

Distributed and Localized Covariance Control of Coupled Systems: A System Level Approach

Ahmed Khalil, Yoonjae Lee, and Efstathios Bakolas

Abstract—This work is concerned with the finite-horizon optimal covariance steering of networked systems governed by discrete-time stochastic linear dynamics. In contrast to existing work that has only considered systems with dynamically decoupled ‘agents,’ we consider a dynamically coupled system composed of interconnected subsystems. In particular, we propose a *distributed* algorithm to compute the *localized* covariance steering policy for each individual subsystem. To this end, we first formulate the localized covariance steering problem and, leveraging the system-level synthesis (SLS) framework, recast this problem as a convex optimization problem over system responses. We then examine the problem’s separability and introduce a problem transformation to address instances with nonseparable objective functions. Finally, we develop a distributed algorithm based on the consensus alternating direction method of multipliers (ADMM) to distribute computation across subsystems based on their local information and communication constraints. We demonstrate the effectiveness of our proposed algorithm on a power system with 36 interconnected subsystems.

Index Terms—Distributed control, networked control systems, stochastic optimal control.

I. INTRODUCTION

THE widespread emergence of large-scale, highly interconnected systems highlights the critical need for scalable controllers that rely only on local information. In complex, safety-critical systems such as power grids, these controllers must ensure the network’s reliable operation and optimal performance despite inherent uncertainties and noise. A promising control paradigm addressing these challenges is Covariance Steering (CS), which focuses on guiding the evolution of the system’s state distribution while enforcing distributional constraints at the terminal state. By explicitly accounting for stochastic disturbances, CS ensures that the controller achieves its desired mean and covariance targets. CS differs from standard stochastic optimal control methods, such as linear-quadratic-Gaussian (LQG) control [1], where uncertainty is only controlled implicitly through shaping the cost function.

Infinite-horizon variants of CS problems were first formulated and studied in the 1980s [2], [3]. Recently, the finite-horizon counterparts have gained traction, particularly within the controls community [4]–[10]. CS has been successfully applied in areas such as path planning [8], [11], [12], trajectory optimization [6], [7], [9], and robotic manipulation [13]. For discrete-time systems, previous works have presented ways to formulate finite-horizon CS problems as convex optimization problems [14] by utilizing various parameterizations, such as the state history feedback policy [15], the disturbance feedback policy [10], and the auxiliary variable policy [11].

Distributed CS problems for multi-agent systems have been addressed using similar parametrizations [16]–[18]. However,

existing approaches overlook the dynamic coupling and the localized nature of interconnected systems, treating the latter as decoupled ‘agents.’ In tightly coupled systems like power grids, however, distributed controllers require accounting for interactions between subsystems, as disturbances can propagate and amplify across the network, potentially leading to undesirable effects. Designing such a controller, however, is known to be challenging due to nonconvexity issues that arise when attempting to apply parameterizations [19]. This motivates our use of the System Level Synthesis (SLS) framework [19]–[22], a method to synthesize a controller over the closed-loop behavior of the system rather than the controller itself. This approach particularly allows for the enforcement of local communication constraints while preserving the problem’s convex structure, admitting the use of distributed optimization techniques such as the Alternating Direction Method of Multipliers (ADMM) [23].

In this work, we address the finite-horizon *localized* CS problem involving dynamically coupled systems subject to discrete-time stochastic linear dynamics. In sharp contrast to previous work that has utilized the SLS framework to formulate distributed and localized Model Predictive Control (MPC) problems under the case of no driving noise [21] or deterministic and bounded noise [22], this work focuses on the distributed and localized steering of state distributions in the presence of Gaussian random disturbances. Our contributions are three-fold. First, we show that the localized CS problem can be formulated as a convex optimization problem over system responses via the SLS framework. Second, we introduce a problem transformation to address instances with nonseparable objectives while maintaining the separable structures of certain constraints. Lastly, we propose a distributed algorithm with provable convergence, based on consensus ADMM [23], to distribute computation across subsystems while accounting for their local communication and information constraints.

The rest of the paper is structured as follows. Section II introduces the localized covariance steering problem in coupled systems. Section III presents a system-level parametrization approach and its distributed solution method. Section IV contains a numerical simulation of the proposed method. Section V concludes this work with directions for future research.

Notation: Let \mathbb{N} denote the set of non-negative integers, and let \mathbb{S}_+^n and \mathbb{S}_{++}^n denote the convex cones of $n \times n$ (real symmetric) positive semidefinite and positive definite matrices, respectively. Let $\mathbf{1}$ denote the vector of ones. The mean and covariance of a random vector x are denoted as $\mathbb{E}[x]$ and $\text{Cov}[x]$, respectively. A Gaussian random vector with mean μ and covariance Σ is denoted as $x \sim \mathcal{N}(\mu, \Sigma)$. Let $\text{diag}(a_1, \dots, a_N)$ denote a diagonal matrix with scalars a_1, \dots, a_N and $\text{blkdiag}(A_1, \dots, A_N)$ denote a block diagonal matrix with matrices A_1, \dots, A_N . The Frobenius norm of a matrix A is denoted as $\|A\|_F := \sqrt{\text{Tr}(A^\top A)}$. Let $A(\mathcal{R}, \mathcal{C})$ denote the submatrix of a matrix A formed by selecting the rows indexed by \mathcal{R} and the columns indexed by \mathcal{C} , with $:$ denoting all indices. Lastly, \circ denotes the elementwise matrix multiplication (Hadamard product).

The first two authors contributed equally as co-first authors.

The authors are with the Department of Aerospace Engineering and Engineering Mechanics, the University of Texas at Austin, Austin, Texas 78712-1221, USA. Email addresses: akhalil@utexas.edu, yol033@utexas.edu, bakolas@austin.utexas.edu.

II. PROBLEM FORMULATION

Consider a discrete-time, linear time-varying (LTV) stochastic system composed of N interconnected subsystems:

$$\begin{aligned} \underbrace{\begin{bmatrix} x_{t+1}^1 \\ \vdots \\ x_{t+1}^N \end{bmatrix}}_{x_{t+1}} &= \underbrace{\begin{bmatrix} A_t^{11} & \dots & A_t^{1N} \\ \vdots & \ddots & \vdots \\ A_t^{N1} & \dots & A_t^{NN} \end{bmatrix}}_{A_t} \underbrace{\begin{bmatrix} x_t^1 \\ \vdots \\ x_t^N \end{bmatrix}}_{x_t} \\ &+ \underbrace{\begin{bmatrix} B_t^1 & & \\ & \ddots & \\ & & B_t^N \end{bmatrix}}_{B_t} \underbrace{\begin{bmatrix} u_t^1 \\ \vdots \\ u_t^N \end{bmatrix}}_{u_t} + \underbrace{\begin{bmatrix} w_t^1 \\ \vdots \\ w_t^N \end{bmatrix}}_{w_t}, \quad (1) \end{aligned}$$

where $x_t^i \in \mathbb{R}^{n_i}$, $u_t^i \in \mathbb{R}^{m_i}$, and $w_t^i \in \mathbb{R}^{n_i}$ denote the local state, control input, and process noise of subsystem i , respectively, with $A_t^{ij} \in \mathbb{R}^{n_i \times n_j}$ and $B_t^i \in \mathbb{R}^{n_i \times m_i}$. The local disturbance w_t^i is assumed to be a white Gaussian noise process with $\mathbb{E}[w_t^i] = 0$ and $\mathbb{E}[w_t^i w_{t'}^{i\top}] = \delta(t, t') S_i$, where $S_i \in \mathbb{S}_{++}^{n_i}$, and $\delta(t, t') = 1$ if $t = t'$, and $\delta(t, t') = 0$ otherwise. Additionally, $x_t \in \mathbb{R}^n$, $u_t \in \mathbb{R}^m$, and $w_t \in \mathbb{R}^n$ denote the global state, control input, and process noise of the system, respectively, with $A_t \in \mathbb{R}^{n \times n}$ and $B_t \in \mathbb{R}^{n \times m}$, where $n := \sum_{i=1}^N n_i$ and $m := \sum_{i=1}^N m_i$.

Let $\mathcal{G} := (\mathcal{V}, \mathcal{E})$ represent the system graph, where $\mathcal{V} := \{1, 2, \dots, N\}$ is a set of vertices (i.e., subsystem indices) and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is a set of edges such that $\forall i, j \in \mathcal{V}$, $(i, j) \in \mathcal{E} \iff \exists t : A_t^{ji} \neq 0$. To capture communication constraints between subsystems, we adopt the following definition [19].

Definition 1 (*d*-Hop Neighbors). For a locality parameter $d \in \mathbb{N}$, the sets of *d*-outgoing neighbors and *d*-incoming neighbors of subsystem i are defined as $\mathcal{O}_i(d) := \{j \in \mathcal{V} : \text{dist}(i, j) \leq d\}$ and $\mathcal{I}_i(d) := \{j \in \mathcal{V} : \text{dist}(j, i) \leq d\}$, respectively, where $\text{dist}(i, j)$ denotes the distance between (i.e., the number of edges in a shortest directed path connecting) vertices $i, j \in \mathcal{V}$.

The local control policy γ^i of subsystem i is said to be *admissible* if each u_t^i depends on t and the past global state histories $\{x_0, x_1, \dots, x_t\}$. Furthermore, an admissible policy γ^i is said to be *d-localized* if each u_t^i depends only on t and the past local state histories of its *d*-incoming neighbors:

$$u_t^i = \gamma^i(t, \{x_0^j, x_1^j, \dots, x_t^j\}_{j \in \mathcal{I}_i(d)}). \quad (2)$$

Problem 1. Consider the coupled system (1) with corresponding system graph \mathcal{G} . Let $\mu_0, \mu_f \in \mathbb{R}^n$, $\Sigma_0, \Sigma_f \in \mathbb{S}_{++}^n$, and $d, T \in \mathbb{N}$. Let $\{Q_t, R_t\}_{t=0}^{T-1}$ be given such that $Q_t \in \mathbb{S}_{++}^n$ and $R_t \in \mathbb{S}_{++}^m$ for all t . Assume, without loss of generality, $\mu_f = 0$. Find a set of admissible control policies $\{\gamma^i\}_{i \in \mathcal{V}}$ that solves the following localized covariance steering problem:¹

$$\begin{aligned} \underset{\{\gamma^i\}_{i \in \mathcal{V}}}{\text{minimize}} \quad & \sum_{t=0}^{T-1} \mathbb{E}[x_t^\top Q_t x_t + u_t^\top R_t u_t] \quad (3a) \\ \text{subject to} \quad & (1), \quad \forall t \in \{0, 1, \dots, T-1\}, \quad (3b) \\ & (2), \quad \forall t \in \{0, 1, \dots, T-1\}, \quad \forall i \in \mathcal{V}, \quad (3c) \\ & x_0 \sim \mathcal{N}(\mu_0, \Sigma_0), \quad (3d) \\ & \mathbb{E}[x_T] = \mu_f, \quad (3e) \\ & (\Sigma_f - \text{Cov}[x_T]) \in \mathbb{S}_{+}^n. \quad (3f) \end{aligned}$$

¹Without (3c), problem (3) reduces to the standard optimal covariance steering problem, which can be transformed into and solved as a convex semi-definite programming problem; see [5].

Note that the constraint in (3f) is a convex relaxation of the non-convex equality constraint, $\text{Cov}[x_T] = \Sigma_f$, which establishes an upper bound on the system's uncertainty in reaching the desired terminal mean μ_f [5].

Remark 1. Existing work on distributed CS overlooks subsystem coupling and instead focuses on decoupled ‘agents’ whose local state evolves according to [16, Equation (17a)]:

$$x_{t+1}^i = A_t^i x_t^i + B_t^i u_t^i + w_t^i, \quad (4)$$

with local state matrices $A_t^i \in \mathbb{R}^{n_i \times n_i}$. Unlike the interconnected subsystem dynamics in (1), each agent’s dynamics (4) evolves based only on its own state, input, and noise, thereby making it incapable of modeling tightly coupled systems such as power grids.

Remark 2. Additionally, related work imposes terminal covariance constraints of the form [16, Equation (19)]:

$$(\Sigma_f^i - \text{Cov}[x_T^i]) \in \mathbb{S}_{+}^{n_i}, \quad \forall i \in \mathcal{V}, \quad (5)$$

for some $\Sigma_f^i \in \mathbb{S}_{++}^{n_i}$, $\forall i \in \mathcal{V}$. These ‘local’ constraints ensure that each agent i meets an individual covariance requirement on their own states. In contrast, in Problem 1, we consider a more general constraint, namely (3f), which acts on the global state of the system. Clearly, this global constraint naturally subsumes the local constraints in (5).

III. MAIN RESULTS

In this work, we will restrict our attention to admissible control policies associated with causal LTV controllers: $\forall t$, $u_t = K_{t,0}x_0 + K_{t,1}x_1 + \dots + K_{t,t}x_t$, where $K_{t,t'} \in \mathbb{R}^{m \times n}$, $\forall t' \in \{0, 1, \dots, t\}$. One natural way to impose the communication constraints (3c) on the structure of this controller is to enforce the sparsity constraints $[K_{t,t'}]^{ij} = 0$ (for all t, t' and all i, j with $i \notin \mathcal{O}_j(d)$), where $[K_{t,t'}]^{ij}$ is a submatrix of $K_{t,t'}$ such that $u_t^i = \sum_{t'=0}^t \sum_{j=1}^N [K_{t,t'}]^{ij} x_{t'}^j$. However, it is known that for certain (e.g., strongly connected²) systems, such a constraint can lead to a nonconvex formulation after applying parameterizations [19, Section 3.5]. To circumvent this issue, we will utilize the SLS framework [19], in which the closed-loop response of the system, rather than the controller itself, is directly designed.

A. System Level Synthesis and Locality Constraints

Let $\mathbf{x} := [x_0^\top, x_1^\top, \dots, x_T^\top]^\top$, $\mathbf{u} := [u_0^\top, u_1^\top, \dots, u_{T-1}^\top]^\top$, and $\mathbf{w} := [w_0^\top, w_1^\top, \dots, w_{T-1}^\top]^\top$. Additionally, let $\mathbf{A} := [\text{blkdiag}(A_0, A_1, \dots, A_{T-1}) \ 0_{Tn \times n}]$, and $\mathbf{B} := [\text{blkdiag}(B_0, B_1, \dots, B_{T-1})]$. Using these definitions, the system dynamics (1) can be compactly written as

$$\mathbf{x} = \mathbf{Z}\mathbf{A}\mathbf{x} + \mathbf{Z}\mathbf{B}\mathbf{u} + \mathbf{w}, \quad (6)$$

where \mathbf{Z} is the block-lower shift matrix, which has identity matrices along its first block subdiagonals and zeros elsewhere. Let \mathbf{K} be the block-lower-triangular (BLT) matrix defined by

$$\mathbf{K} := \begin{bmatrix} K_{0,0} & 0 & \dots & 0 & 0 \\ K_{1,0} & K_{1,1} & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ K_{T-1,0} & \dots & K_{T-1,T-2} & K_{T-1,T-1} & 0 \end{bmatrix}.$$

²That is, the corresponding system graph \mathcal{G} is *strongly connected*, meaning there exists a directed path from any vertex to any other vertex in the graph.

The closed-loop behavior of the system (1) under the feedback gain \mathbf{K} can then be characterized as follows:

$$\mathbf{x} = \underbrace{(I - \mathbf{Z}(\mathbf{A} + \mathbf{BK}))^{-1}}_{=: \Phi_x} \mathbf{w}, \quad (7a)$$

$$\mathbf{u} = \underbrace{\mathbf{K}(I - \mathbf{Z}(\mathbf{A} + \mathbf{BK}))^{-1}}_{=: \Phi_u} \mathbf{w}, \quad (7b)$$

where Φ_x and Φ_u are BLT matrices called the state and control *system responses*, respectively. In the sequel, we will write $\Phi_{\bullet}^{t,t'}$ to denote the (t, t') th entry of Φ_{\bullet} , $\forall \bullet \in \{x, u\}$.

By virtue of the following theorem, instead of optimizing over the feedback gain \mathbf{K} , the controller synthesis can be performed by optimizing over the system responses $\{\Phi_x, \Phi_u\}$.

Theorem 1. [19, Theorem 2.1] For the LTV system dynamics (1) evolving under the feedback control law $\mathbf{u} = \mathbf{K}\mathbf{x}$, the following statements are true:

- 1) The affine subspace defined by

$$\mathbf{Z}_{\mathbf{A}, \mathbf{B}} \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} = I \quad (8)$$

parametrizes all possible system responses (7), where $\mathbf{Z}_{\mathbf{A}, \mathbf{B}} := [I - \mathbf{Z}\mathbf{A} \quad -\mathbf{Z}\mathbf{B}]$.

- 2) For any pair $\{\Phi_x, \Phi_u\}$ of BLT matrices satisfying (8), the feedback gain $\mathbf{K} = \Phi_u \Phi_x^{-1}$ achieves the desired closed-loop response.

Remark 3. If constraint (8) holds, the diagonal entries of Φ_x (i.e., $\Phi_x^{t,t}, \forall t$) are identity matrices, implying Φ_x is invertible.

The construction of the feedback gain \mathbf{K} requires the inversion of the matrix Φ_x . However, it is known that this can be avoided via the following controller realization (see [19, Section 2.1]):

$$u_t = \sum_{t'=1}^{T-1} \Phi_u^{t,t'} \hat{w}_{t-t'}, \quad (9)$$

where $\hat{w}_t = x_{t+1} - \hat{x}_{t+1}$ and $\hat{x}_{t+1} = \sum_{t'=1}^{T-1} \Phi_x^{t+1,t'+1} \hat{w}_{t-t'}$.

The main advantage of the SLS framework is that the communication constraints (3c) can be cast as affine (subspace) constraints on the system responses $\{\Phi_x, \Phi_u\}$, as summarized in the following definitions from [19].³

Definition 2 (*d*-Localized System Responses). Let $[\Phi_x]^{ij}$ be the submatrix of Φ_x that maps the local noise history $\mathbf{w}^j := [x_0^j, w_0^j, w_1^j, \dots, w_{T-1}^j]^\top$ of subsystem j to the local state history $\mathbf{x}^i := [x_0^i, x_1^i, \dots, x_T^i]^\top$ of subsystem i . Then, Φ_x is said to be *d*-localized if $[\Phi_x]^{ij} = 0, \forall j \in \mathcal{V}, \forall i \notin \mathcal{O}_j(d)$. An analogous definition holds for Φ_u .

Definition 3 (*d*-Locality Constraints). A subspace $\mathcal{L}(d)$ is said to be a *d*-locality constraint if

$$\begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} \in \mathcal{L}_d \quad (10)$$

implies that Φ_x is *d*-localized and Φ_u is $(d+1)$ -localized.

³Note that the *d*-locality constraint (10) does not necessarily ensure that the feedback gain $\mathbf{K} = \Phi_u \Phi_x^{-1}$ itself is *d*-localized. However, the controller realization in (9) is ensured to be *d*-localized, as any structural constraints imposed on the system responses $\{\Phi_x, \Phi_u\}$ are directly translated; see [19, Section 2.1] or [22, Section III-B] for further details.

B. Localized SLS Covariance Steering Problem

We will now employ the SLS framework to transform the localized CS problem (3) into an SLS problem over system responses. Excluding the constraints (3c) at the moment, problem (3) can be written as

$$\underset{\mathbf{x}, \mathbf{u}}{\text{minimize}} \quad \mathbb{E} \left[\left\| \mathbf{F}^{\frac{1}{2}} \begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} \right\|_2^2 \right] \quad (11a)$$

$$\text{subject to} \quad (6), \quad (11b)$$

$$\mathbf{P}_0 \mathbf{x} \sim \mathcal{N}(\mu_0, \Sigma_0), \quad (11c)$$

$$\mathbb{E}[\mathbf{P}_T \mathbf{x}] = \mu_f, \quad (11d)$$

$$(\Sigma_f - \text{Cov}[\mathbf{P}_T \mathbf{x}]) \in \mathbb{S}_+^n, \quad (11e)$$

where $\mathbf{Q} := \text{blkdiag}(Q_0, Q_1, \dots, Q_{T-1}, 0)$, $\mathbf{R} := \text{blkdiag}(R_0, R_1, \dots, R_{T-1})$, $\mathbf{F} := \text{blkdiag}(\mathbf{Q}, \mathbf{R})$, and \mathbf{P}_t is a $1 \times (T+1)$ block matrix, where all blocks are $n \times n$ zero matrices except for the $(t+1)$ th block which is an identity matrix. After some algebraic manipulations, the mean and covariance of \mathbf{x} can be written in Φ_x as

$$\mathbb{E}[\mathbf{x}] = \Phi_x \boldsymbol{\mu}_w, \quad \text{Cov}[\mathbf{x}] = \Phi_x \boldsymbol{\Sigma}_w \Phi_x^\top, \quad (12)$$

where $\boldsymbol{\mu}_w := [\mu_0^\top, 0^\top]^\top$ and $\boldsymbol{\Sigma}_w := \text{Cov}[\mathbf{w}] = \text{blkdiag}(\Sigma_0, W_0, W_1, \dots, W_{T-1})$ with $W_t := \text{Cov}[w_t]$. In view of (12) and Remark 3, and with the addition of (10) in place of (3c), problem (11) can be formulated as the following convex localized SLS problem:^{4,5}

$$\underset{\Phi_x, \Phi_u}{\text{minimize}} \quad f(\Phi_x, \Phi_u) := \left\| \mathbf{F}^{\frac{1}{2}} \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} \boldsymbol{\Theta}^{\frac{1}{2}} \right\|_F^2 \quad (13a)$$

$$\text{subject to} \quad (8), (10), \quad (13b)$$

$$\mathbf{P}_T \Phi_x \mathbf{P}_0^\top \mu_0 = \mu_f, \quad (13c)$$

$$\begin{bmatrix} \Sigma_f & \mathbf{P}_T \Phi_x \boldsymbol{\Sigma}_w^{\frac{1}{2}} \\ (\mathbf{P}_T \Phi_x \boldsymbol{\Sigma}_w^{\frac{1}{2}})^\top & I \end{bmatrix} \succeq 0, \quad (13d)$$

where $\boldsymbol{\Theta} := \boldsymbol{\Sigma}_w + \boldsymbol{\mu}_w \boldsymbol{\mu}_w^\top$. The linear-matrix-inequality (LMI) constraint (13d) is obtained by applying Schur's complement formula to the terminal covariance constraint:

$$(\Sigma_f - \mathbf{P}_T \Phi_x \boldsymbol{\Sigma}_w \Phi_x^\top \mathbf{P}_T) \in \mathbb{S}_+^n.$$

For the convexity of constraint (13d), see [5, Proposition 3]. Note that the localized SLS CS problem (13) includes constraint (13c) that couples the initial and target means and constraint (13d) that couples the initial and target covariances.

C. Separability of the Localized SLS CS Problem

Solving the convex but centralized SLS problem (13) requires knowledge of all system and problem parameters, including the system matrices $\{\mathbf{A}, \mathbf{B}\}$, cost matrices $\{\mathbf{Q}, \mathbf{R}\}$, and the target mean μ_f and covariance Σ_f of the global system state. However, this assumption can be overly restrictive in practice, as each subsystems i may have access to only local information on the aforementioned variables. To address this scenario, we will investigate the separable structure of problem (13) by utilizing the notion of *separability* [20]. This structural insight will be utilized in the subsequent section to

⁴Note that each system response is taken from a set of BLT matrices, which is a convex subset of the set of real matrices of the same dimension.

⁵While *d*-locality constraints are always convex, for some d , there may not exist system responses satisfying both (8) and (10). Hence, the value of d must be selected with great care (see, for instance, [24]). In this work, similar to [21], [22], we assume that, for a given d , the localized SLS problem (13) is feasible.

distribute computation across subsystems based on their local information and communication constraints.

Let $\mathcal{C} := \{\mathcal{C}^1, \dots, \mathcal{C}^N\}$ and $\mathcal{R}_\bullet := \{\mathcal{R}_\bullet^1, \dots, \mathcal{R}_\bullet^N\}$ be partitions of the sets of indices of the columns of Φ_x (or Φ_u) and the rows of Φ_\bullet , respectively, $\forall \bullet \in \{x, u\}$. Additionally, let $\mathcal{S} := \{\mathcal{S}^1, \dots, \mathcal{S}^N\}$ be a partition of the set $\{1, \dots, n\}$.

First, both the system-level parametrization constraint (8) and the locality constraint (10) are known to be *column-wise separable* with respect to the column-wise partition \mathcal{C} (see [20, Section III-A]): (8) $\iff \mathbf{Z}_{\mathbf{A}, \mathbf{B}} \begin{bmatrix} \Phi_x(:, \mathcal{C}^i) \\ \Phi_u(:, \mathcal{C}^i) \end{bmatrix} = \mathbf{I}(:, \mathcal{C}^i)$,

$$\forall i \in \mathcal{V}; \text{ and (10) } \iff \begin{bmatrix} \Phi_x(:, \mathcal{C}^i) \\ \Phi_u(:, \mathcal{C}^i) \end{bmatrix} \in \mathcal{L}_d(:, \mathcal{C}^i), \forall i \in \mathcal{V}.$$

Subsequently, the terminal mean constraint (13c) is *row-wise separable* with respect to the partition \mathcal{S} , namely (13c) $\iff [\mathbf{P}_T \Phi_x \mathbf{P}_0^\top](\mathcal{S}^i, :) \mu_0 = \mu_f(\mathcal{S}^i, :)$, $\forall i \in \mathcal{V}$.⁶ In contrast, the global terminal covariance constraint (13d) is generally not separable due to the existence of nondiagonal elements of Σ_f ; as such, exceptions include the special instance with local terminal covariance constraints as described in Remark 2, in which case Σ_f is diagonal.

Lastly, the objective function f in (13a) is generally nonseparable. However, if \mathbf{F} is diagonal, then f is *row-wise separable* with respect to the row-wise partitions \mathcal{R}_x and \mathcal{R}_u : $f(\Phi_x, \Phi_u) = \sum_{i \in \mathcal{V}} (\sum_{r \in \mathcal{R}_x^i} \mathbf{Q}(r, r) \Phi_x(r, :)) \Theta \Phi_x(r, :)^T + \sum_{s \in \mathcal{R}_u^i} \mathbf{R}(s, s) \Phi_u(s, :)) \Theta \Phi_u(s, :)^T$. Further, if Θ is diagonal, then f is *column-wise separable* with respect to the column-wise partition \mathcal{C} : $f(\Phi_x, \Phi_u) = \sum_{i \in \mathcal{V}} \sum_{c \in \mathcal{C}^i} (\Theta(c, c) \Phi_x(:, c)^T \mathbf{Q} \Phi_x(:, c) + \Phi_u(:, c)^T \mathbf{R} \Phi_u(:, c))$. To address general instances in which these conditions do not necessarily hold, the proposition below provides a transformation that will allow one to recast problem (13) as an equivalent problem with a separable objective function while maintaining the separable structures of the aforementioned constraints.

Proposition 1. Let the pair $\{\Psi_x^*, \Psi_u^*\}$ of system responses be a solution of the following convex SLS problem:

$$\underset{\Psi_x, \Psi_u}{\text{minimize}} \quad \tilde{f}(\Psi_x, \Psi_u) := \left\| \mathbf{F}^{\frac{1}{2}} \begin{bmatrix} \Psi_x \\ \Psi_u \end{bmatrix} \Lambda^{\frac{1}{2}} \right\|_F^2 \quad (14a)$$

$$\text{subject to} \quad \mathbf{Z}_{\mathbf{A}, \mathbf{B}} \begin{bmatrix} \Psi_x \\ \Psi_u \end{bmatrix} = \mathbf{V}, \quad (14b)$$

$$\begin{bmatrix} \Psi_x \\ \Psi_u \end{bmatrix} \mathbf{V}^\top \in \mathcal{L}_d, \quad (14c)$$

$$\mathbf{P}_T \Psi_x \mathbf{V}^\top \mathbf{P}_0^\top \mu_0 = \mu_f, \quad (14d)$$

$$\begin{bmatrix} \Sigma_f & \mathbf{P}_T \Psi_x \mathbf{V}^\top \Sigma_w^{\frac{1}{2}} \\ (\mathbf{P}_T \Psi_x \mathbf{V}^\top \Sigma_w^{\frac{1}{2}})^\top & \mathbf{I} \end{bmatrix} \succeq 0, \quad (14e)$$

where Λ is a diagonal matrix whose entries are the eigenvalues of Θ and \mathbf{V} is an orthogonal matrix satisfying $\Theta = \mathbf{V} \Lambda \mathbf{V}^\top$. Then, the pair $\{\Phi_x^*, \Phi_u^*\}$ of system responses, where $\Phi_x^* := \Psi_x^* \mathbf{V}^\top$ and $\Phi_u^* := \Psi_u^* \mathbf{V}^\top$, is a solution of problem (13).

Proof. Since Θ is a real symmetric matrix, it can always be diagonalized as $\Theta = \mathbf{V} \Lambda \mathbf{V}^\top$ with some orthogonal matrix \mathbf{V} . Since \mathbf{V} is orthogonal, we have that the transformation $\varphi(\Psi) = \Psi \mathbf{V}^\top$ is one-to-one, and that its inverse is given by $\varphi^{-1}(\Phi) = \Phi \mathbf{V}$. Under the map φ , the transformed problem (14) is equivalent to the original problem (13) in the sense that if $\{\Phi_x^*, \Phi_u^*\}$ solves problem (13), then $\{\Phi_x^* \mathbf{V}, \Phi_u^* \mathbf{V}\}$ solves problem (14), whereas if $\{\Psi_x^*, \Psi_u^*\}$ solves problem

(14), then $\{\Psi_x^* \mathbf{V}^\top, \Psi_u^* \mathbf{V}^\top\}$ solves problem (13) (see [25, Section 4.1.3]). Note that Ψ_x^* is invertible as both Φ_x^* and \mathbf{V} are invertible. It follows from the linearity of φ and the orthogonality of \mathbf{V} that

$$\mathbf{K}^* = \Psi_u^* \Psi_x^{*-1} = \Phi_u^* \mathbf{V} \mathbf{V}^\top \Phi_x^{*-1} = \Phi_u^* \Phi_x^{*-1}.$$

This proves the desired result. \square

Remark 4. Since Λ is diagonal, whether \mathbf{F} is diagonal or not, the objective function \tilde{f} in (14a) is column-wise separable with respect to \mathcal{C} : $\tilde{f}(\Psi_x, \Psi_u) = \sum_{i \in \mathcal{V}} \sum_{c \in \mathcal{C}^i} \Lambda(c, c) (\Psi_x(:, c)^\top \mathbf{Q} \Psi_x(:, c) + \Psi_u(:, c)^\top \mathbf{R} \Psi_u(:, c))$. It can be shown that the constraints (14b) and (14c) remain column-wise separable with respect to \mathcal{R}_x and \mathcal{R}_u , while (14d) remains row-wise separable with respect to \mathcal{S} . The separability of the LMI constraint (14e) still depends on whether it can be decomposed into local constraints, as described in Remark 2.

D. Distributed Solution using Consensus ADMM

To solve in a distributed manner the localized SLS problem (13), which as shown generally includes nonseparable components, we employ the consensus ADMM framework [23]. To ensure consensus across subsystems, we let each subsystem $i \in \mathcal{V}$ maintain its local copies of system responses $\{\Phi_x^{(i)}, \Phi_u^{(i)}\}$, while introducing global variables Φ_x^g and Φ_u^g representing the consensus system responses such that the consensus constraints are: $\Phi_x^{(i)} = \Phi_x^g$ and $\Phi_u^{(i)} = \Phi_u^g$, $\forall i \in \mathcal{V}$. The subsystem communication topology is assumed to be determined by the system graph \mathcal{G} and a locality parameter d . That is, each subsystem can only send information to its d -outgoing neighbors and receive information from its d -incoming neighbors. To ensure that all subsystems arrive at a consensus for all $d \geq 1$, we make the following assumption.

Assumption 1. The system graph \mathcal{G} is strongly connected.

Algorithm 1 Distributed and localized covariance steering for coupled systems using consensus ADMM

- 1: $\Phi_{x,0}^{(i)}, \Phi_{u,0}^{(i)}, \Omega_{x,0}^{(i)}, \Omega_{u,0}^{(i)}$, and $k \leftarrow 0$ \triangleright Initialization
- 2: **while** not converged **do**
- 3: **for** $i \in \mathcal{V}$ **do in parallel**
- 4: Send $\Phi_{x,k}^{(i)}, \Phi_{u,k}^{(i)}$ to neighbors in $\mathcal{O}_i(d)$
- 5: Receive $\Phi_{x,k}^{(j)}, \Phi_{u,k}^{(j)}$ from neighbors in $\mathcal{I}_i(d)$
- 6: $\Phi_{x,k+1}^{(i)}, \Phi_{u,k+1}^{(i)} \leftarrow$ Use (15) \triangleright Primal update
- 7: $\Omega_{x,k+1}^{(i)}, \Omega_{u,k+1}^{(i)} \leftarrow$ Use (16) \triangleright Dual update
- 8: **end for**
- 9: $k \leftarrow k + 1$
- 10: **end while**

The proposed consensus ADMM algorithm for the localized SLS problem (13) is detailed in Algorithm 1. First, the local system responses, $\Phi_{x,0}^{(i)}, \Phi_{u,0}^{(i)}$, are initialized to random matrices, while the local dual matrices, $\Omega_{x,0}^{(i)}, \Omega_{u,0}^{(i)}$, are set to zero matrices (line 1). Each subsystem i then sends its local system responses to its d -outgoing neighbors (line 4) and receives local system responses from its d -incoming neighbors (line 5). Each subsystem then performs primal and dual update procedures (lines 6 and 7), as specified below.

Primal Update: In the case where \mathbf{Q} and \mathbf{R} are diagonal, each subsystem i updates their local system responses $\Phi_x^{(i)}$

⁶Notice that $\mathbf{P}_T \Phi_x \mathbf{P}_0^\top$ is the first block of $\mathbf{P}_T \Phi_x$, or equivalently $\Phi_x^{T,0}$.

and $\Phi_u^{(i)}$ by solving the following convex SLS subproblem:

$$\begin{aligned} & \underset{\Phi_x^{(i)}, \Phi_u^{(i)}}{\text{minimize}} \quad f_i(\Phi_x^{(i)}, \Phi_u^{(i)}) \\ & \quad + \mathbf{1}^\top (\Omega_{x,k}^{(i)} \circ \Phi_x^{(i)} + \Omega_{u,k}^{(i)} \circ \Phi_u^{(i)}) \mathbf{1} \\ & \quad + \rho \sum_{j \in \mathcal{I}_i(d)} \left\| \Phi_x^{(i)} - \frac{1}{2} (\Phi_{x,k}^{(i)} + \Phi_{x,k}^{(j)}) \right\|_F^2 \\ & \quad + \rho \sum_{j \in \mathcal{I}_i(d)} \left\| \Phi_u^{(i)} - \frac{1}{2} (\Phi_{u,k}^{(i)} + \Phi_{u,k}^{(j)}) \right\|_F^2 \end{aligned} \quad (15a)$$

$$\text{subject to } \mathbf{Z}_{\mathbf{A}, \mathbf{B}} \begin{bmatrix} \Phi_x^{(i)}(:, \mathcal{C}^i) \\ \Phi_u^{(i)}(:, \mathcal{C}^i) \end{bmatrix} = \mathbf{I}(:, \mathcal{C}^i), \quad (15b)$$

$$\begin{bmatrix} \Phi_x^{(i)}(:, \mathcal{C}^i) \\ \Phi_u^{(i)}(:, \mathcal{C}^i) \end{bmatrix} \in \mathcal{L}_d(:, \mathcal{C}^i), \quad (15c)$$

$$[\mathbf{P}_T \Phi_x^{(i)} \mathbf{P}_0^\top] (\mathcal{S}^i, :) \mu_0 = \mu_f (\mathcal{S}^i, :), \quad (15d)$$

$$\begin{bmatrix} \Sigma_f & \mathbf{P}_T \Phi_x^{(i)} \Sigma_w^{\frac{1}{2}} \\ (\mathbf{P}_T \Phi_x^{(i)} \Sigma_w^{\frac{1}{2}})^\top & \mathbf{I} \end{bmatrix} \succeq 0, \quad (15e)$$

where $\rho > 0$ is a penalty parameter, and

$$\begin{aligned} f_i(\Phi_x, \Phi_u) &:= \sum_{r \in \mathcal{R}_x^i} \mathbf{Q}(r, r) \Phi_x(r, :) \Theta \Phi_x(r, :)^T \\ & \quad + \sum_{s \in \mathcal{R}_u^i} \mathbf{R}(s, s) \Phi_u(s, :) \Theta \Phi_u(s, :)^T. \end{aligned}$$

For the case of not necessarily diagonal \mathbf{F} , an analogous convex SLS subproblem can be constructed by decomposing the transformed problem (14), which is omitted here due to limited space.

Dual Update: Each subsystem i updates their dual variables according to: $\forall \bullet \in \{x, u\}$,

$$\Omega_{\bullet, k+1}^{(i)} = \Omega_{\bullet, k}^{(i)} + \rho \sum_{j \in \mathcal{I}_i(d)} \left(\Phi_{\bullet, k+1}^{(i)} - \Phi_{\bullet, k+1}^{(j)} \right). \quad (16)$$

Algorithm 1 keeps iterating until all N subsystems reach a consensus, i.e., when the average consensus constraints' residual norms for both system responses reaches a tolerance

$$\varepsilon > 0: \frac{1}{N} \sum_{i \in \mathcal{V}} \sum_{j \in \mathcal{I}_i(d)} \left\| \Phi_{\bullet}^{(i)} - \Phi_{\bullet}^{(j)} \right\|_F^2 \leq \varepsilon, \quad \forall \bullet \in \{x, u\}.$$

Proposition 2. Let $\mathcal{F}_i(d)$ denote the feasible set of the SLS subproblem (15) for subsystem i . If Assumption 1 holds and, for all $i \in \mathcal{V}$, Slater's condition [25] holds (i.e., the relative interior of $\mathcal{F}_i(d)$ is nonempty), then the sequence of local system responses $(\{\Phi_{x,k}^{(i)}, \Phi_{u,k}^{(i)}\})_{k \in \mathbb{N}}$ in Algorithm 1 satisfies, as $k \rightarrow \infty$, 1) $\left\| \Phi_{\bullet, k}^{(i)} - \Phi_{\bullet, k}^{(j)} \right\|_F^2 \rightarrow 0$, $\forall j \in \mathcal{I}_i(d)$ and $\forall \bullet \in \{x, u\}$, and 2) $f(\Phi_{x,k}^{(i)}, \Phi_{u,k}^{(i)}) \rightarrow p^*$, where p^* denotes the optimal value of the localized SLS problem (13).

Proof. Let h_i be the extended real-valued function defined by:

$$h_i(\Phi_x, \Phi_u) := \begin{cases} f_i(\Phi_x, \Phi_u), & \text{if } \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} \in \mathcal{F}_i(d), \\ +\infty, & \text{otherwise.} \end{cases}$$

Let $\text{epi } h_i := \{(\Phi_x, \Phi_u, \tau) : h_i(\Phi_x, \Phi_u) \leq \tau\}$ be the epigraph of h_i . From the definition of h_i , it follows that $\text{epi } h_i = \{(\Phi_x, \Phi_u, \tau) \in \mathcal{F}_i(d) \times \mathbb{R} : f_i(\Phi_x, \Phi_u) \leq \tau\}$. Since $\mathcal{F}_i(d)$ is a nonempty closed convex set and f_i is a real-valued convex function, $\text{epi } h_i$ is a nonempty closed convex

set, which is both necessary and sufficient for h_i to be a closed proper convex function [23, Section 3.2]. Now, consider the distributed form of the SLS problem (13):

$$\underset{\{\Phi_x^{(i)}, \Phi_u^{(i)}\}_{i \in \mathcal{V}}}{\text{minimize}} \quad \sum_{i \in \mathcal{V}} h_i(\Phi_x^{(i)}, \Phi_u^{(i)}) \quad (17a)$$

$$\text{subject to} \quad \Phi_x^{(i)} = \Phi_x^{(j)}, \quad \forall j \in \mathcal{I}_i(d), \quad \forall i \in \mathcal{V}, \quad (17b)$$

$$\Phi_u^{(i)} = \Phi_u^{(j)}, \quad \forall j \in \mathcal{I}_i(d), \quad \forall i \in \mathcal{V}. \quad (17c)$$

Clearly, problem (17) is equivalent to its centralized counterpart (13) with the same optimal solution and optimal value, as Assumption 1 ensures that $\Phi_{\bullet}^{(i)} = \Phi_{\bullet}^{(j)}$, $\forall \bullet \in \{x, u\}$, $\forall i, j \in \mathcal{V}$. Lastly, Slater's condition implies that strong duality is achieved. In light of all these results, the rest of the proof follows directly from [23, Appendix A]. \square

Remark 5. The computational complexity of Algorithm 1 is determined by the primal update procedure (line 6), in which each subsystem solves per iteration the SLS subproblem (15) over $O((n^2 + nm)T^2)$ variables subject to $O((n + m)T)$ constraints. While this complexity matches that of the centralized SLS problem (13) due to the use of the local copies of the global system responses, it is worth highlighting that solving problem (15) only requires each subsystem i to know, besides the global target covariance Σ_f , its assigned portion of the system matrices $\{\mathbf{A}, \mathbf{B}\}$, the cost matrices $\{\mathbf{Q}, \mathbf{R}\}$, and the target mean μ_f . Additionally, in the special instance of local terminal covariance constraints as described in Remark 2, this complexity can further be reduced.

IV. NUMERICAL SIMULATIONS

In this section, we verify the performance of Algorithm 1 on a power system modeled as a randomized spanning tree within a 6×6 grid (i.e., $N = 36$), as shown in Figure 1. Each subsystem is governed by the discretized swing equations: $m^i \ddot{\theta}^i + d^i \dot{\theta}^i = -\sum_{j \in \mathcal{N}_i} k^{ij} (\theta^i - \theta^j) + w^i + u^i$, where θ^i and $\dot{\theta}^i$ denote the phase angle and frequency deviations; m^i and d^i are the inertia and damping; w^i and u^i represent external disturbances and control inputs; k^{ij} is the coupling term between subsystems i and j ; and \mathcal{N}_i is the set of neighboring vertices. Defining the local state vector of each subsystem as $x^i := [\theta^i \quad \dot{\theta}^i]^\top$, the discretized swing dynamics can be expressed as

$$x_{t+1}^i = A^{ii} x_t^i + \sum_{j \in \mathcal{N}_i \setminus \{i\}} A^{ij} x_t^j + B^i u_t^i + w_t^i, \quad (18)$$

where $A^{ii} = \begin{bmatrix} 1 & \Delta t \\ -\frac{k^i}{m^i} \Delta t & 1 - \frac{d^i}{m^i} \Delta t \end{bmatrix}$, $A^{ij} = \begin{bmatrix} 0 & 0 \\ \frac{k^{ij}}{m^i} \Delta t & 0 \end{bmatrix}$, $B^i = [1 \quad 0]^\top$, $k^i = \sum_{j \in \mathcal{N}_i} k^{ij}$, and $\Delta t = 0.2$. The values of k^{ij} , d^i , and m^i are sampled uniformly at random from $[0.5, 1]$, $[0.2, 0.8]$, and $[0.5, 1]$, respectively. Note that the system graph \mathcal{G} coincides with the network topology displayed in Figure 1.

The localized CS problem (3) is initialized with $T = 10$, $Q = \text{diag}(100, 500, \dots, 100, 500)$, $R = 0.01I$, $W_t = 0.2I$, $\mu_f = 0$, and $\Sigma_f = MM^\top$ with the diagonal (resp., non-diagonal) elements of M sampled from $\mathcal{N}(0.5, 0.1)$ (resp., $\mathcal{N}(0, 0.1)$). The initial state is sampled as $x_0 \sim \mathcal{N}(\mu_0, \Sigma_0)$, where μ_0 follows the standard normal distribution scaled by 30, and Σ_0 is a diagonal matrix with entries sampled from a uniform distribution over $(0, 60]$. The locality parameter is chosen as $d = 1$, i.e., the communication graph of subsystems is also equivalent to the system graph \mathcal{G} . We execute Algorithm 1 using CLARABEL [26], an interior point conic solver available in CVXPY [27]. With $\rho = 0.01$ and $\varepsilon = 10^{-4}$, the algorithm converges within approximately 1,000 iterations.




Fig. 1: Network topology of the 6×6 power grid.

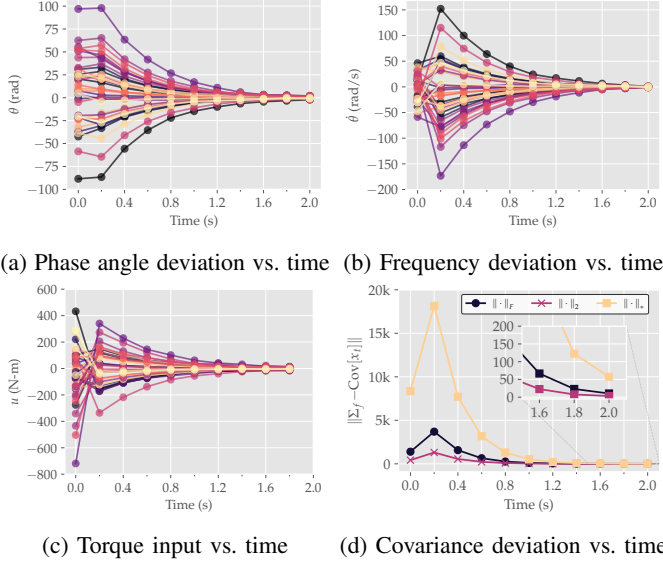


Fig. 2: Results for the simulation example with $d = 1$.

Figure (2) displays results for the numerical simulation. Figures 2a and 2b display each subsystem's phase angle and frequency deviations over the time horizon, respectively. It can be seen that both state deviations converge to 0 at the final time. Figure 2c shows the torque input at each time step. Figure 2d shows the difference between the final covariance and the state covariance at each time step, evaluated using three different norms⁷. In all cases, the state covariance at the final stage either converges to or closely approximates the target covariance.

V. CONCLUSION

This work addressed the problem of localized covariance steering in coupled stochastic linear systems. We employed the system-level synthesis framework to transform this problem into a convex optimization problem over system responses. A consensus-based algorithm was then proposed to distribute computation among subsystems. The effectiveness of this approach was demonstrated on a power system. In future work, we plan to address how clustering agents can accelerate the convergence of the proposed algorithm, as in [28], and to investigate potential data-driven extensions of distributed and localized covariance steering.

⁷Here, $\|\cdot\|_2$ and $\|\cdot\|_*$ denote the spectral norm and the nuclear (or trace) norm of a matrix, respectively.

ACKNOWLEDGMENT

The first author would like to thank Dr. Takashi Tanaka for the insightful discussions that have led to this work. The authors would also like to thank the editors and anonymous reviewers for their constructive reviews and recommendations for potential extensions.

REFERENCES

- [1] R. Stengel. *Optimal control and estimation*. Courier Corporation, 1994.
- [2] E. Collins and R. Skelton. Covariance control discrete systems. In *1985 IEEE Conf. Decis. Contr. (CDC)*, pages 542–547, 1985.
- [3] A. Hotz and R. Skelton. Covariance Control Theory. *Int. J. Control.*, 46(1):13–32, 1987.
- [4] Y. Chen, T. Georgiou, and M. Pavon. Optimal steering of a linear stochastic system to a final probability distribution, Part I. *IEEE Trans. Automat. Contr.*, 61(5):1158–1169, 2015.
- [5] E. Bakolas. Finite-horizon covariance control for discrete-time stochastic linear systems subject to input constraints. *Automatica*, 91:61–68, 2018.
- [6] I. Balci, E. Bakolas, B. Vlahov, and E. Theodorou. Constrained Covariance Steering Based Tube-MPPI. In *2022 Amer. Contr. Conf. (ACC)*, pages 4197–4202, 2022.
- [7] I. Balci and E. Bakolas. Covariance control of discrete-time gaussian linear systems using affine disturbance feedback control policies. In *2021 IEEE Conf. Decis. Contr. (CDC)*, pages 2324–2329, 2021.
- [8] I. Balci and E. Bakolas. Covariance Steering of Discrete-Time Stochastic Linear Systems Based on Wasserstein Distance Terminal Cost. *IEEE Contr. Syst. Lett.*, 5(6):2000–2005, 2021.
- [9] K. Okamoto, M. Goldshtein, and P. Tsiotras. Optimal Covariance Control for Stochastic Systems Under Chance Constraints. *IEEE Contr. Syst. Lett.*, 2(2):266–271, 2018.
- [10] I. Balci and E. Bakolas. Constrained minimum variance and covariance steering based on affine disturbance feedback control parameterization. *Int. J. Robust Nonlinear Control*, 34(11):7332–7370, 2024.
- [11] K. Okamoto and P. Tsiotras. Optimal Stochastic Vehicle Path Planning Using Covariance Steering. *IEEE Robot. Autom. Lett.*, 4(3):2276–2281, 2019.
- [12] J. Knaup, K. Okamoto, and P. Tsiotras. Safe High-Performance Autonomous Off-Road Driving Using Covariance Steering Stochastic Model Predictive Control. *IEEE Trans. Control Syst. Technol.*, 31(5):2066–2081, 2023.
- [13] J. Lee, E. Bakolas, and L. Sentis. Hierarchical Task-Space Optimal Covariance Control With Chance Constraints. *IEEE Contr. Syst. Lett.*, 6:2359–2364, 2022.
- [14] F. Liu, G. Rapakoulias, and P. Tsiotras. Optimal covariance steering for discrete-time linear stochastic systems. *IEEE Trans. Automat. Contr.*, pages 1–16, 2024.
- [15] J. Skaf and S. Boyd. Design of Affine Controllers via Convex Optimization. *IEEE Trans. Automat. Contr.*, 55(11):2476–2487, 2010.
- [16] A. Saravanos, A. Tzolovikos, E. Bakolas, and E. Theodorou. Distributed Covariance Steering with Consensus ADMM for Stochastic Multi-Agent Systems. In *2021 Robot.: Sci. Syst. (RSS)*, Virtual, July 2021.
- [17] A. Saravanos, I. Balci, E. Bakolas, and E. Theodorou. Distributed Model Predictive Covariance Steering. In *2024 IEEE Int. Conf. Intell. Robots Syst. (IROS)*, pages 5740–5747, 2024.
- [18] A. Saravanos, Y. Li, and E. Theodorou. Distributed Hierarchical Distribution Control for Very-Large-Scale Clustered Multi-Agent Systems. In *2023 Robot.: Sci. Syst. (RSS)*, Daegu, Republic of Korea, July 2023.
- [19] J. Anderson, J. Doyle, S. Low, and N. Matni. System level synthesis. *Annu. Rev. Control*, 47:364–393, 2019.
- [20] Y. Wang, N. Matni, and J. Doyle. Separable and localized system-level synthesis for large-scale systems. *IEEE Trans. Automat. Contr.*, 63(12):4234–4249, 2018.
- [21] C. Alonso and N. Matni. Distributed and localized closed loop model predictive control via system level synthesis. In *2020 IEEE Conf. Decis. Contr. (CDC)*, pages 5598–5605, 2020.
- [22] C. Alonso, J. Li, J. Anderson, and N. Matni. Distributed and Localized Model-Predictive Control—Part I: Synthesis and Implementation. *IEEE Trans. Control Netw. Syst.*, 10(2):1058–1068, 2023.
- [23] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3(1):1–122, 2011.
- [24] Y. Wang, N. Matni, and J. Doyle. Localized LQR control with actuator regularization. In *2016 Amer. Contr. Conf. (ACC)*, pages 5205–5212, 2016.
- [25] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, March 2004.
- [26] P. Goulart and Y. Chen. Clarabel: An Interior-Point Solver for Conic Programs with Quadratic Objectives. *arXiv preprint arXiv:2405.12762*, 2024.
- [27] S. Diamond and S. Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *J. Mach. Learn. Res.*, 17(83):1–5, 2016.
- [28] Y. Lee, A. Khalil, E. Bakolas, and G. Gremillion. A two-phase algorithm for joint task assignment and coordination in hierarchical multi-vehicle systems. In *AIAA SCITECH 2025 Forum*, 2025.