Multi-Model Covariance Steering for Continuous-Time Stochastic Linear Systems

Efstathios Bakolas

Abstract— **In this paper, we consider the finite-horizon covariance steering problem for continuous-time stochastic linear systems subject to white noise in the special case in which the exact state space realization of the (uncertain) system is unknown but belongs necessarily to a finite collection of known models. We refer to the latter problem as the multi-model covariance steering (MMCS) problem. Because a common controller that would steer the state covariance of each possible system realization to the same positive definite matrix may likely not exist in general, we consider a different problem formulation that is based on the a posteriori interpretation of the standard covariance steering problem as a Linear Quadratic Gaussian (LQG) control problem with a special terminal cost. This interpretation allows us to formulate the multi-model covariance steering problem as a minimax (robust) stochastic optimal control problem that can be addressed by means of the Robust Maximum Principle. In particular, we show that the feedback control law that solves the MMCS problem can be fully characterized by (1) the solution to an LQG-type problem for a system with an augmented state space and (2) a tractable, yet non-convex, optimization problem. Simulations results that illustrate some key ideas of this paper are also included.**

I. Introduction

The covariance steering (CS) problem corresponds to a non-trivial extension of the Linear Quadratic Gaussian (LQG) control problem in which the terminal state covariance of a stochastic linear system is prescribed and equal to a known positive definite matrix. In this work, we consider the finite-horizon CS problem for continuous-time uncertain linear systems whose realization is not known a priori but belongs necessarily to a finite collection of known state space models (*multi-model* representation of uncertain systems).

Covariance Steering (CS), or covariance control, problems for both continuous-time and discrete-time linear systems were first studied for the infinite horizon case [1], [2]. Finite-horizon problems for the continuous-time case are presented in [3]–[5] whereas the discrete-time case is considered in many recent papers that study different extensions / variations of the CS problem, including semi-definite programming approaches [6]–[9], data-driven approaches [10], [11], problems with parametric uncertainty [12], maximum entropy based formulations [13], and distributed approaches [14], [15], to name but a few.

The CS problem for the continuous-time case can be reduced to a system of two coupled differential Riccati / Lyapunov equations which admits a closed-form solution provided that the input and noise channels of the system coincide [3]. An interesting interpretation of the result in [3] is that the solution to the CS problem can be interpreted *a posteriori* as the solution to an LQG problem under a certain technical assumption. The important nuance here is that the terminal cost of the LQG problem which is equivalent to the CS problem cannot be defined without first solving the CS problem. Although the latter interpretation does not really offer an alternative or easier way to solve the CS problem, it can be used to find a systematic way to solve the MMCS problem, as shown in this work.

Despite the plethora of variations of the CS problem studied in the relevant literature, robust versions of covariance steering problems that explicitly account for model uncertainty have not received a lot of attention. In this paper, we study CS problems in the case of model uncertainty that can be described by a finite collection of a priori known state space models. The first step in our approach is to place the MMCS problem under the umbrella of minimax stochastic control problems and subsequently address it via the Robust Maximum Principle (RMP) [16]. Instead of seeking a single (common) optimal control law that will simultaneously steer the state covariance of each possible realization of the uncertain system to the same positive definite matrix (such a control law may not even exist in many cases), we propose a soft-constrained problem formulation. In particular, based on the interpretation of the CS problem as a specialized LQG-type problem, we replace the hard constraint on the terminal state covariance with a weighted sum of quadratic terminal costs with weights that are determined by the solution of a single tractable, but non-convex, optimization problem.

The rest of the paper is organized as follows. The multi-model covariance steering (MMCS) problem is formulated in Section II. The analysis and solution of the MMCS problem are presented in Section III. Numerical simulations are presented in Section IV. Finally, Section V concludes the paper with a summary of remarks.

E. Bakolas is an Associate Professor 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

II. Problem Setup

A. Notation

We denote by $\mathbb{S}_{+}^{n}$ and $\mathbb{S}_{++}^{n}$ the cones of (symmetric) positive semi-definite and positive definite matrices, respectively. Given a matrix $M \in \mathbb{S}_{+}^{n}$, we denote by $\sqrt{M}$ its (unique) square root matrix in $\mathbb{S}_{+}^{n}$. We denote by $\mathbf{1}$ the vector of ones and by $\mathbf{\Delta}_n$ the $n-1$ dimensional simplex, where $\mathbf{\Delta}_n := \left\{\lambda \in \mathbb{R}^n : \mathbf{1}^{\mathrm{T}}\lambda = 1 \text{ and } \lambda \geq\geq 0\right\}$. We denote by $\otimes$ and $\mathrm{tr}(\cdot)$ the Kronecker product operator and trace operator, respectively. The identity and zero matrices are denoted by I and $\mathbf{0}$, respectively. Given a matrix $N \in \mathbb{R}^{m\times n}$, we denote by $\|N\|_F$ its Frobenius norm, where $\|N\|_F := \sqrt{\mathrm{tr}(NN^{\mathrm{T}})} = \sqrt{\mathrm{tr}(N^{\mathrm{T}}N)}$. We denote by $\mathrm{bdiag}(M_1, \ldots, M_k)$ the block diagonal matrix with diagonal blocks $M_1, \ldots, M_k$. We write $[A; B]$ to denote the vertical concatenation of matrices A and B. The set of continuous mappings from $\mathcal{D}$ to $\mathcal{R}$ is denoted as $\mathcal{C}^0(\mathcal{D}; \mathcal{R})$. Finally, $\mathbb{E}\{\cdot\}$ denotes the expectation operator (or functional).

B. Problem Formulation

We consider a stochastic linear system described by a one-parameter family of stochastic differential equations (SDEs):

$$\mathrm{d}x^{\gamma}(t) = (A^{\gamma}(t)x^{\gamma}(t) + B^{\gamma}(t)u(t))\mathrm{d}t + B^{\gamma}(t)\mathrm{d}w(t), \quad (1)$$

for $t \in [0, T]$ ($T > 0$ is a given time horizon), where $x^{\gamma}(t) \in \mathbb{R}^n$ is the state, $u(t) \in \mathbb{R}^m$ is the input and $\{w(t) : \ t \in [0, T]\}$ is an $\mathbb{R}^m$-valued standard Wiener process which is defined in a filtered probability space $(\Omega, \mathfrak{F}, \mathfrak{F}_{t\in[0,T]}, \mathbb{P})$ where $\mathfrak{F}_{t\in[0,T]}$ is the natural filtration (ordered sequence of sub-σ-algebras) generated by the noise process $\{w(t) : \ t \in [0, T]\}$ and augmented with $\mathbb{P}$-null subsets of $\mathfrak{F}$ [17]. Informally, the noise process $\{w(t) : \ t \in [0, T]\}$ is a Wiener process with zero *incremental* mean and *incremental* covariance equal to $I\mathrm{d}t$ (that is, $\mathbb{E}\{\mathrm{d}w(t)\} = 0$ and $\mathbb{E}\{\mathrm{d}w(t)\mathrm{d}w(t)^{\mathrm{T}}\} = I\mathrm{d}t$). The initial state $x^{\gamma}(0)$ is a zero mean Gaussian random vector with covariance matrix $\Sigma_0 \in \mathbb{S}^n_{++}$ that is independent of the noise process $\{w(t) : \ t \in [0, T]\}$. In addition, $A^{\gamma}(\cdot) \in \mathcal{C}^0([0, T]; \mathbb{R}^{n\times n})$ and $B^{\gamma}(\cdot) \in \mathcal{C}^0([0, T]; \mathbb{R}^{n\times m})$, for all $\gamma \in \Gamma$, where we assume that the parameter set Γ is a finite set and denote by $|\Gamma|$ its cardinality. The corresponding finite collection $\mathcal{M}$ of all possible state space realizations of the uncertain system given in (1), $\mathcal{M} := \{(A^{\gamma}(t), B^{\gamma}(t)) : \ \gamma \in \Gamma\}$, is known a priori, although its exact (or true) realization with corresponding pair $(A^{\gamma^\star}(t), B^{\gamma^\star}(t))$ and parameter $\gamma^\star \in \Gamma$ are unknown.

Note that, in general, γ can be a vector that determines the values of $A^{\gamma}(t)$ and $B^{\gamma}(t)$ for $t \geq 0$. For instance, γ can be an n-dimensional vector and $A^{\gamma}(t) = \mathrm{diag}(\gamma)$, $\forall t \in [0, T]$ (structured model uncertainty). By contrast, in the special case in which γ is taken to be an index (positive integer) with corresponding index set Γ, then the model uncertainty has no specific structure. We will say that the collection $\mathcal{M}$ determines a *multi-model* representation of the uncertain system. In our problem setup, the unknown parameter $\gamma \in \Gamma$ is constant for all $t \in [0, T]$ and therefore, switching between different system models cannot take place. In addition, the input and the noise are independent of the parameter γ whereas the state evolution depends on the particular state space realization in $\mathcal{M}$ and thus the parameter γ.

Given $\gamma \in \Gamma$, the standard Covariance Steering (CS) problem seeks the minimum effort control input that will generate a state process for the system described by the SDE (1) that satisfies the following boundary conditions:

$$\mathbb{E}\{x^{\gamma}(0)\} = 0, \qquad \mathbb{E}\{x^{\gamma}(0)x^{\gamma}(0)^{\mathrm{T}}\} = \Sigma_0 \quad (2a)$$

$$\mathbb{E}\{x^{\gamma}(T)\} = 0, \qquad \mathbb{E}\{x^{\gamma}(T)x^{\gamma}(T)^{\mathrm{T}}\} = \Sigma_T \quad (2b)$$

where Σ_0 and Σ_T are given positive definite matrices. Next, we provide the formulation of the standard CS problem.

Problem 1: Let $\gamma \in \Gamma$ be given and let $\mathcal{U}^{\gamma}_{\mathrm{CS}}$ denote the set of admissible control inputs which consists of all adapted inputs of finite energy that will produce a state process for the uncertain system (1) that satisfies the boundary conditions (2a)-(2b). Then, find an admissible control input $u^{\star}_{\mathrm{CS}}(\cdot) \in \mathcal{U}^{\gamma}_{\mathrm{CS}}$ that minimizes the performance index

$$J^{\gamma}_{\mathrm{CS}}(u(\cdot)) := \mathbb{E}\{\int_0^T u(t)^{\mathrm{T}}u(t)\mathrm{d}t\}. \quad (3)$$

The MMCS seeks an adapted control input of finite energy that will generate state trajectories that satisfy (2a)-(2b) for all $\gamma \in \Gamma$. We denote the set of admissible inputs by $\mathcal{U}_{\mathrm{MMCS}}$, where $\mathcal{U}_{\mathrm{MMCS}} := \cap_{\gamma\in\Gamma}\mathcal{U}^{\gamma}_{\mathrm{CS}}$. Next, we formulate the MMCS problem as a minimax stochastic control problem.

Problem 2: Find an admissible control input $u^{\star}_{\mathrm{MMCS}}(\cdot) \in \mathcal{U}_{\mathrm{MMCS}}$ that minimizes the functional $\mathfrak{J}_{\mathrm{MMCS}}(u(\cdot))$, where

$$\mathfrak{J}_{\mathrm{MMCS}}(u(\cdot)) := \max_{\gamma\in\Gamma} J^{\gamma}_{\mathrm{CS}}(u(\cdot)) \quad (4)$$

subject to the dynamic constraints given in (1) and the boundary conditions given in (2a)-(2b) for all $\gamma \in \Gamma$.

Remark 1 The goal of Problem 2 is to find a control input that will furnish the minimal worst value of the control effort metric in the time interval $[0, T]$ while steering the state covariance from Σ_0 to Σ_T for all possible state space realizations of the uncertain system in $\mathcal{M}$. Therefore, the solvability of Problem 2 requires a strong form of controllability (essentially, we seek a common control law that can enforce the same terminal constraint on the state covariance (2b) for all models from $\mathcal{M}$) that it is hard to satisfy in practice (i.e., the set $\mathcal{U}_{\mathrm{MMCS}}$ may turn out to be empty).

III. Problem Analysis

In this section, we first review the solution to the standard CS problem and subsequently, we propose a soft-constrained formulation of the MMCS problem. Finally, we will characterize the structure of the solution to the MMCS problem based on the Robust Maximum Principle (RMP).

A. Solution to the Standard Covariance Steering Problem

Next, we present the basic elements of the solution to the standard (non-robust) CS problem (Problem 1) for the case of continuous-time Gaussian linear systems following [3]. In the subsequent discussion, we will drop the superscript γ.

In a nutshell, the optimal control law $u^{\star}_{\mathrm{CS}}(\cdot)$ that solves the standard CS problem is given by

$$u^{\star}_{\mathrm{CS}}(t, x) = -B(t)^{\mathrm{T}}\Pi(t)x, \quad (5)$$

where $\Pi(t)$ satisfies the Differential Riccati Equation (DRE):

$$-\dot{\Pi}(t) = A(t)^{\mathrm{T}}\Pi(t) + \Pi(t)A(t) - \Pi(t)B(t)B(t)^{\mathrm{T}}\Pi(t). \quad (6)$$

Note that in the CS problem there are no prescribed boundary conditions for (6) and terminal constraints on the state covariance (2a)-(2b) are enforced instead. This is in sharp contrast with the LQG problem, in which a terminal condition on $\Pi(T)$ is determined by a relevant terminal (quadratic) cost. A similar terminal cost is absent from the formulation of the CS problem and thus one cannot obtain a boundary condition as in the standard LQG problem. The following theorem, which summarizes key results from [3], completely characterizes the solution to the standard CS problem.

Theorem 1: The optimal control law $u^\star_{\mathrm{CS}}(\cdot) \in \mathcal{U}_{\mathrm{CS}}$ that solves the standard CS problem (Problem 1) can be written as the feedback control law given in (5) where $\Pi(t)$ satisfies the DRE (6) with initial condition $\Pi(0) = \Pi_0$, where

$$\Pi_0^{-1} = \sqrt{N(T,0)}\sqrt{S_0}\Big(S_0 + (1/2)I - \sqrt{(\sqrt{S_0}S_T\sqrt{S_0} + (1/4)I)}\Big)^{-1}\sqrt{S_0}\sqrt{N(T,0)}, \tag{7}$$

with

$$S_0 = \big(\sqrt{N(T,0)}\big)^{-1}\Sigma_0\big(\sqrt{N(T,0)}\big)^{-1} \tag{8a}$$

$$S_T = \big(\sqrt{N(T,0)}\big)^{-1}\Phi(0,T)\Sigma_T\Phi(0,T)^{\mathrm{T}}\big(\sqrt{N(T,0)}\big)^{-1} \tag{8b}$$

and $N(\cdot,\cdot)$ denotes the controllability Grammian corresponding to the pair $(A(t), B(t))$, that is,

$$N(T,0) := \int_0^T \Phi(0,\sigma)B(\sigma)B(\sigma)^{\mathrm{T}}\Phi(0,\sigma)^{\mathrm{T}}\mathrm{d}t, \tag{9}$$

where $\Phi(t,\sigma)$ denotes the fundamental matrix solution of the homogeneous linear system $\dot{x} = A(t)x$.

The state covariance of the closed-loop system satisfies the following Lyapunov (matrix) differential equation:

$$\dot{\Sigma}(t) = B(t)B(t)^{\mathrm{T}} + (A(t) - B(t)B(t)^{\mathrm{T}}\Pi(t))\Sigma(t) + \Sigma(t)(A(t) - B(t)B(t)^{\mathrm{T}}\Pi(t))^{\mathrm{T}}. \tag{10}$$

Proposition 1: The forward solution to (10), $\Sigma(t; 0, \Sigma_0)$, where $\Sigma(0; 0, \Sigma_0) = \Sigma_0$, satisfies the boundary condition $\Sigma(T; 0, \Sigma_0) = \Sigma_T$ provided that the matrix $\Pi(t)$ that appears in the right hand side of (10) corresponds to the forward solution $\Pi(t; 0, \Pi_0)$ to the DRE (6) with initial condition $\Pi(0; 0, \Pi_0) = \Pi_0$, where Π_0 is defined as in (7).

Remark 2 Both Theorem 1 and Proposition 1 imply that knowledge of matrix Π_0 unlocks the CS problem by allowing us to compute the forward solution to DRE (6) which, in turn, determines the optimal control law that steers the state covariance from Σ_0, at $t = 0$, to Σ_T, at $t = T$.

Next, we associate the CS problem with a standard LQG problem based on the observation that the optimal controller of the CS problem has exactly the same structure as the optimal controller of a relevant LQG problem. In particular, we will associate the MMCS problem with a more tractable problem which can be addressed by means of the Robust Maximum Principle (RMP) [16].

Proposition 2: Assume that the pair $(A(t), B(t))$ is controllable and let $\mathcal{S} \in \mathbb{S}^n_{++}$, where $\mathcal{S} := S_0 + I - S_T$, and S_0, S_T are defined as in (8a)-(8b). Then, the optimal control input $u^\star_{\mathrm{CS}}(\cdot) \in \mathcal{U}_{\mathrm{CS}}$ that solves the standard CS problem (Problem 1) corresponds to the optimal control law that solves the finite-horizon LQG problem for the same system described by the SDE (1) and performance index

$$J_{\mathrm{LQG}}(u(\cdot)) := \mathbb{E}\big\{x(T)^{\mathrm{T}}\Pi(T; 0, \Pi_0)x(T) + \int_0^T u(t)^{\mathrm{T}}u(t)\mathrm{d}t\big\}, \tag{11}$$

where $\Pi(T; 0, \Pi_0)$ denotes the value of the (forward) solution $\Pi(t; 0, \Pi_0)$ to the DRE equation (6) with initial condition $\Pi(0; 0, \Pi_0) = \Pi_0$, where Π_0 is given in (7), at $t = T$. In the latter LQG problem, there are no hard constraints on the terminal state covariance and the set of admissible inputs, which is denoted by $\mathcal{U}_{\mathrm{LQG}}$, consists of all adapted inputs of finite energy (with no requirement that the terminal state covariance satisfies (2b)).

Proof: First we show that the quadratic form $x(T)^{\mathrm{T}}\Pi(T; 0, \Pi_0)x(T)$ is convex, that is, $\Pi(T; 0, \Pi_0) \in \mathbb{S}^n_{+}$. Under the assumption that $(A(t), B(t))$ is controllable, which implies that $N(0,T) \in \mathbb{S}^n_{++}$, it follows that the matrix $\Psi := S_0 + (1/2)I - \sqrt{(\sqrt{S_0}S_T\sqrt{S_0} + (1/4)I)}$ is invertible (refer to [13, Proposition 3]) and thus, under the assumption that $\mathcal{S} \in \mathbb{S}^n_{++}$, where $\mathcal{S} := S_0 + I - S_T$, it follows that $\Psi \in \mathbb{S}^n_{++}$. We immediately conclude that $\Pi_0 \in \mathbb{S}^n_{++}$, where Π_0 is defined in (7). In view of [13, Proposition 3], the forward solution of DRE (6), $\Pi(t) = \Pi(t; 0, \Pi_0)$, with initial condition $\Pi(0) = \Pi_0$ will remain non-singular for all times $t \in [0, T]$. Given that the eigenvalues of $\Pi(t; 0, \Pi_0)$ will change continuously with time and $\Pi(0; 0, \Pi_0) = \Pi_0 \in \mathbb{S}^n_{++}$, we conclude that $\Pi(t; 0, \Pi_0) \in \mathbb{S}^n_{++}$, $\forall t \in [0, T]$.

From standard results on stochastic control [18], [19], we have that the optimal control law $u^\star_{\mathrm{LQG}}(t, x)$ that solves the LQG problem with performance index as in (11) is defined as follows: $u^\star_{\mathrm{LQG}}(t, x) = -B(t)^{\mathrm{T}}\Pi_-(t)x$, where $\Pi_-(t)$ denotes the backward solution of the DRE (6) with terminal condition $\Pi_-(T) = \Pi(T; 0, \Pi_0)$. We also write $\Pi_-(t) = \Pi_-(t; T, \Pi(T; 0, \Pi_0))$ from which it follows that $\Pi_-(T) = \Pi_-(T; T, \Pi(T; 0, \Pi_0)) = \Pi(T; 0, \Pi_0) = \Pi(T)$. Therefore, $\Pi(t)$ and $\Pi_-(t)$ satisfy the same DRE and the same boundary conditions at $t = 0$ and $t = T$. We conclude that $\Pi_-(t) = \Pi_-(t; T, \Pi(T; 0, \Pi_0)) = \Pi(t; 0, \Pi_0) = \Pi(t)$ $\forall t \in [0, T]$, from which it also follows that $u^\star_{\mathrm{CS}}(t, x) = u^\star_{\mathrm{LQG}}(t, x)$. ■

Remark 3 Proposition 2 does not imply that one can solve the standard CS problem by solving instead an "equivalent" LQG problem, which is a well-studied stochastic optimal control problem [18]. This is because one cannot solve the LQG problem defined in Proposition 2 without having solved first the CS problem (finding the matrix Π_0 is essentially the key ingredient of the solution to the CS problem as implied by Theorem 1), leading to a cyclic argument. The whole point of Proposition 2 is that the solution to the CS problem has the same underlying structure as an LQG problem and this observation will be leveraged later on to address the multi-model (robust) CS problem by associating it with an LQG problem with a special terminal cost together with a tractable optimization problem.

Proposition 3: The optimal cost of the CS problem, $J^\star$, where $J^\star := J(u^\star_{\mathrm{CS}}(\cdot))$, is given by

$$J^\star_{\mathrm{CS}} = \mathrm{tr}(\Sigma_0\Pi_0) - \mathrm{tr}(\Sigma_T\Pi(T; 0, \Pi_0)) + \int_0^T \mathrm{tr}\big(B(t)^{\mathrm{T}}\Pi(t; 0, \Pi_0)B(t)\big)\mathrm{d}t, \tag{12}$$

where Π_0 is defined as in (7).

Proof: First, we observe that, in view of (2a)-(2b), (3) and (11), the performance indices $J(u(\cdot))$ and $J_{\text{LQG}}(u(\cdot))$ are related to each other as follows:

$$\begin{aligned} J(u(\cdot)) &= J_{\text{LQG}}(u(\cdot)) - \mathbb{E}\{\text{tr}(\Pi(T;0,\Pi_0)x(T)x(T)^{\mathrm{T}})\} \\ &= J_{\text{LQG}}(u(\cdot)) - \text{tr}(\Sigma_T\Pi(T;0,\Pi_0)). \end{aligned} \tag{13}$$

Note that the terms $\text{tr}(\Sigma_0\Pi_0)$ and $\text{tr}(\Sigma_T\Pi(T;0,\Pi_0))$ are constant for the CS problem in view of the boundary conditions on the state covariance given in (2a)-(2b) (the latter boundary conditions do not appear in the LQG problem). Therefore, it follows from (13) that

$$J^\star_{\text{CS}} = J^\star_{\text{LQG}} - \text{tr}\left(\Sigma_T\Pi(T;0,\Pi_0)\right), \tag{14}$$

where $J^\star_{\text{LQG}}$ is the optimal value of J_{LQG} that is given by

$$J^\star_{\text{LQG}} = \text{tr}(\Sigma_0\Pi_0) + \int_0^T \text{tr}(B(t)^{\mathrm{T}}\Pi(t;0,\Pi_0)B(t))\mathrm{d}t. \tag{15}$$

To derive the last equality we have used the fact that the driving noise $\tilde{w}(t) = B(t)w(t)$ that appears in the right hand side of the SDRE (1) is a Wiener process with incremental covariance $\mathbb{E}\{\mathrm{d}\tilde{w}(t)\mathrm{d}\tilde{w}(t)^{\mathrm{T}}\} = B(t)B(t)^{\mathrm{T}}\mathrm{d}t$. The result follows readily from (14) and (15). ∎

B. Reformulation of the Multi-Model Covariance Steering

The goal of the multi-model (robust) covariance steering problem (Problem 2) is to find a single (common) control law that will steer the state covariance of each possible state space realization of the uncertain linear system from $\mathcal{M}$ to the same positive definite matrix while rendering a minimax to the control effort metric over the finite horizon $[0,T]$. In this section, we will propose an alternative, and more practical, problem formulation.

Before we proceed any further, we introduce new notation that will be used throughout this section. In particular, let $\boldsymbol{x} := \left[x^1;\ldots;x^{|\Gamma|}\right]$, $\mathbf{A}(t) := \text{bdiag}(A^1(t),\ldots,A^{|\Gamma|}(t))$, $\mathbf{B}(t) := \left[B^1(t);\ldots;B^{|\Gamma|}(t)\right]$, $\boldsymbol{\Sigma}_0 := \text{bdiag}(\Sigma_0,\ldots,\Sigma_0)$, and $\boldsymbol{\Sigma}_T := \text{bdiag}(\Sigma_T,\ldots,\Sigma_T)$. The dynamics of the whole collection of models, $\mathcal{M}$, can be described in a compact way by the following SDE:

$$\mathrm{d}\boldsymbol{x} = (\mathbf{A}(t)\boldsymbol{x} + \mathbf{B}(t)u(t))\mathrm{d}t + \mathbf{B}(t)\mathrm{d}w, \tag{16}$$

where $\boldsymbol{x} \in \mathbb{R}^{n|\Gamma|}$ denotes the augmented state (concatenation of the states of each state space model in $\mathcal{M}$ driven from the same input), whereas the input $u(t)$ is an $\mathbb{R}^m$-valued function and $w(t)$ is a $\mathbb{R}^m$-valued Wiener process with zero incremental mean and incremental covariance $I\mathrm{d}t$. Actually, the (m-dimensional) input of the ($n|\Gamma|$-dimensional) augmented system will be the same input that will be applied to the actual (n-dimensional) uncertain system regardless of its "correct" state space realization, and the same holds true for the noise process.

The multi-model covariance steering problem (Problem 2) can be formulated more precisely using the new notation as the following minimax or robust stochastic optimal control problem for the augmented system given in (16).

Problem 3 (MMCS): Find a control input $u^\star_{\text{MMCS}}(\cdot) \in \mathcal{U}_{\text{MMCS}}$ that minimizes the following performance index:

$$J_{\text{MMCS}}(u(\cdot)) := \max_{\gamma\in\Gamma} J^\gamma_{\text{CS}}(u(\cdot)), \tag{17}$$

where $J^\gamma_{\text{CS}}(u(\cdot))$ is defined as in (3), subject to the dynamics given in (16) and the following boundary conditions:

$$\mathbb{E}\{\boldsymbol{x}(0)\} = 0, \qquad \mathbb{E}\{\boldsymbol{x}(0)\boldsymbol{x}(0)^{\mathrm{T}}\} = \boldsymbol{\Sigma}_0, \tag{18}$$

$$\mathbb{E}\{\boldsymbol{x}(T)\} = 0, \qquad \mathbb{E}\{\boldsymbol{x}(T)\boldsymbol{x}(T)^{\mathrm{T}}\} = \boldsymbol{\Sigma}_T. \tag{19}$$

Problem 3 is not practical given that its goal is to find a common control input that will simultaneously steer the state covariance of all $|\Gamma|$ state space models in $\mathcal{M}$ to the same matrix $\Sigma_T \in \mathbb{S}^n_{++}$. The problem becomes more challenging (in the sense that the set of admissible inputs $\mathcal{U}_{\text{MMCS}}$ can be empty) as $|\Gamma|$ increases. To simplify the control synthesis, we will next propose a relaxed version of the latter problem.

Our approach is based on the interpretation of the CS program for any $\gamma \in \Gamma$ as an equivalent LQG problem with a special terminal cost that is determined by the forward solution to the DRE (6) for a particular initial condition. By using the latter interpretation, we will replace the boundary conditions on the terminal state covariance given in (19) with an equivalent terminal cost for each state space realization in $\mathcal{M}$. In particular, let

$$\begin{aligned} J^\gamma_{\text{LQG}}(u(\cdot)) = \mathbb{E}\{&x^\gamma(T)^{\mathrm{T}}\Pi^\gamma(T;0,\Pi^\gamma_0)x^\gamma(T) \\ &+ \int_0^T u(t)^{\mathrm{T}}u(t)\mathrm{d}t\}, \end{aligned} \tag{20}$$

where $\Pi^\gamma(T;0,\Pi^\gamma_0)$ denotes the solution to the DRE given in (6) with initial condition $\Pi^\gamma(0;0,\Pi^\gamma_0) = \Pi^\gamma_0$, where Π^γ_0 is defined as in (7), after replacing $(A(t),B(t))$ with $(A^\gamma(t),B^\gamma(t))$ in (6) and (7). Then, we consider the relaxed version of the MMCS problem based on a terminal cost in lieu of the hard constraint on the terminal state covariance.

Problem 4 (relaxed MMCS): Find an input $u^\star_{\text{rMMCS}}(\cdot) \in \mathcal{U}_{\text{LQG}}$, where $\mathcal{U}_{\text{LQG}}$, consists of all adapted inputs of finite energy, that minimizes the following performance index:

$$J_{\text{rMMCS}}(u(\cdot)) := \max_{\gamma\in\Gamma} J^\gamma_{\text{LQG}}(u(\cdot)) \tag{21}$$

where $J^\gamma_{\text{LQG}}(u(\cdot))$ is defined in (20), subject to (16).

Remark 4 Problem 4 is always well-posed without requiring any special form of controllability in contrast with Problem 3 (the latter requires the existence of a common controller that can steer the state covariance of all the models in $\mathcal{M}$ to the same matrix $\Sigma_T \in \mathbb{S}^n_{++}$). Note that the performance index defined in (21) depends on the terminal costs associated (a posteriori) with each of the $|\Gamma|$ standard CS problems corresponding to the n-dimensional state space models from $\mathcal{M}$, which can be solved a priori in order to obtain the matrices $\Pi^\gamma(T;0,\Pi^\gamma_0)$, for all $\gamma \in \Gamma$.

C. Solution to the Relaxed Problem

Next, we present the solution to the relaxed MMCS problem (Problem 4) by utilizing the RMP for multi-model LQG-type problems [16, Chapter 16] along with new insights on the structure of the weight selection problem.

Theorem 2: The optimal control law $u^\star_{\text{rMMCS}}(\cdot) \in \mathcal{U}_{\text{LQG}}$ that solves Problem 4 (relaxed MMCS problem) is given by

$$u^\star_{\text{rMMCS}}(t,\boldsymbol{x}) = -\mathbf{B}(t)^{\mathrm{T}}\sqrt{\boldsymbol{\Lambda}^\star}\mathbf{P}(t)\sqrt{\boldsymbol{\Lambda}^\star}\,\boldsymbol{x} \tag{22}$$

where $\mathbf{P}(t)$ satisfies the following DRE:

$$-\dot{\mathbf{P}}(t) = \mathbf{A}(t)^{\mathrm{T}}\mathbf{P}(t) + \mathbf{P}(t)\mathbf{A}(t) - \mathbf{P}(t)\sqrt{\mathbf{\Lambda}^{\star}}\mathbf{B}(t)\mathbf{B}(t)^{\mathrm{T}}\sqrt{\mathbf{\Lambda}^{\star}}\mathbf{P}(t) \tag{23}$$

with boundary condition $\mathbf{P}(T) = \mathbf{P}_T$, where

$$\mathbf{P}_T := \mathrm{bdiag}\big(\Pi^{\gamma_1}(T; 0, \Pi_0^{\gamma_1}), \ldots, \Pi^{\gamma_{|\Gamma|}}(T; 0, \Pi_0^{\gamma_{|\Gamma|}})\big) \tag{24}$$

where $\Pi^{\gamma}(t; 0, \Pi_0^{\gamma})$ denotes the forward solution to the DRE given in (6) with initial condition Π_0^{γ} which is defined as in Theorem 1 for a given $\gamma \in \Gamma$ (in (24), we assume that $\{\gamma_1, \ldots, \gamma_{|\Gamma|}\}$ corresponds to a particular enumeration of the finite parameter set Γ). In addition, $\mathbf{\Lambda}^{\star} := \mathrm{diag}(\lambda^{\star}) \otimes I$, where $\lambda^{\star} := [\lambda_1^{\star}, \ldots, \lambda_{|\Gamma|}^{\star}]^{\mathrm{T}} \in \mathbf{\Delta}_{|\Gamma|}$ corresponds to the minimizer of the following constrained optimization problem:

$$\underset{\lambda \in \mathbf{\Delta}_{|\Gamma|}}{\text{minimize}} \quad \mathcal{J}(\lambda) := \sum_{\gamma \in \Gamma} \lambda(\gamma) J_{\mathrm{LQG}}^{\gamma}(u_{\mathrm{rMMCS}}^{\star}(\cdot)) \tag{25}$$

where $\lambda(\gamma)$ denotes the component of λ that corresponds to a given $\gamma \in \Gamma$.

Proof: Refer to Chapter 16 from [16]. ■

Theorem 2 relies on the application of a quite general result from multi-model stochastic optimal control. Next, we will further characterize the structure of the optimization problem and reason that it can be put under the umbrella of a special class of nonlinear programs.

Proposition 4: The optimal weight vector $\lambda^{\star}$ defined in Theorem 2 corresponds to a (global) minimizer of the following constrained optimization problem:

$$\underset{\lambda \in \mathbf{\Delta}_{|\Gamma|}}{\text{minimize}} \quad \mathcal{J}(\lambda) := \mathrm{tr}\left(\mathbf{\Sigma}_0\sqrt{\mathbf{\Lambda}}\mathbf{P}(0)\sqrt{\mathbf{\Lambda}}\right) + \int_0^T \mathrm{tr}\left(\mathbf{B}(t)^{\mathrm{T}}\sqrt{\mathbf{\Lambda}}\mathbf{P}(t)\sqrt{\mathbf{\Lambda}}\mathbf{B}(t)\right)\mathrm{d}t. \tag{26}$$

In addition, the optimization problem in (26) is a non-convex quadratic program with a convex objective function subject to a single quadratic equality (non-convex) constraint.

Proof: The characterization of the optimization problem for the selection of the optimal weight vector $\lambda^{\star}$ is based on [16, Theorem 16.2]. Next, we show that the latter problem admits the structure claimed in the theorem. To this aim, let $\mu := [\sqrt{\lambda_1}, \ldots, \sqrt{\lambda_{|\Gamma|}}]^{\mathrm{T}}$ (note that by definition, $\mu \geq\geq 0$) and let $\mathbf{M} := \sqrt{\mathbf{\Lambda}} = \mathrm{diag}(\mu) \otimes I$. Then, the optimization problem in (26) can be formulated equivalently as follows:

$$\text{minimize} \quad \tilde{\mathcal{J}}(\mu) := \tilde{\mathcal{J}}_0(\mu) + \int_0^T \tilde{\mathcal{J}}_t(\mu)\mathrm{d}t \tag{27}$$

$$\text{subject to} \quad \mu^{\mathrm{T}}\mu = 1 \quad \text{and} \quad \mu \geq\geq 0 \tag{28}$$

where $\tilde{\mathcal{J}}_0(\mu) := \mathrm{tr}(\mathbf{\Sigma}_0\mathbf{M}\mathbf{P}(0)\mathbf{M})$ and $\tilde{\mathcal{J}}_t(\mu) := \mathrm{tr}\left(\mathbf{B}(t)^{\mathrm{T}}\mathbf{M}\mathbf{P}(t)\mathbf{M}\mathbf{B}(t)\right)$. First, we show that $\tilde{\mathcal{J}}_0(\mu)$ is a convex function. Indeed, $\tilde{\mathcal{J}}_0(\mu) = \|\sqrt{\mathbf{P}(0)}(\mathrm{diag}(\mu) \otimes I)\sqrt{\mathbf{\Sigma}_0}\|_F^2$. Note that the function $\phi(\mu) = \sqrt{\mathbf{P}(0)}(\mathrm{diag}(\mu) \otimes I)\sqrt{\mathbf{\Sigma}_0}$ is linear. Indeed, for any $\mu_1, \mu_2 \in \mathbb{R}^n$ and α, $\beta \in \mathbb{R}$, it follows from the associativity and bilinearity properties of the Kronecker product that $\mathrm{diag}(\alpha\mu_1 + \beta\mu_2) \otimes I = \alpha(\mathrm{diag}(\mu_1) \otimes I) + \beta(\mathrm{diag}(\mu_2) \otimes I)$, from which it follows that $\phi(\alpha\mu_1 + \beta\mu_2) = \alpha\phi(\mu_1) + \beta\phi(\mu_2)$. In addition, $\varphi(\cdot) : \mathcal{Q} \to \|\mathcal{Q}\|_F^2$ is a convex quadratic function and thus, $\tilde{\mathcal{J}}_0(\mu) = (\varphi \circ \phi)(\mu) = \varphi(\phi(\mu))$ is a convex quadratic function as the composition of a convex (quadratic) function with a linear function. Therefore, there exist $\mathcal{R}_0 \in \mathbb{S}_+^{|\Gamma|}$, $c_0 \in \mathbb{R}^{|\Gamma|}$ and $\kappa_0 \in \mathbb{R}$ such that $\mathcal{J}_0(\mu) = \mu^{\mathrm{T}}\mathcal{R}_0\mu + c_0^{\mathrm{T}}\mu + \kappa_0$. Similarly, we have $\mathcal{J}_t(\mu) = \|\sqrt{\mathbf{P}(t)}(\mathrm{diag}(\mu) \otimes I)\mathbf{B}(t)\|_F^2$, $\forall t \in [0, T]$; we conclude that $\mathcal{J}_t(\mu)$ can also be written as a convex quadratic function of μ at every $t \in [0, T]$, i.e., $\mathcal{J}_t(\mu) = \mu^{\mathrm{T}}\mathcal{R}(t)\mu + c(t)^{\mathrm{T}}\mu + \kappa(t)$, where $\mathcal{R}(\cdot) \in \mathcal{C}([0, T]; \mathbb{S}_+^{|\Gamma|})$, $c(\cdot) \in \mathcal{C}([0, T], \mathbb{R}^{|\Gamma|})$ and $\kappa(\cdot) \in \mathcal{C}([0, T]; \mathbb{R})$. Therefore, we can write $\mathcal{J}(\mu) = \mu^{\mathrm{T}}\tilde{\mathcal{R}}\mu + \tilde{c}^{\mathrm{T}}\mu + \tilde{\kappa}$, where $\tilde{\mathcal{R}} := (\mathcal{R}_0 + \int_0^T \mathcal{R}(t)\mathrm{d}t)$, $\tilde{c} := \big(c_0 + \int_0^T c(t)\mathrm{d}t\big)$ and $\tilde{\kappa} := \big(\kappa_0 + \int_0^T \kappa(t)\mathrm{d}t\big)$. Given that $\mathcal{R}(t) \in \mathbb{S}_+^{|\Gamma|}$, $\forall t \in [0, T]$, it follows that $\int_0^T \mathcal{R}(t)\mathrm{d}t \in \mathbb{S}_+^{|\Gamma|}$ and thus, we conclude that $\tilde{\mathcal{R}} \in \mathbb{S}_+^{|\Gamma|}$. Therefore, $\mathcal{J}(\mu)$ is a convex quadratic function.

Finally, the requirement that $\lambda \in \mathbf{\Delta}_{|\Gamma|}$ is equivalent to $\lambda \geq\geq 0$ and $\mathbf{1}^{\mathrm{T}}\lambda = 1$, where the latter equality can be written as $\sum_{\gamma \in \Gamma} \lambda_{\gamma} = \sum_{\gamma \in \Gamma} \mu_{\gamma}^2 = \|\mu\|^2 = 1$. The last equality together with the requirement $\mu \geq\geq 0$, which follows from the definition of μ (i.e., $\mu_{\gamma} = \sqrt{\lambda_{\gamma}}$, $\forall \gamma \in \Gamma$), establishes the constraint given in (28). This completes the proof. ■

Remark 5 The objective function in (27) is a convex quadratic function and the inequality constraint in (28) is linear; however, the constraint $\mu^{\mathrm{T}}\mu = 1$ is a quadratic equality constraint which is non-convex (the optimization problem in (27)-(28) is a special type of a non-convex quadratic program [20], [21]).

IV. Numerical Simulations

In this section, we present numerical simulations to illustrate the key ideas of this work. We consider a simple mechanical system, namely a harmonic oscillator subject to white noise, for which we have two possible models with $A^{\gamma} = \left[\begin{smallmatrix} 0 & \gamma \\ -\gamma & 0 \end{smallmatrix}\right]$ and $B^{\gamma} = \left[\begin{smallmatrix} 0 \\ 1 \end{smallmatrix}\right]$, where the uncertain parameter $\gamma \in \Gamma = \{\gamma_1, \gamma_2\}$. For our simulations, we take $\Sigma_0 = \left[\begin{smallmatrix} 18 & 3 \\ 3 & 8 \end{smallmatrix}\right]$, $\Sigma_T = \left[\begin{smallmatrix} 1.5534 & 0.0906 \\ 0.0906 & 2.7801 \end{smallmatrix}\right]$, $\gamma_1 = 1$, $\gamma_2 = 2.2$ and $T = 2$ (note that the assumption $\mathcal{S} \in \mathbb{S}_{++}^2$ from Prop. 2 is satisfied).

Figure 1 shows the evolution of the ellipsoid $\mathcal{E}(t) := \{x \in \mathbb{R}^2 : x^{\mathrm{T}}\Sigma(t)^{-1}x \leq 1\}$, for $t \in [0, T]$, which we will refer to as the covariance path (the blue and red ellipses correspond to, respectively, Σ_0 and Σ_T). The covariance paths generated by three different controllers are illustrated in the left (respectively, right) sub-figures in Fig. 1 for the state space model with $\gamma = \gamma_1$ (respectively, $\gamma = \gamma_2$). In particular, Fig. 1a shows the covariance paths corresponding to the solutions to the standard CS problem for the two models (in the absence of model uncertainty), whereas Fig. 1b and Fig. 1c illustrate the covariance paths in the presence of model uncertainty for $\lambda_1 = 0.90$ (and $\lambda_2 = 0.10$) and $\lambda_1 = 0.10$ (and $\lambda_2 = 0.90$), respectively. Note that the covariance paths in Fig. 1b and Fig. 1c were generated by controllers that solve the robust CS problem with a pair of a priori fixed mixing weights that do not correspond to the minimax to the robust optimal control problem. These figures are included to illustrate the role of the mixing weight $\lambda = [\lambda_1,\ \lambda_2]^{\mathrm{T}} \in \mathbf{\Delta}_2$ in the solution to the robust CS problem. In particular, the covariance path in Fig. 1b for both models is essentially based on the controller of the system with $\gamma = \gamma_1$ (since the weight λ_1 takes a value close to unity) and this is why the left sub-figures in Fig. 1a and Fig. 1b are very similar. By contrast, the right sub-figures in Fig. 1a and Fig. 1b differ significantly from each other because the covariance path shown in the right sub-figure in Fig. 1b, which refers to the model with $\gamma = \gamma_2$, was generated by a controller that was primarily based on the other model ($\gamma = \gamma_1$) given that $\lambda_2 = 0.10$. The situation is reversed in Fig. 1c, where the covariance paths are generated by a controller that is mainly based on the second model ($\gamma = \gamma_2$).

It turns out that the minimax solution in this example corresponds to the CS problem for the first model ($\lambda_1 = 1$ and $\lambda_2 = 0$). Similar results would have been obtained if in our simulation study we have considered more than two models (typically, the minimax solution would be determined by the solution to the CS problem for one particular model corresponding to one of the vertices of the simplex). These observations are in agreement with the interpretation of the solutions to the minimax linear quadratic optimal control problems and supporting simulations reported in [16], [22].

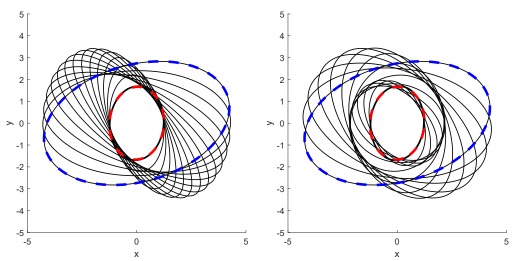

(a) Covariance path with no model uncertainty for $\gamma = \gamma_1$ (left sub-figure) and $\gamma = \gamma_2$ (right sub-figure).

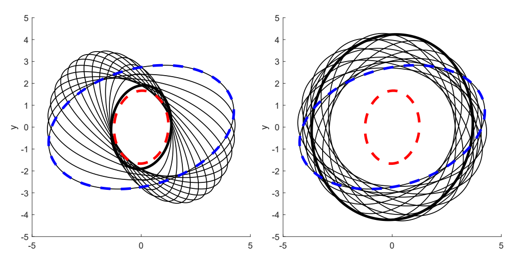

(b) Covariance path when $\lambda_1 = 0.90$ ($\lambda_2 = 0.10$), for $\gamma = \gamma_1$ (left sub-figure) and $\gamma = \gamma_2$ (right sub-figure).

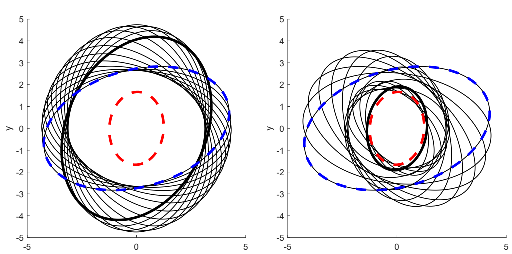

(c) Covariance path when $\lambda_1 = 0.10$ and $\lambda_2 = 0.90$, for $\gamma = \gamma_1$ (left sub-figure) and $\gamma = \gamma_2$ (right sub-figure).

Fig. 1: Covariance paths for the two models of the uncertain system driven by their nominal CS controllers (Fig. 1a) and (suboptimal) robust CS-based controllers with mixing weights $\lambda_1 = 0.90$, $\lambda_2 = 0.10$ (Fig. 1b) and $\lambda_1 = 0.10$, $\lambda_2 = 0.90$ (sub-figure 1c).

V. Concluding Remarks

In this work, we have addressed the multi-model CS problem based on its formulation as a robust stochastic linear quadratic problem whose solution relies on the solution to an LQG-type problem defined in an augmented state space and a tractable, yet non-convex, optimization problem. In our future work, we will explore different problem formulations that can lead to practical controllers which can, for instance, steer the state covariance of the uncertain system to a positive definite matrix that corresponds to the "best" compromise, in an appropriate sense, solution for all possible models.

References

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

[2] E. Collins and R. Skelton, "Covariance control discrete systems," in *1985 24th IEEE CDC*, pp. 542–547, 1985.

[3] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal steering of a linear stochastic system to a final probability distribution, part i," *IEEE Trans. Automat. Contr*, vol. 61, no. 5, pp. 1158–1169, 2015.

[4] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal steering of a linear stochastic system to a final probability distribution, part ii," *IEEE Trans. Automat. Contr*, vol. 61, no. 5, pp. 1170–1180, 2016.

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

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

[7] G. Kotsalis, G. Lan, and A. S. Nemirovski, "Convex optimization for finite-horizon robust covariance control of linear stochastic systems," *SIAM J. Control Optim.*, vol. 59, no. 1, pp. 296–319, 2021.

[8] G. Rapakoulias and P. Tsiotras, "Discrete-time optimal covariance steering via semidefinite programming," in *IEEE CDC*, pp. 1802–1807, 2023.

[9] I. M. Balci and E. Bakolas, "Constrained minimum variance and covariance steering based on affine disturbance feedback control parameterization," *International Journal of Robust and Nonlinear Control*, vol. 34, no. 11, pp. 7332–7370, 2024.

[10] A. Tsolovikos and E. Bakolas, "Cautious nonlinear covariance steering using variational gaussian process predictive models," *IFAC-PapersOnLine*, vol. 54, no. 20, pp. 59–64, 2021.

[11] J. Pilipovsky and P. Tsiotras, "Data-driven covariance steering control design," in *IEEE CDC*, pp. 2610–2615, 2023.

[12] J. W. Knaup and P. Tsiotras, "Adaptive dual covariance steering with active parameter estimation," in *IEEE CDC*, pp. 659–664, IEEE, 2024.

[13] K. Ito and K. Kashima, "Maximum entropy optimal density control of discrete-time linear systems and schrödinger bridges," *IEEE Transactions on Automatic Control*, vol. 69, no. 3, pp. 1536–1551, 2023.

[14] A. D. Saravanos, I. M. Balci, E. Bakolas, and E. A. Theodorou, "Distributed model predictive covariance steering," in *IEEE/RSJ IROS*, pp. 5740–5747, 2024.

[15] A. Khalil, Y. Lee, and E. Bakolas, "Distributed and localized covariance control of coupled systems: A system level approach," *IEEE Control Systems Letters*, vol. 9, pp. 180–185, 2025.

[16] V. G. Boltyanski and A. S. Poznyak, *The robust maximum principle: theory and applications*. Springer Science & Business Media, 2011.

[17] A. S. Poznyak, *Advanced Mathematical Tools for Automatic Control Engineers: Volume 2: Stochastic Systems*, vol. 2. Elsevier, 2009.

[18] K. J. Åström, *Introduction to Stochastic Control Theory*. Courier Corporation, 2012.

[19] J. L. Speyer and W. H. Chung, *Stochastic processes, estimation, and control*. SIAM, 2008.

[20] H. Hmam, "Quadratic optimization with one quadratic equality constraint," technical report, DSTO, Australia, 2010.

[21] C. J. Albers, F. Critchley, and J. C. Gower, "Quadratic minimisation problems in statistics," *Journal of Multivariate Analysis*, vol. 102, no. 3, pp. 698–713, 2011.

[22] A. Poznyak, T. Duncan, B. Pasik-Duncan, and V. Boltyanski, "Robust maximum principle for multi-model lq-problem," *International Journal of Control*, vol. 75, no. 15, pp. 1170–1177, 2002.