau, t_f)]^{-1} (\Phi(t_f, \tau) x_\tau - \mu_f), \quad (11)$$

where $\tau := \max\{\tau_i, \text{s.t. } \tau_i \leq t\}$ is the time of the most recent observation of the state and x_τ is the corresponding observed value.

Proof: Define the following variables

$$\bar{x}_s := \mathbb{E}_{\mathcal{F}_t^{x,d}}^{[u^*]}[x_s], \quad \bar{u}_s := \mathbb{E}_{\mathcal{F}_t^{x,d}}^{[u^*]}[u_s], \quad (12)$$

$$\tilde{x}_s := x_s - \bar{x}_s, \quad \tilde{u}_s := u_s - \bar{u}_s, \quad (13)$$

and recall that under $\mathcal{F}_t^{st} = \mathcal{F}_t^{x,d}$, the terminal constraint (3) imposes a constraint on the $\bar{x}$ process, i.e.,

$$\bar{x}_{t_f} = \mu_f, \quad (14)$$

whereas $\tilde{x}_{t_f}$ is free (unconstrained). Noting the zero expectation terms $\mathbb{E}[\tilde{x}] = 0$, $\mathbb{E}[\tilde{x}^\top \tilde{x}] = 0$, etc., the cost (4) can be decomposed as

$$\begin{aligned} J(t, x_t; [u]) &= \frac{1}{2} \mathbb{E}_{\mathcal{F}_t^o}^{[u]} \left[\int_t^{t_f} (\bar{u}_s + \tilde{u}_s)^\top R_s (\bar{u}_s + \tilde{u}_s) ds \right. \\ &\quad \left. + (\bar{x}_{t_f} + \tilde{x}_{t_f} - \mu_f)^\top H_f (\bar{x}_{t_f} + \tilde{x}_{t_f} - \mu_f) \right] \\ &\stackrel{(14)}{=} \frac{1}{2} \int_t^{t_f} \bar{u}_s^\top R_s \bar{u}_s ds + \frac{1}{2} \mathbb{E}_{\mathcal{F}_t^o}^{[u]} \left[\int_t^{t_f} \tilde{u}_s^\top R_s \tilde{u}_s ds + \tilde{x}_{t_f}^\top H_f \tilde{x}_{t_f} \right] \\ &\equiv J(t, \bar{x}_t; [\bar{u}]) + J(t, \tilde{x}_t; [\tilde{u}]). \end{aligned} \quad (15)$$

From the definitions of $\bar{x}_s$ and $\tilde{x}_s$ and the dynamics (1), it follows that

$$d\bar{x}_s = (A_s \bar{x}_s + B_s \bar{u}_s) ds, \quad (16)$$

$$d\tilde{x}_s = (A_s \tilde{x}_s + B_s \tilde{u}_s) ds + G_s dw_s, \quad (17)$$

hold over the interval $[t, t_f]$, subject to the initial conditions

$$\bar{x}_t = \mathbb{E}_{\mathcal{F}_t^{x,d}}^{[u^*]}[x_t] = \hat{x}_t, \quad \tilde{x}_t = \check{x}_t, \quad (18)$$

where the process $\hat{x}$ is defined on the interval $[\tau, t]$ as the solution of

$$d\hat{x}_s = (A_s \hat{x}_s + B_s u_s) ds, \quad \hat{x}_\tau = x_\tau, \quad (19)$$

and the (unobserved) process $\check{x}$ is the solution of

$$d\check{x}_s = A_s \check{x}_s ds + G_s dw_s, \quad \check{x}_\tau = 0. \quad (20)$$

In other words, the minimization of the cost (4) under the dynamics (1) and the terminal state constraint (3) is decomposed into

$$\inf_{[u]} J(t, x_t; [u]) = \inf_{[\bar{u}]} J(t, \bar{x}_t; [\bar{u}]) + \inf_{[\tilde{u}]} J(t, \tilde{x}_t; [\tilde{u}]), \quad (21)$$

where, on the right hand side, the first term corresponds to the fixed end point deterministic optimal control problem with the dynamics (16), the cost $J(t, \bar{x}_t; [\bar{u}]) = \frac{1}{2} \int_t^{t_f} \bar{u}_s^\top R_s \bar{u}_s ds$, the initial condition (18) and the terminal condition (14), and the second term corresponds to the free end point stochastic linear quadratic regulator problem with the dynamics (17), the cost $J(t, \tilde{x}_t; [\tilde{u}]) = \frac{1}{2} \mathbb{E}_{\mathcal{F}_t^o}^{[u]} \left[\int_t^{t_f} \tilde{u}_s^\top R_s \tilde{u}_s ds + \tilde{x}_{t_f}^\top H_f \tilde{x}_{t_f} \right]$, and the initial condition (18).

In order to obtain the solution of these problems, consider the two auxiliary problems over the interval $[\tau, t_f]$, i.e.,

$$\begin{aligned} \inf_{[\bar{u}]} \frac{1}{2} \int_\tau^{t_f} \bar{v}_s^\top R_s \bar{v}_s ds, \\ \frac{d}{ds} \bar{z}_s = A_s \bar{z}_s + B_s \bar{v}_s, \quad \bar{z}_\tau = x_\tau, \quad \bar{z}_{t_f} = \mu_f, \end{aligned} \quad (P1)$$

and

$$\begin{aligned} \inf_{[\tilde{u}]} \frac{1}{2} \mathbb{E}_{\mathcal{F}_\tau^o}^{[\tilde{u}]} \left[\int_\tau^{t_f} \tilde{v}_s^\top R_s \tilde{v}_s ds + \tilde{z}_{t_f}^\top H_f \tilde{z}_{t_f} \right], \\ d\tilde{z}_s = (A_s \tilde{z}_s + B_s \tilde{v}_s) ds + G_s dw_s, \quad \tilde{z}_\tau = 0. \end{aligned} \quad (P2)$$

The optimal input for the auxiliary problem (P1) has the form

$$\bar{v}_s^* = -R_s^{-1} B_s^\top \Phi(t_f, s)^\top [\mathcal{G}(\tau, t_f)]^{-1} (\Phi(t_f, \tau) x_\tau - \mu_f), \quad (22)$$

for $s \in [\tau, t_f]$ and the optimal input for the auxiliary problem (P2) over $[\tau, t_f]$ is

$$\tilde{v}_s^* = 0, \quad (23)$$

since any nontrivial input process adapted to $\mathcal{F}_t^o = \mathcal{F}_t^{x,d}$ yields a higher cost.

It remains to show that the optimal solution for the $x = \bar{x} + \tilde{x}$ process coincides with the optimal solution for the $z = \bar{z} + \tilde{z}$ process, i.e. $u^* \stackrel{a.s.}{=} v^*$. This is simply shown by the method of contradiction by, first, arguing that under the Assumption 3.1, v^* uniquely exists and second, by arguing that if $\bar{u}^* \neq \bar{v}^*$ over an interval $[t', t''] \subset [t, t_f]$ then it violates the uniqueness of v^* , i.e. u^* must also be an optimal input for the process z because $\bar{z}_{t'}^* = \bar{x}_{t'}$ due to that fact that $v_s^* \equiv \bar{v}_s^*$ (and $\tilde{v}_s^* = 0$) for $s \in [\tau, t']$ are pure functions of time (notice also that for $t' = t$, we obtain $\bar{z}_t^* = \bar{x}_t \equiv \hat{x}_t$).

Thus, $u^* \stackrel{a.s.}{=} v^*$ must hold, which together with $u_s^* = \bar{u}_s^* + \tilde{u}_s^*$, the expression (11) is obtained. ■

C. DFO-CNO: Discrete-Time Full Observation of the State together with Continuous-Time Noisy Observations

For this case, full observations of the state are taken over a discrete set of times, as in Section IV-B, but in between these times, we have access to a noisy observation of the state, throughout the process

$$dy_s = C_s x_s ds + D_s dv_s, \quad (24)$$

where the observation noise $v_s \in \mathbb{R}^{n_v}$, $s \in [t_0, t_f]$ is a standard n_v -dimensional Wiener process, independent of the process noise w_s . Hence, $\mathcal{F}_t^o = \sigma\{\mathcal{F}_t^{x,d} \cup \mathcal{F}_t^y\}$ where $\mathcal{F}_t^{x,d}$ is defined in (10), and $\mathcal{F}_t^y \equiv \mathcal{F}_t^{dy}$ are defined as

$$\mathcal{F}_t^y := \sigma\{y_s; s \in [t_0, t]\} \equiv \mathcal{F}_t^{dy} := \sigma\{dy_s; s \in [t_0, t]\}. \quad (25)$$

Assumption 4.3: We assume that the nullity of the matrices C and D are zero, i.e., C_s and D_s are full-rank for all $s \in [t_0, t_f]$.

This case often arises in applications where a digital measurement device provides accurate information about the state at countable time instances, and a physical device with

indirect measurements provides noisy observations continuously in time. We consider the case where the steering requirement (3) is enforced based on the un-noisy (exact) subset of the observation information.

Theorem 4.4: For the enforcement of (3) under the filtration $\mathcal{F}_t^{st} = \mathcal{F}_t^{x,d}$, the minimum value of the cost (4) under filtration $\mathcal{F}_t^o = \sigma\{\mathcal{F}_t^{x,d} \cup \mathcal{F}_t^y\}$ is achieved by

$$\begin{aligned} u_s^* = & -R_s^{-1}B_s^\top \Phi(t_f, s)^\top [\mathcal{G}(\tau, t_f)]^{-1} (\Phi(t_f, \tau)x_\tau - \mu_f) \\ & - R_s^{-1}B_s^\top \Pi(s; t_f) \left(\hat{x}_s - \Phi(s, \tau)x_\tau \right. \\ & \left. + \mathcal{G}(\tau, s)\Phi(t_f, s)^\top [\mathcal{G}(\tau, t_f)]^{-1} (\Phi(t_f, \tau)x_\tau - \mu_f) \right), \end{aligned} \quad (26)$$

where x_τ is the most recent observation of the state, and $\hat{x}_s := \mathbb{E}_{\mathcal{F}_t^o}^u[x_s]$ is obtained from the Kalman-Bucy filter

$$d\hat{x}_s = (A_s\hat{x}_s + B_s u_s)ds + \Sigma_s C_s^\top (D_s D_s^\top)^{-1} (dy_s - C_s \hat{x}_s ds), \quad (27)$$

with $\Sigma_s := \mathbb{E}_{\mathcal{F}_t^o}^u[(x_s - \hat{x}_s)(x_s - \hat{x}_s)^\top]$ governed by the following Lyapunov equation

$$\begin{aligned} \dot{\Sigma}_s = & (A_s - B_s B_s^\top \Pi(s; t_f))\Sigma_s + \Sigma_s (A_s - B_s B_s^\top \Pi(s; t_f))^\top \\ & - \Sigma_s C_s^\top (D_s D_s^\top)^{-1} C_s \Sigma_s + G_s G_s^\top, \end{aligned} \quad (28)$$

subject to the (re)initialization condition $\Sigma_\tau = 0$.

Proof: Due to space limitations, only a sketch of the proof is provided here.

With the definition of $\hat{x}_s := \mathbb{E}_{\mathcal{F}_t^o}^u[x_s]$, it follows that $\hat{x}_s \stackrel{(12)}{=} \mathbb{E}_{\mathcal{F}_t^o}^u[\bar{x}_s + \tilde{x}_s] \stackrel{(13)}{=} \mathbb{E}_{\mathcal{F}_t^o}^u[\mathbb{E}_{\mathcal{F}_t^{x,d}}^u[x_s]] + \mathbb{E}_{\mathcal{F}_t^o}^u[\tilde{x}_s] = \bar{x}_s + \mathbb{E}_{\mathcal{F}_t^o}^u[\tilde{x}_s]$. Moreover, $\mathbb{E}_{\mathcal{F}_t^o}^u[\tilde{x}_s] = \mathbb{E}_{\sigma\{\mathcal{F}_t^{x,d} \cup \mathcal{F}_t^y\}}^u[\tilde{x}_s] = \mathbb{E}_{\mathcal{F}_t^y}^u[\tilde{x}_s]$, where the last equality is obtained using the fact that $\mathbb{E}_{\mathcal{F}_t^{x,d}}^u[\tilde{x}_s] = 0$.

Thus, with the definition of $\tilde{x}_s := x_s - \hat{x}_s \equiv \mathbb{E}_{\mathcal{F}_t^y}^u[\tilde{x}_s]$, and following similar steps as in the proof of Theorem 4.2, we obtain (P1) together with

$$\begin{aligned} \inf_{\tilde{u}} \frac{1}{2} \mathbb{E}_{\mathcal{F}_\tau^y}^u \left[\int_\tau^{t_f} \tilde{u}_s^\top R_s \tilde{u}_s ds + \tilde{x}_{t_f}^\top H_f \tilde{x}_{t_f} \right], \\ d\tilde{x}_s = (A_s \tilde{x}_s + B_s \tilde{u}_s)ds + G_s dw_s, \quad \tilde{x}_\tau = 0, \\ d\tilde{y}_s = C_s \tilde{x}_s ds + D_s dv_s, \end{aligned} \quad (P3)$$

whose solution is obtained from an LQG observer based controller with a Kalman-Bucy filtration [25] for the state estimator. ■

V. NUMERICAL ILLUSTRATIONS

Consider the system governed by

$$dx_s = \left(\begin{bmatrix} 0 & 1 \\ 1 & 2 \end{bmatrix} \begin{bmatrix} x_s^{(1)} \\ x_s^{(2)} \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u_s \right) ds + \begin{bmatrix} 0 \\ 1 \end{bmatrix} dw_s, \quad (29)$$

over the time horizon $[t_0, t_f] = [0, 1]$, with the initial condition $x_0 = [1, 1]^\top$, and consider the problem of steering its state towards the desired terminal state by enforcing

$$\mathbb{E}_{\mathcal{F}_t^{st}}^u[x_1] = \begin{bmatrix} -1 \\ -1 \end{bmatrix}, \quad (30)$$

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

$$J(t, x_t; [u]) := \mathbb{E}_{\mathcal{F}_t^o}^u \left[\int_t^{t_f} \frac{1}{2} u_s^2 ds + \frac{1}{2} \|x_{t_f} - \mu_f\|^2 \right].$$

For the case of $\mathcal{F}_t^o = \mathcal{F}_t^{st} = \mathcal{F}_t^{x,c}$, i.e., the case of continuous full state observation, the results of Theorem 4.1 are illustrated in Figure 1 for 50 samples paths, and the associated distributions to 2000 samples paths for the terminal state and the terminal input are displayed in Figure 2. It can be observed that the terminal state matches with the desired value almost surely, i.e., the terminal state distribution is a Dirac delta distribution with its mean matching the desired state and a zero covariance¹. However, the use of large input values is inevitable since the process noise pushes the system away from the satisfaction of this constraint, and the counterbalance of this effect requires larger and larger input values as the time approaches the terminal time. This can be observed in Figures 1 and 2 and it is also anticipated due to the appearance of $[\mathcal{G}(t, t_f)]^{-1}$ in the control law (9).

For the case of $\mathcal{F}_t^o = \mathcal{F}_t^{st} = \mathcal{F}_t^{x,d}$ with discrete observations every 0.05 sec, the optimal trajectories are displayed in Figure 3 and the associated terminal distributions are illustrated in Figure 4. Moreover, in order to illustrate the effect of inter-measurement time, the optimal trajectories and the associated terminal distributions for the case of discrete observations every 0.01 sec are illustrated, respectively, in Figure 5 and Figure 6. It can be observed from these examples that the enforcement of the terminal state constraint (3) under the filtration corresponding to discrete observations of the state results in a terminal state distribution with its mean matching with the desired state value, but non-zero covariance. This distribution is a function of the length of inter-observation times, and in particular, the penultimate observation preceding the terminal time. However, this is accompanied by the benefit of avoiding large input values the statistics of which is also a function of the length of inter-observation times. This can be observed in Figures 4 and 6, and is also anticipated from the control law (11) and, in particular, from the fact that $[\mathcal{G}(\tau_N, t_f)]^{-1} \equiv [\mathcal{G}(\tau_N, \tau_{N+1})]^{-1}$ possesses upper bounds as function of $\tau_{N+1} - \tau_N$.

For the case of $\mathcal{F}_t^o = \sigma\{\mathcal{F}_t^{x,d} \cup \mathcal{F}_t^y\}$ and $\mathcal{F}_t^{st} = \mathcal{F}_t^{x,d}$, we consider the observation process

$$\begin{aligned} dy_s \equiv d \begin{bmatrix} y_s^{(1)} \\ y_s^{(2)} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_s^{(1)} \\ x_s^{(2)} \end{bmatrix} ds \\ + \begin{bmatrix} 0.005 & 0 \\ 0 & 0.005 \end{bmatrix} d \begin{bmatrix} v_s^{(1)} \\ v_s^{(2)} \end{bmatrix}, \end{aligned} \quad (31)$$

to accompany discrete observations every 0.01 sec. The optimal trajectories and the associated terminal distributions are displayed, respectively, in Figure 7 and Figure 8. It can be observed that the presence of continuous noisy observations

¹The reason that the experimental distribution displayed in Figure 2 is not exactly a delta distribution is due to the facts that (i) the time-discretization in the numerical solution and (ii) the space discretization in the derivation of the experimental distribution. However, in mathematical terms, the equality is achieved almost surely, i.e. the probability of inequality is zero.




Fig. 1: State evolution and the associated optimal inputs with continuous-time full state observation.




Fig. 5: State evolution and the associated optimal inputs with discrete full state observations every 0.01 sec.




Fig. 2: Distribution of the terminal state and input for 2000 sample paths with continuous-time full state observation.




Fig. 6: Distribution of the terminal state and input for 2000 samples with discrete full state observations every 0.01 sec.




Fig. 3: State evolution and the associated optimal inputs with discrete full state observations every 0.05 sec.




Fig. 7: State evolution and the associated optimal inputs with discrete full state observations every 0.01 sec and noisy observations in between.




Fig. 4: Distribution of the terminal state and input for 2000 samples with discrete full state observations every 0.05 sec.




Fig. 8: Distribution of the terminal state and input for 2000 samples with discrete full state observations every 0.01 sec and noisy observations in between.

in addition to discrete exact observations yields similar properties for the terminal state and terminal inputs whenever the terminal state constraint (3) under the filtration corresponding to discrete observations of the state. This can be deduced by comparing Figures 7 and 8 with Figures 5 and 6. The significant difference, however, lies in the fact the cost for this case is necessarily lower than the cost for the case with only discrete observations.

VI. CONCLUDING REMARKS

The enforcement of terminal state constraints under different information structures is studied and closed form expressions of the optimal input for the steering of the state towards the desired state are established. Since the controllability Gramian and solutions of Riccati and Lyapunov equations are fundamental characteristics of the system which can be computed and stored in advance of implementation, the required online computations are minimal and are limited to substitutions of the state observations.

In this study, the considered terminal state constraints are defined over the first moment (expectation) of the terminal state and, therefore, the second moment (covariance) of the terminal state becomes a byproduct of the assumed information structure. To be specific, the enforcement of the terminal state constraints under the filtration associated with continuous observations of the state yield a zero covariance for the terminal state in the expense of the appearance of large input values. In contrast, the enforcement under the filtration associated with discrete observations of the state yield a non-zero but finite fixed covariance which depends on inter-observation times. The presence of continuous noisy observations of the state does not significantly change this covariance, and only serves in cost minimization. A major practical benefit of enforcement under discrete observations, in addition to requiring fewer observations, is that it yields a bound over the input values required to enforce the terminal state constraint.

Future work includes the study of other scenarios for the state observation and the constraint enforcement; in particular, the case where the discrete observations are themselves incomplete and noisy. Further avenues of research also include the extension of the results for nonlinear and hybrid stochastic 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] A. Pakniyat and P. Tsiotras, "Steering the State of Linear Stochastic Systems: A Constrained Minimum Principle Formulation," in *Proceedings of the IEEE American Control Conference, New Orleans, USA*, 2021, pp. 1296–1301.
- [22] S. Peng, "A General Stochastic Maximum Principle for Optimal Control Problems," *SIAM Journal on Control and Optimization*, vol. 28, no. 4, pp. 966–979, June 1990.
- [23] C.-T. Chen, "Linear System Theory and Design," *Rinhardt and Winston, Orlando*, 1984.
- [24] A. Bensoussan, *Lectures on Stochastic Control*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1982, pp. 1–62.
- [25] R. S. Bucy and P. D. Joseph, *Filtering for Stochastic Processes with Applications to Guidance*. American Mathematical Soc., 2005, vol. 326.