

Regret-Optimal Defense Against Stealthy Adversaries: A System Level Approach

Hiroyasu Tsukamoto^{*†§}, Joudi Hajar^{*‡}, Soon-Jo Chung^{‡§}, and Fred Y. Hadaegh^{‡§}

Abstract—Modern control designs in robotics, aerospace, and cyber-physical systems rely heavily on real-world data obtained through system outputs. However, these outputs can be compromised by system faults and malicious attacks, distorting critical system information needed for secure and reliable operation. In this paper, we introduce a novel regret-optimal control framework for designing controllers that make a linear system robust against stealthy attacks, including both sensor and actuator attacks. Specifically, we present (a) a convex optimization-based system metric to quantify the regret under the worst-case stealthy attack (the difference between actual performance and optimal performance with hindsight of the attack), which adapts and improves upon the \mathcal{H}_2 and \mathcal{H}_∞ norms in the presence of stealthy adversaries, (b) an optimization problem for minimizing the regret of (a) in system-level parameterization, enabling localized and distributed implementation in large-scale systems, and (c) a rank-constrained optimization problem equivalent to the optimization of (b), which can be solved using convex rank minimization methods. We also present numerical simulations that demonstrate the effectiveness of our proposed framework.

I. INTRODUCTION

Recent advances in autonomous control, often enhanced by machine learning, highlight the critical role of real-world data in building resilient autonomous systems and their decision-making processes. When using such data in robotics, aerospace, and cyber-physical systems, a critical question arises: “What if the system outputs are compromised and fail to capture important real-world data, potentially indicating a system anomaly?”

This paper addresses this issue by examining discrete-time linear time-varying systems affected by stealthy actuator and sensor disturbances. We define a disturbance as stealthy when the difference between the system outputs, with and without disturbances, stays below a certain threshold. These disturbances can lead to the loss of crucial adversarial information in real-world data, occurring in scenarios such as (1) malicious sensor and actuator attacks, (2) system faults like actuator or sensor failures, and (3) disturbances or uncertainties masked by the system’s output. Developing a framework to detect and optimally mitigate these stealthy

disturbances is crucial for tackling problems in cybersecurity, Fault Detection, Identification, and Recovery (FDIR), thereby ensuring reliable autonomy in robotics and aerospace systems.

Contributions: We propose a novel rank minimization framework to derive the regret-optimal system-level parameterization, specifically designed to enhance resilience against stealthy actuator and sensor disturbances. The regret in the presence of adversarial disturbances is defined as $\text{Regret} = J_t - J_c$, where J_t represents the performance using only causal information, and J_c represents the performance with non-causal information [1]–[3], i.e.,

J_t = true performance with access only to past attacks
 J_c = clairvoyant performance with access to past, present, and future attacks.

Note that the resultant regret-optimal controller does not require any knowledge of future attacks, although the clairvoyant J_c depends on them: it measures the control performance based on what it could have done if it had known about future attacks. Our method further leverages the System-Level Synthesis (SLS) parameterization introduced in [4]–[7], which provides a necessary and sufficient characterization of the achievable system response for the entire closed-loop system. It reframes the conventional control synthesis problem as designing the entire closed-loop system, enabling the integration of a wide range of system-level constraints, such as distributed constraints, localizability, and robustness, within a convex formulation [6]. We achieve the following contributions by combining the approaches of regret optimality and SLS in systems with stealthy adversaries.

- (a) We introduce a new convex optimization-based system metric that quantifies regret under the worst-case stealthy attack. This metric adjusts between the average and robust performances of \mathcal{H}_2 and \mathcal{H}_∞ norms in the presence of stealthy adversaries.
- (b) We formulate an optimization problem for minimizing the regret of (a) based on system-level parameterization, enabling scalable implementation in distributed and localized control systems.
- (c) We reformulate the optimization problem of (b) as a rank-constrained problem, featuring a convex objective with convex constraints and rank restrictions.

The optimization problem in the third contribution is convex except for the rank constraints, which can be addressed using established convex rank minimization techniques [8]–[12].

* These authors contributed equally to this work.

† Department of Aerospace Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, hiroyasu@illinois.edu.

‡ Division of Engineering and Applied Science, Caltech, Pasadena, CA, {jhajar, sjchung, hadaegh}@caltech.edu.

§ Jet Propulsion Laboratory, Pasadena, CA.

This project is funded by the Technology Innovation Institution (TII) with a contract with Caltech. The first author is funded by the University of Illinois at Urbana-Champaign. Part of the research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Related Work: Regret-optimal control [1]–[3] is an adaptive framework that minimizes the performance gap, or regret, between a causal controller and an ideal clairvoyant (non-causal) controller. Building on concepts from online learning and adaptive control [13]–[19], it dynamically adjusts control inputs based on real-time observations, aiming to minimize worst-case regret. While inspired by the minimax approach of \mathcal{H}_∞ control [20], regret-optimal control is less conservative due to its consideration of clairvoyant information, striking a balance between the robust performance of \mathcal{H}_∞ and the average performance of \mathcal{H}_2 [7], [21], [22].

Recent work [7] shows that dynamic regret-optimal controllers can be synthesized using system-level parameterization [4]–[6]. This method enables the design of the entire closed-loop system response, expanding the range of controllers that can be implemented through convex optimization.

This paper extends these methods to handle more challenging scenarios involving stealthy adversaries (see [23], [24] and [25]–[29] for details on stealthy adversaries). In particular, the security metric we derive generalizes the output-to-output ℓ_2 gain [25], [26] by incorporating regret-optimal control and system-level parameterization. This approach retains the adaptive and robust defense capabilities, making it well-suited for large-scale linear systems facing stealthy attacks. In Sec. VI, we propose ways to enhance this approach by incorporating the useful properties of system-level parameterization and regret-optimal control, such as sparsity and localizability, within an online learning framework. This fact demonstrates the strong potential for data-driven applications [30]–[33].

II. PRELIMINARIES

For $A \in \mathbb{R}^{n \times n}$, we use $A > 0$, $A \geq 0$, $A < 0$, and $A \leq 0$ to denote its definiteness. For $x \in \mathbb{R}^n$, $\|x\|$ represents the Euclidean norm. We also define blkdiag , rank , and vec as the block diagonal, matrix rank, and vectorization functions, respectively. Lastly, \mathbb{O} and \mathbb{I} represent the zero and identity matrices, respectively. In this paper, we study the following discrete-time, time-varying, networked linear control system:

$$x_{k+1} = A_k x_k + B_{u,k} u(y_k, v_k) + B_{a,k} a_k \quad (1a)$$

$$y_k = C_{y,k} x_k + D_{ya,k} a_k \quad (1b)$$

$$z_k = C_{z,k} x_k + D_{zu,k} u_k \quad (1c)$$

where v_k represents the network topology (e.g., adjacency matrix), x_k is the system state, u is the control input, a_k is the attack input where $B_{a,k} a_k$ represents the actuator attack and $D_{ya,k} a_k$ accounts for the sensor attack, y_k is the system measurement, and z_k is the regulated output used to evaluate control performance.

Operator Formalism: As in [6], we use the following notations for the system signals: $\mathbf{x} = [x_0^\top, \dots, x_T^\top]^\top$, $\mathbf{u} = [u_0^\top, \dots, u_T^\top]^\top$, $\mathbf{a} = [a_0^\top, a_0^\top, \dots, a_{T-1}^\top]^\top$, $\mathbf{y} = [y_0^\top, \dots, y_T^\top]^\top$, and $\mathbf{z} = [z_0^\top, \dots, z_T^\top]^\top$, where $u_k = u(y_k, v_k)$ and T is the time horizon. We also use the following notations for the system matrices: $\mathcal{A} = \text{blkdiag}(A_0, \dots, A_{T-1}, \mathbb{O})$, $\mathcal{B}_u = \text{blkdiag}(B_{u,0}, \dots, B_{u,T-1}, \mathbb{O})$, $\mathcal{B}_a = \text{blkdiag}(\mathbb{I}, B_{a,0}, \dots, B_{a,T-1})$, $C_p = \text{blkdiag}(C_{p,0}, \dots, C_{p,T})$, and $\mathcal{D}_{pq} = \text{blkdiag}($

$D_{pq,0}, \dots, D_{pq,T}$), where $q \in \{u, a\}$, $p \in \{y, z\}$, and blkdiag denotes the block diagonalization. Then (1) can be written as

$$\mathbf{x} = \mathcal{Z} \mathcal{A} \mathbf{x} + \mathcal{Z} \mathcal{B}_u \mathbf{u} + \mathcal{B}_a \mathbf{a} \quad (2a)$$

$$\mathbf{y} = C_y \mathbf{x} + \mathcal{D}_{ya} \mathbf{a} \quad (2b)$$

$$\mathbf{z} = C_z \mathbf{x} + \mathcal{D}_{zu} \mathbf{u} \quad (2c)$$

where \mathcal{Z} is the block-downshift operator, i.e., a matrix with identity matrices along its first block sub-diagonal and zeros elsewhere.

Control: We adopt the following output feedback control, parameterized by the network topology $\mathbf{v} = (v_0, \dots, v_T)$:

$$\mathbf{u} = \mathbf{K}(\mathbf{v}) \mathbf{y} \quad (3)$$

This framework includes state estimation-based feedback control [5], defined by $\xi_{k+1} = A_{e,k} \xi_k + B_{e,k} y_k$ and $u_k = C_{e,k} \xi_k + D_{e,k} y_k$, where ξ_k represents the internal state of the controller.

A. System Level Parameterization

Applying the feedback control from (3), we get the following system dynamics:

$$\begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} \mathbf{R} \mathcal{B}_a + \mathbf{N} \mathcal{D}_{ya} \\ \mathbf{M} \mathcal{B}_a + \mathbf{L} \mathcal{D}_{ya} \end{bmatrix} \mathbf{a} = \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} \mathbf{a} \quad (4a)$$

$$\mathbf{R} = (\mathbb{I} - \mathcal{Z} \mathcal{A} - \mathcal{Z} \mathcal{B}_u \mathbf{K} C_y)^{-1} \quad (4b)$$

$$\mathbf{N} = \mathbf{R} \mathcal{Z} \mathcal{B}_u \mathbf{K} \quad (4c)$$

$$\mathbf{M} = \mathbf{K} C_y \mathbf{R} \quad (4d)$$

$$\mathbf{L} = \mathbf{K} + \mathbf{K} C_y \mathbf{R} \mathcal{Z} \mathcal{B}_u \mathbf{K} \quad (4e)$$

These equations satisfy the following affine relations:

$$\begin{bmatrix} \mathbb{I} - \mathcal{Z} \mathcal{A} & -\mathcal{Z} \mathcal{B}_u \end{bmatrix} \begin{bmatrix} \mathbf{R} & \mathbf{N} \\ \mathbf{M} & \mathbf{L} \end{bmatrix} = \begin{bmatrix} \mathbb{I} & \mathbb{O} \end{bmatrix}$$

$$\begin{bmatrix} \mathbf{R} & \mathbf{N} \\ \mathbf{M} & \mathbf{L} \end{bmatrix} \begin{bmatrix} \mathbb{I} - \mathcal{Z} \mathcal{A} \\ -C_y \end{bmatrix} = \begin{bmatrix} \mathbb{I} \\ \mathbb{O} \end{bmatrix}.$$

ensuring both necessary and *sufficient* conditions for the existence of an output feedback controller. This formulation allows us to leverage the system-level parameterization of (4) for control design [4]–[6].

Lemma 1. *The affine subspace defined by (5) parameterizes all possible system responses (4) achievable by the output feedback control (3). Furthermore, for any matrices $(\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L})$ satisfying (5), the feedback control (3) with $\mathbf{K} = \mathbf{L} - \mathbf{M} \mathbf{R}^{-1} \mathbf{N}$ achieves the desired response (4).*

Proof. The first part can be shown just by verifying (5) for $(\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L})$ of (4). The second part follows from substituting $\mathbf{K} = \mathbf{L} - \mathbf{M} \mathbf{R}^{-1} \mathbf{N}$ into (2) with the conditions of (5), which results in (4). See [4]–[6] for details. \square

Thanks to Lemma 1, we can directly work with the system level parameterization (4) and the affine conditions to look for a controller (3) with a desired property. It also enables localized and distributed control implementation, which is suitable for large-scale networked control systems in this paper.

B. Regret-Optimal Control

Let $J(\mathbf{u}, \mathbf{a})$ be some performance objective to be minimized for the control and attack input defined in (2). The dynamic regret of J , which measures the gap between the actual optimal performance and the optimal performance in hindsight, is defined as:

$$\text{Regret}_J(\mathbf{u}, \mathbf{a}) = J(\mathbf{u}, \mathbf{a}) - J^*(\mathbf{a}) \quad (6a)$$

$$J^*(\mathbf{a}) = J(\mathbf{u}_{nc}(\mathbf{a}), \mathbf{a}) \quad (6b)$$

where \mathbf{u}_{nc} represents the optimal non-causal control policy with access to past and future attacks \mathbf{a} .

Regret-optimal control seeks to minimize the worst-case regret, denoted $\text{Regret}_J^*(\mathbf{u})$ [1]–[3], as:

$$\mathbf{u}^* = \arg \min_{\mathbf{u}} \text{Regret}_J^*(\mathbf{u}) \quad (7a)$$

$$\text{Regret}_J^*(\mathbf{u}) = \max_{\|\mathbf{a}\| \leq 1} \text{Regret}_J(\mathbf{u}, \mathbf{a}) \quad (7b)$$

This approach balances two extremes: the robustness of \mathcal{H}_∞ , which may be overly conservative by focusing on worst-case scenarios, and the efficiency of \mathcal{H}_2 , which can be too optimistic under adversarial disturbances [1], [3]. It is also applicable in the output feedback context considered in this work, as discussed in [21], [22], [34]. Additionally, recent results [7] demonstrate the compatibility of regret-optimal control with the system-level parameterization in Sec. II-A.

III. PROBLEM FORMULATION

We define a class of stealthy attacks called α -stealthy attacks, consistent with the approach in [26], which broadens the concept of stealthiness. When $\alpha = 0$, it describes conventional stealthy attacks where the detector's output is always zero, such as zero dynamics attacks [27], [28]. For $\alpha > 0$, the detector's output remains below a specified threshold.

Definition 1. For a given time horizon T , an attack \mathbf{a} of (2a) is α -stealthy in $k = 1, \dots, T$ if the following holds:

$$\exists \alpha \in \mathbb{R}_+ \text{ s.t. } \|\mathbf{y} - \mathbf{y}_n\|^2 \leq \alpha$$

where \mathbf{y} is given by (2b) and \mathbf{y}_n is by \mathbf{y} of (2b) with $\mathbf{a} = 0$.

We summarize the key challenges in designing a general defense strategy against α -stealthy network adversaries as follows:

- How can we adapt the regret-optimal control approach (7) to handle the α -stealthy attacks defined in Definition 1, instead of using $\|\mathbf{a}\| \leq 1$ (see Theorem 1)?
- How can we minimize the metric from (a) using system-level parameterization for scalable, distributed control in large networks (see Proposition 1)?
- How can we solve (b) for the network topology $\boldsymbol{\nu}$ in (3) in a computationally efficient way to optimally detect and mitigate α -stealthy attacks (see Theorem 2)?

IV. REGRET-OPTIMAL DEFENSE AGAINST STEALTHY ATTACKS

Building on (a), we formulate the following optimization problem to evaluate the worst-case regret under α -stealthy attacks, as defined in Definition 1:

$$\sigma(\mathbf{u}) = \max_{\mathbf{a}} \text{Regret}_J(\mathbf{u}, \mathbf{a}) \text{ s.t. } \mathbf{a} \text{ is } \alpha \text{ stealthy.} \quad (8)$$

This formulation enables the simultaneous consideration of both regret-based robustness and the stealth characteristics of the attack sequence a_k from (1a), analogous to the output-to-output gain approach used in \mathcal{H}_∞ robustness analysis [26].

A. Optimal Non-Causal Control

We define the performance objective $J(\mathbf{u}, \mathbf{a})$ from (6) as $J(\mathbf{u}, \mathbf{a}) = \|\mathbf{z}\|^2 + \|\mathbf{u}\|^2$, where \mathbf{z} is the regulated output for the control performance given in (2c), and \mathbf{u} is the control input. The non-causal optimal performance $\mathcal{J}^*(\mathbf{a})$ as in $J^*(\mathbf{a})$ of (6) (where \mathcal{J} is introduced to distinguish it from the performance objective J in Sec. II-B) can be expressed as follows:

$$\mathcal{J}^*(\mathbf{a}) = \min_{\mathbf{u}} \|\mathcal{E}\mathbf{u} + \mathcal{F}\mathcal{B}_a\mathbf{a}\|^2 + \|\mathbf{u}\|^2 \quad (9)$$

where $\mathcal{E} = C_z(\mathbb{I} - \mathcal{Z}\mathcal{A})^{-1}\mathcal{Z}\mathcal{B}_u + \mathcal{D}_{zu}$, $\mathcal{F} = C_z(\mathbb{I} - \mathcal{Z}\mathcal{A})^{-1}$. This problem is clearly convex and can be solved as [3], [7]

$$\mathcal{J}^*(\mathbf{a}) = \mathbf{a}^\top \mathcal{B}_a^\top \mathcal{F}^\top (\mathbb{I} + \mathcal{E}\mathcal{E}^\top)^{-1} \mathcal{F}\mathcal{B}_a\mathbf{a} = \mathbf{a}^\top \mathcal{Q}\mathbf{a} \quad (10)$$

where \mathcal{Q} is defined appropriately. The optimal non-causal control policy \mathbf{u}_{nc} of (6) is then given as

$$\mathbf{u}_{nc}(\mathbf{a}) = -(\mathbb{I} + \mathcal{E}^\top \mathcal{E})^{-1} \mathcal{E}^\top \mathcal{F}\mathcal{B}_a\mathbf{a}.$$

B. System Level Approach to Worst-Case Stealthy Attack

Let $\boldsymbol{\Omega} = (\mathbf{R}, \mathbf{N}, \mathbf{M}, \mathbf{L})$ for the system level parameterization (4). Then the constraint of (8) can be written as

$$\mathbf{a}^\top \boldsymbol{\Phi}(\boldsymbol{\Omega})^\top \boldsymbol{\Phi}(\boldsymbol{\Omega})\mathbf{a} \leq \alpha, \quad (11)$$

$$\begin{aligned} \boldsymbol{\Phi}(\boldsymbol{\Omega}) &= [C_y(\mathbf{R}\mathcal{B}_a + \mathbf{N}\mathcal{D}_{ya}) + \mathcal{D}_{ya}] = [C_y\phi_x + \mathcal{D}_{ya}] \\ &= [C_y \quad 0] \begin{bmatrix} \phi_x \\ \phi_u \end{bmatrix} + \mathcal{D}_{ya} \end{aligned}$$

by Definition 1. Also, $J(\mathbf{u}, \mathbf{a}) = \mathcal{J}(\boldsymbol{\Omega}, \mathbf{a}) = \|\mathbf{z}\|^2 + \|\mathbf{u}\|^2$ is given as

$$\mathcal{J}(\boldsymbol{\Omega}, \mathbf{a}) = \mathbf{a}^\top (\boldsymbol{\Psi}(\boldsymbol{\Omega})^\top \boldsymbol{\Psi}(\boldsymbol{\Omega}) + \boldsymbol{\Phi}_u(\boldsymbol{\Omega})^\top \boldsymbol{\Phi}_u(\boldsymbol{\Omega}))\mathbf{a} \quad (12)$$

$$\begin{aligned} \boldsymbol{\Psi}(\boldsymbol{\Omega}) &= [C_z\mathbf{R}\mathcal{B}_a + \mathcal{D}_{zu}\mathbf{M}\mathcal{B}_a + C_z\mathbf{N}\mathcal{D}_{ya} + \mathcal{D}_{zu}\mathbf{L}\mathcal{D}_{ya}] \\ &= [C_z\phi_x + \mathcal{D}_{zu}\phi_u] = [C_z \quad \mathcal{D}_{zu}] \begin{bmatrix} \phi_x \\ \phi_u \end{bmatrix} \end{aligned}$$

and $\boldsymbol{\Phi}_u(\boldsymbol{\Omega})$ is as in (4a). The metric in (8) can now be formulated as the following quadratically constrained quadratic program (QCQP), using the system-level parameterization from (4) and (5):

$$\zeta(\boldsymbol{\Omega}) = \max_{\mathbf{a}} \mathbf{a}^\top (\boldsymbol{\Psi}^\top \boldsymbol{\Psi} + \boldsymbol{\Phi}_u^\top \boldsymbol{\Phi}_u - \mathcal{Q})\mathbf{a} \text{ s.t. } \mathbf{a}^\top \boldsymbol{\Phi}^\top \boldsymbol{\Phi}\mathbf{a} \leq \alpha. \quad (13)$$

For simplicity, the argument $\boldsymbol{\Omega}$ for $\boldsymbol{\Phi}$, $\boldsymbol{\Psi}$, and $\boldsymbol{\Phi}_u$ is omitted here. The terms \mathcal{Q} , $\boldsymbol{\Phi}$, $\boldsymbol{\Psi}$, and $\boldsymbol{\Phi}_u$ are defined in (10), (11), (12), and (4a), respectively. QCQPs exhibit strong duality when Slater's condition holds [35, pp. 653-658], [7], leading to the following theorem.

Theorem 1. *The QCQP for the regret-optimal metric for α -stealthy attacks (13) is equivalent to the following convex optimization problem:*

$$\begin{aligned} \mu(\Omega) &= \min_{\lambda \geq 0} \lambda \alpha & (14) \\ \text{s.t. } & \lambda \Phi(\Omega)^\top \Phi(\Omega) - \Psi(\Omega)^\top \Psi(\Omega) - \Phi_u(\Omega)^\top \Phi_u(\Omega) + Q \geq 0 \end{aligned}$$

and thus $\mu(\Omega) = \varsigma(\Omega)$ for ς of (13).

Proof. The Lagrangian of the QCQP (13) is given as follows:

$$L(\mathbf{a}, \lambda) = \lambda(\mathbf{a}^\top \Phi^\top \Phi \mathbf{a} - \alpha) + \mathbf{a}^\top (-\Psi^\top \Psi - \Phi_u^\top \Phi_u + Q) \mathbf{a}$$

which gives the Lagrange dual function $g(\lambda)$ defined as

$$g(\lambda) = \min_{\mathbf{a}} L = \begin{cases} -\lambda \alpha & \lambda \Phi^\top \Phi - \Psi^\top \Psi - \Phi_u^\top \Phi_u + Q \geq 0 \\ -\infty & \text{otherwise} \end{cases}.$$

As a result, we get the following convex optimization problem:

$$\min_{\lambda \geq 0} \lambda \alpha \text{ s.t. } \lambda \Phi^\top \Phi - \Psi^\top \Psi - \Phi_u^\top \Phi_u + Q \geq 0.$$

Since the original QCQP (13) is strictly feasible for $\mathbf{a} = 0$ when $\alpha \in \mathbb{R}_+$, strong duality holds [35, pp. 653-658]. \square

The metric introduced in (14) provides a novel framework for quantifying worst-case regret in controllers, leveraging the system-level parameterization outlined in Sec.II-A. Similar to the output-to-output gain discussed in [26], this metric can be interpreted as a system gain from the α -stealthy attack \mathbf{a} in (2a) to the regret of the regulated performance output \mathbf{z} in (2c), thereby addressing problem (a). Theorem 1 introduces a more refined and flexible method for assessing system resilience against stealthy attacks through convex optimization, employing regret as an alternative to traditional system gains such as \mathcal{H}_∞ and \mathcal{H}_2 .

C. Regret-Optimal Defense via Topology Switching

Using the system-level parameterization in Sec.II-A, we now design a feedback control policy to minimize the metric μ from (14) for regret-optimal stealthy attack mitigation.

Proposition 1. *Consider the following optimization problem defined by the system-level parameterization $\Omega = (\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L})$ of (4):*

$$\mu^* = \min_{\Omega, \lambda} \lambda \alpha \quad (15a)$$

$$\text{s.t. } \lambda \Phi(\Omega)^\top \Phi(\Omega) - \Psi(\Omega)^\top \Psi(\Omega) - \Phi_u(\Omega)^\top \Phi_u(\Omega) + Q \geq 0 \quad (15b)$$

$$\begin{bmatrix} \mathbf{I} - \mathcal{Z}\mathcal{A} & -\mathcal{Z}\mathcal{B}_u \end{bmatrix} \begin{bmatrix} \mathbf{R} & \mathbf{N} \\ \mathbf{M} & \mathbf{L} \end{bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \end{bmatrix} \quad (15c)$$

$$\begin{bmatrix} \mathbf{R} & \mathbf{N} \\ \mathbf{M} & \mathbf{L} \end{bmatrix} \begin{bmatrix} \mathbf{I} - \mathcal{Z}\mathcal{A} \\ -\mathcal{C}_y \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix} \quad (15d)$$

$$\Omega \in \mathfrak{S} (= \text{set for system level constraints}), \lambda \geq 0 \quad (15e)$$

where Q , Φ , and Ψ are as given in (10), (11), and (12), respectively, and other terms are defined in (2). The output feedback control of (3) with $\mathbf{K} = \mathbf{K}^* = \mathbf{L}^* - \mathbf{M}^* \mathbf{R}^{*-1} \mathbf{N}^*$ achieves the desired response (4) and ensures regret-optimal

resilience against α -stealthy attacks (Definition 1). The optimal solution $\Omega^* = (\mathbf{R}^*, \mathbf{M}^*, \mathbf{N}^*, \mathbf{L}^*)$ guarantees regret-optimal performance.

In particular, for the dynamic regret $\text{Regret}_{\mathcal{J}}(\Omega, \mathbf{a}) = \mathcal{J}(\Omega, \mathbf{a}) - \mathcal{J}^*(\mathbf{a})$ where \mathcal{J}^* and \mathcal{J} are defined in (9) and (12), the optimal value μ^* of (15) satisfies

$$\text{Regret}_{\mathcal{J}}(\Omega^*, \mathbf{a}) \leq \mu^*, \forall \alpha\text{-stealthy attacks } \mathbf{a} \quad (16)$$

Moreover, for the worst-case α -stealthy attack $\mathbf{d}_{\text{wc}}(\Omega)$, which minimizes (13) (the optimal value given in Theorem 1), we have

$$\mu^* = \text{Regret}_{\mathcal{J}}(\Omega^*, \mathbf{a}^*) \leq \text{Regret}_{\mathcal{J}}(\Omega, \bar{\mathbf{a}}), \forall \Omega \text{ with (5)} \quad (17)$$

where $\mathbf{a}^* = \mathbf{d}_{\text{wc}}(\Omega^*)$ and $\bar{\mathbf{a}} = \mathbf{d}_{\text{wc}}(\Omega)$.

Proof. As discussed in Lemma 1, the conditions (15c) and (15d) ensure that the feedback controller with $\mathbf{K} = \mathbf{K}^* = \mathbf{L}^* - \mathbf{M}^* \mathbf{R}^{*-1} \mathbf{N}^*$ achieves the desired response (4). Since the problem (15) is formulated to find Ω that minimizes the metric μ of (14) in Theorem 1, the relations (16) and (17) naturally follow by construction of ς in (13). \square

Due to the affine nature of the system-level parameterization in Lemma 1, the system-level constraints \mathfrak{S} can accommodate a wide range of system properties, including temporal and spatial locality, distributed control, and scalability [4]–[6]. In particular, for networked control systems, the control gain can be adjusted by switching the topology as in (3), and distributed control constraints can be incorporated directly into (4) in affine terms based on the network topology. This leads to the following assumption.

Assumption 1. *Given the history of topologies of \mathbf{v} (4) and their adjacency matrices, we can define a convex constraint $\Omega \in \mathfrak{C}(\mathbf{v})$ for restricting Ω to be in a set consistent with the network topologies and the system's causality.*

However, despite using $\mathfrak{S} = \mathfrak{C}(\mathbf{v})$ from Assumption 1, the optimization problem (15) in Proposition 1 remains nonlinear. Before proceeding, we introduce the following lemma useful for convexification (see also, e.g., [8], [9]).

Lemma 2. *For any real column vector $a \in \mathbb{R}^b$ and real symmetric matrix $S \in \mathbb{S}^{b \times b}$, we have the following:*

$$\{(a, S) | S = aa^\top\} = \left\{ (a, S) \left| \begin{bmatrix} S & a \\ a^\top & 1 \end{bmatrix} \geq 0 \ \& \ \text{rank}(S) = 1 \right. \right\} \quad (18)$$

where $\text{rank}(S)$ denotes the matrix rank of S .

Proof. Let (LHS) = \mathfrak{L} and (RHS) = \mathfrak{R} .

($\mathfrak{L} \subset \mathfrak{R}$): Since $\text{rank}(S) = \text{rank}(aa^\top) = 1$ and $S = aa^\top \Rightarrow S \geq aa^\top$, applying Schur's complement lemma [35, pp. 650-651] shows $\mathfrak{L} \subset \mathfrak{R}$.

($\mathfrak{R} \subset \mathfrak{L}$): Since S is real symmetric, we can always write $S = V \text{diag}(\varpi_1, \dots, \varpi_b) V^\top$, where V is an orthogonal matrix of the orthonormal eigenvectors and ϖ_i are the real eigenvalues of S . Also, since $\text{rank}(S) = 1$, S has at most one non-zero eigenvalue ϖ , which implies $S = \varpi v v^\top$ for the corresponding unit eigenvector v . Defining a as $a = \sqrt{\varpi} v$ shows $\mathfrak{R} \subset \mathfrak{L}$. \square

The following theorem provides one way to reformulate the problem (15) of Proposition 1 into an almost equivalent optimization problem in a computationally efficient form.

Theorem 2. *Let Assumption 1 hold and let $\mathfrak{S} = \mathfrak{C}(\mathbf{v})$ as in (15). Consider the following optimization problem, where the objective and constraints are convex, except for the rank constraints in (19e):*

$$\bar{\mu}^* = \min_{\Omega, \lambda, \mathcal{X}, \Lambda} \lambda \alpha \quad (19a)$$

$$\text{s.t.} \begin{bmatrix} \lambda_{\text{inv}} \mathcal{Q} + \ell(\mathcal{X}) & \Psi(\Omega)^\top & \Phi_u(\Omega)^\top \\ \Psi(\Omega) & \lambda \mathbb{I} & \mathbb{O} \\ \Phi_u(\Omega) & \mathbb{O} & \lambda \mathbb{I} \end{bmatrix} \geq 0, \quad (15c), (15d) \quad (19b)$$

$$\begin{bmatrix} \mathcal{X} & \text{vec}(\Phi(\Omega)) \\ \text{vec}(\Phi(\Omega))^\top & 1 \end{bmatrix} \geq 0, \quad \begin{bmatrix} \Lambda & \lambda \\ \lambda^\top & 1 \end{bmatrix} \geq 0, \quad (19c)$$

$$\Omega \in \mathfrak{C}(\mathbf{v}), \quad \lambda = [\lambda, \lambda_{\text{inv}}]^\top, \quad \lambda > 0, \quad \Lambda_{12} = \Lambda_{21} = 1 \quad (19d)$$

$$\text{rank}(\mathcal{X}) = 1, \text{rank}(\Lambda) = 1 \quad (19e)$$

where ℓ is a linear function of the decision variable \mathcal{X} defined by $\ell(\mathcal{X}) = \Phi(\Omega)^\top \Phi(\Omega)$, and $\text{vec}(\Phi(\Omega))$ denotes the vectorization of the matrix $\Phi(\Omega)$. Other notations follow those in Proposition 1. If $\lambda^* \neq 0$ for the optimizer (λ^*, Ω^*) of (15), which is the case when the optimal regret $\mu^* = \text{Regret}_{\mathcal{J}}(\Omega^*, \mathbf{a}^*) > 0$ in (17), then we have

$$\bar{\mu}^* = \mu^*. \quad (20)$$

If $\lambda^* = 0$, which is the case when the optimal regret $\mu^* = \text{Regret}_{\mathcal{J}}(\Omega^*, \mathbf{a}^*) = 0$, then we have

$$\bar{\mu}^* \geq \mu^*. \quad (21)$$

Proof. If $\lambda^* \neq 0$ in (15), i.e., $\mu^* = \text{Regret}_{\mathcal{J}}(\Omega^*, \mathbf{a}^*) > 0$ in (17), the constraint (15b) can be equivalently expressed as

$$\begin{bmatrix} \lambda_{\text{inv}} \mathcal{Q} + \mathcal{P} & \Psi(\Omega)^\top & \Phi_u(\Omega)^\top \\ \Psi(\Omega) & \lambda \mathbb{I} & \mathbb{O} \\ \Phi_u(\Omega) & \mathbb{O} & \lambda \mathbb{I} \end{bmatrix} \geq 0, \quad (22)$$

$$\lambda \lambda_{\text{inv}} = 1, \quad \mathcal{P} = \Phi(\Omega)^\top \Phi(\Omega) \quad (23)$$

using Schur's complement lemma [35, pp. 650-651], which is nonlinear due to the equality constraints (23). Still, since the nonlinearity is at most quadratic/bilinear, these constraints can be written as a function of additional decision variables (i.e., liftings) for $\Lambda = \lambda \lambda^\top$ and $\mathcal{X} = \text{vec}(\Phi(\Omega)) \text{vec}(\Phi(\Omega))^\top$, where $\lambda = [\lambda, \lambda_{\text{inv}}]^\top$. This observation leads to the following equivalent constraints of (23) due to (18) of Lemma 2:

$$(23) \Leftrightarrow \begin{bmatrix} \mathcal{X} & \text{vec}(\Phi(\Omega)) \\ \text{vec}(\Phi(\Omega))^\top & 1 \end{bmatrix} \geq 0, \quad \begin{bmatrix} \Lambda & \lambda \\ \lambda^\top & 1 \end{bmatrix} \geq 0, \quad \text{rank}(\mathcal{X}) = 1, \text{rank}(\Lambda) = 1.$$

Rewriting the constraints (22) and (23) using these additional lifted variables completes the proof for (20).

The relation (21) can be immediately obtained by the fact that $\{\lambda \in \mathbb{R} | \lambda > 0\} \subset \{\lambda \in \mathbb{R} | \lambda \geq 0\}$. \square

Since the optimization problem derived in Theorem 2 is convex except for the rank constraint, we can apply various convex rank minimization techniques from the literature, including, but not limited to, [8]–[12].

D. Relaxed Regret-Optimal Control

As in the relaxed, suboptimal \mathcal{H}_∞ control (see, e.g., [36]), we can consider a suboptimal, bounded regret-optimal gain for stealthy attacks in Proposition 1, which results in a simpler optimization problem as follows.

Corollary 1. *If we consider a relaxed problem with a fixed $\lambda = \bar{\lambda} \in \mathbb{R}_+$ in (15) of Proposition 1, then the problem (19) simplifies to the following feasibility problem:*

$$\text{Find } (\Omega, \mathcal{X}) \quad (24a)$$

$$\text{s.t.} \begin{bmatrix} \mathcal{Q} + \lambda \ell(\mathcal{X}) & \Psi(\Omega)^\top & \Phi_u(\Omega)^\top \\ \Psi(\Omega) & \mathbb{I} & \mathbb{O} \\ \Phi_u(\Omega) & \mathbb{O} & \mathbb{I} \end{bmatrix} \geq 0, \quad (15c), (15d) \quad (24b)$$

$$\begin{bmatrix} \mathcal{X} & \text{vec}(\Phi(\Omega)) \\ \text{vec}(\Phi(\Omega))^\top & 1 \end{bmatrix} \geq 0, \quad \Omega \in \mathfrak{C}(\mathbf{v}), \quad (24c)$$

$$\text{rank}(\mathcal{X}) = 1 \quad (24d)$$

which is convex except for the rank condition (24d).

Proof. This follows from Theorem 2 by fixing $\lambda = \bar{\lambda}$. \square

Also, dropping the rank constraints in the problems (19) and (24) results in a convex optimization problem with a well-known relaxation approach.

Corollary 2. *Applying Shor's relaxation [37] to the problem (15) of Proposition 1 yields the problem (19) of Theorem 2 without the rank constraints (19e). The solution of this problem $\bar{\mu}_{\text{shor}}^*$ gives the lower bound of μ^* in (15a), i.e.,*

$$\bar{\mu}_{\text{shor}}^* \leq \mu^*.$$

Proof. This follows from (18) of Lemma 2, which indicates that dropping the rank constraint there results in Shor's relaxation. See also, e.g., [38, pp. 220-225], [8]. \square

V. NUMERICAL EXAMPLE

This section presents the results of numerical simulation to demonstrate the efficacy of our proposed method.

Two approaches are used: 1) Fix λ and solve the feasibility problem in Corollary 1 using iterative rank minimization [9], adjusting λ via linear search until a satisfactory sub-optimal solution is found. 2) Fix λ and solve the relaxed problem in Corollary 2, compute the worst-case regret cost (14) under stealthy attacks, and iterate with a linear search on λ until the cost (14) is sub-optimally minimized.

Using the above, we implement our proposed controller on a spring-damper system with two masses and a sampling time of $T_s = 0.5$ s, with the stealthiness measure set to $\alpha = 0.1$. We compare its performance to the classical \mathcal{H}_∞ robust controller, with the results shown in this section. A comparison with the typical regret-optimal controller, consistent with (25), is omitted for brevity, as its performance under stealthy attacks closely resembles that of \mathcal{H}_∞ . Our formulation also allows the implementation of controllers resilient to stealthy attacks using any quadratic cost function, such as \mathcal{H}_∞ cost function, along with stealthiness constraints, which would prioritize robustness over adaptability.




Fig. 1. (Left) Stealthiness of the attacks measured as $\|y - y_n\|^2 = a^T \Phi(\Omega)^T \Phi(\Omega) a$, for \mathcal{H}_∞ and our controller. Both attacks satisfy the constraint: $\|y - y_n\|^2 \leq 0.1$. (Right): Worst-case regret cost (13). Our controller achieves a 15 \times lower cost compared to \mathcal{H}_∞ .

The classical \mathcal{H}_∞ controller, which isn't designed specifically for stealthy attacks, solves the following problem:

$$\begin{aligned} \min_{\Omega \in \mathcal{C}(\mathbf{v})} \max_{\|a\| \leq 1} \mathcal{J}(\Omega, \mathbf{a}) \\ = \min_{\Omega, \lambda} \lambda \alpha \end{aligned} \quad (25a)$$

$$\text{s.t.} \quad \begin{bmatrix} \lambda \mathcal{I} & \Psi(\Omega)^T & \Phi_u(\Omega)^T \\ \Psi(\Omega) & \mathbf{I} & \mathbf{O} \\ \Phi_u(\Omega) & \mathbf{O} & \mathbf{I} \end{bmatrix} \geq 0, \quad (15c), (15d) \quad (25b)$$

$$\Omega \in \mathcal{C}(\mathbf{v}) \quad (25c)$$

where $\mathcal{J}(\Omega, \mathbf{a}) = \mathbf{a}^T (\Psi(\Omega)^T \Psi(\Omega) + \Phi_u(\Omega)^T \Phi_u(\Omega)) \mathbf{a}$ as in (12).

We compare the worst-case regret cost of \mathcal{H}_∞ and our controller under stealthy attacks, as quantified by (13), (14). Table I shows that our controller consistently outperforms \mathcal{H}_∞ as the horizon increases.

We also compute the worst-case stealthy attacks for each controller and apply them to our system (2). Figure 1 shows that these attacks remain α -stealthy ($\alpha = 0.1$), with our controller achieving a 15 \times lower cost than \mathcal{H}_∞ . Figure 2 illustrates that only our controller regulates the output near zero. For computational efficiency, we chose a horizon of $T = 5$. A model predictive control scheme, however, can be used to manage longer horizons in practice.

This example shows the potential of our system-level approach in detecting and mitigating stealthy attacks with regret-optimal strategies. Again, as will be detailed in the next section, the framework extends to a wider class of constrained controllers, including predictive control and online learning, while maintaining a convex formulation.

TABLE I

WORST-CASE REGRET UNDER STEALTHY ATTACKS AS IN (13) FOR DIFFERENT TIME HORIZONS ($T=2, T=5$).

	T=2	T=5
\mathcal{H}_∞ controller	41.38	1312.6
Proposed approach	10.19	87.10
Improvement Factor: $\frac{\mathcal{H}_\infty(\text{Cost})}{\text{Proposed approach}(\text{Cost})}$	4.02	15.07

VI. POTENTIAL EXTENSIONS

This section explores various extensions of the topology switching strategy from Proposition 1 and Theorem 2, highlighting the broad applicability of our regret-optimal metric




Fig. 2. Regulated output z (2c) under \mathcal{H}_∞ and our controller. For a horizon of $T = 5$, our controller manages to keep the regulated output near 0, while the regulated output under \mathcal{H}_∞ shows big oscillations.

in more general problem settings under stealthy attacks. The purpose of this section is just to emphasize and outline the potential research directions that this paper would open up.

A. Sparsity on System Responses

If the convex constraint $\Omega \in \mathcal{C}(\mathbf{v})$ of (19d) in Theorem 2 includes sparsity constraints due to, e.g., distributed communications of the network, we could use the following corollary for additional simplification.

Corollary 3. *Suppose that we have sparsity constraints for Ω , which leads to linear constraints of the form $\text{vec}(\Phi(\Omega))_i = 0$. Then all the entries in the i th column and i th row of the matrix \mathcal{X} of (19) in Theorem 2 are also equal to zero.*

Proof. This follows from the definition of \mathcal{X} introduced in the proof of Theorem 2. \square

Another potential direction for addressing sparsity is considering spatial regret (SpRegret) defined in [39] as the regret between the cost of a topology Ω which belongs to a sparse set $\mathcal{C}(\mathbf{v})$, and that of a topology $\hat{\Omega}$ which belongs to a denser set $\mathcal{C}(\hat{\mathbf{v}})$: $\text{SpRegret}(\Omega, \hat{\Omega}) = \max_{\|a\| \leq 1} J(\Omega, a) - J(\hat{\Omega}, a)$. Under no further constraints, ideally, the constraint set should satisfy a quadratic invariance condition to ensure the optimality of the control strategy. In cases where it is not, there exist methods that allow for the synthesis of the controller even when the set is not Quadratic Invariant (QI). As in [39], we could use the denser topology as the one that satisfies the QI condition.

We could further extend the notion of regret to spatiotemporal regret by defining it as the difference in costs of a sparse topology and of a denser topology that also knows the future attacks. In this case, the challenge we face is ensuring that the problem is well-posed (i.e., regret > 0) and finding the closest QI, noncausal sparse set.

B. Data-Driven Control

The system level parameterization is also useful for the case where we have access only to the system's input and output data $u_{[0, T_{\text{data}}-1]} = \{u_0, \dots, u_{T_{\text{data}}-1}\}$ and $y_{[0, T_{\text{data}}-1]} = \{y_0, \dots, y_{T_{\text{data}}-1}\}$ of the system (1) with $a_k = 0$. Let $\mathcal{H}(\sigma_{[0, j]})$ be the Hankel matrix of depth D , associated with the signal

$\sigma_{[0, T_{\text{data}}-1]}$ defined as follows:

$$\mathcal{H}(\sigma_{[0, T_{\text{data}}-1]}) = \begin{bmatrix} \sigma_0 & \cdots & \sigma_{T_{\text{data}}-D} \\ \vdots & \ddots & \vdots \\ \sigma_{D-1} & \cdots & \sigma_{T_{\text{data}}-1} \end{bmatrix} \quad (26)$$

where $D \leq T_{\text{data}}$.

Definition 2. *The signal $\sigma_{[0, T_{\text{data}}-1]}$ is persistently exciting of order D if the Hankel matrix $\mathcal{H}(\sigma_{[0, T_{\text{data}}-1]})$ (26) has full row rank.*

The Hankel matrix of (26) leads to the following lemma essential in the data-driven system level synthesis [40], [41].

Lemma 3. *Suppose that the input data $u_{[0, T_{\text{data}}-1]}$ is persistently exciting of order $n+D$ as in Definition 2, where n is the number of the system state x of (1a). Suppose also that the system (1a) is time-invariant and controllable. Then $(\bar{u}_{[0, D-1]}, \bar{y}_{[0, D-1]})$ is a valid input/output trajectory of (1) with $a_k = 0$ if and only if $\exists g \in \mathbb{R}^{T_{\text{data}}-D+1}$ s.t.*

$$\begin{bmatrix} u_{[0, T_{\text{data}}-1]} \\ y_{[0, T_{\text{data}}-1]} \end{bmatrix} = \begin{bmatrix} \mathcal{H}(u_{[0, T_{\text{data}}-1]}) \\ \mathcal{H}(y_{[0, T_{\text{data}}-1]}) \end{bmatrix} g. \quad (27)$$

Proof. See, e.g., [40], [41]. \square

It is shown in [30] that g of (27) in Lemma 3 can be characterized by the system level parameterization of (4) for the case of full-state feedback. Its robust properties allow for extending this idea to perturbed trajectory data, as also shown in [30]. These observations imply the great potential of our approach to model-free settings as in [30], [32], [33].

C. Additional Constraints

1) *Performance Constraints:* We can clearly see that the control input \mathbf{u} in (4) of Sec. II-A is linear in the system level parameters $\mathbf{\Omega}$; the constraint

$$\|\mathbf{u}\| = \|\mathbf{M}\mathbf{d}_x + \mathbf{D}\mathbf{d}_y\|_* \leq u_{\max}$$

where $u_{\max} \in \mathbb{R}_+$ is some constant, is always a convex constraint for any norm $*$. We can also consider other performance requirements as convex constraints on top of the problems (19) of Theorem 2 and (24) of Corollary 24, the examples of which are given in [4]–[6].

2) *Localizability:* For large-scale networks, we could apply spatiotemporal constraints to further reduce the computational burden thanks to the system level parameterization. In particular, if we assume that 1) the closed-loop disturbance responses have finite-impulse responses and that 2) the effects of disturbances are felt only in local neighborhoods of each agent, then we can decompose the problems (19) of Theorem 2 into small subproblems that can be solved locally [4]–[6]. This allows the model predictive control-like formulation of the system level synthesis [31].

3) *Controllability, Observability, and Sensitivity:* As discussed in [5], [42], [43], the localizability of Sec. VI-C.2 allows for verifying the observability and controllability conditions using a set of affine constraints, which can be readily added to the problem (19) in Theorem 2 while preserving its convex

structure. The minimum attack sensitivity (or detectability in the original paper [29]) could potentially be considered directly as a constraint in this problem using a similar approach. Alternatively, we could enforce the sensitivity constraint by using the \mathcal{H}_- index [25]–[27], [44], [45] as in [46], with the same lifting procedure of Theorem 2.

4) *Stability:* For infinite-horizon problems, we can still ensure the internal stability of our formulation in a convex manner as discussed in [4]–[6] and in [47] for adversarial settings. Even in our finite-horizon setting, we could still enforce the Lyapunov-type stability constraint as in [46] due to the affine nature of the system level parameterization as to be implied in the numerical example.

VII. CONCLUSION

This paper introduced a system-level, regret-optimal framework for designing linear systems that are resilient to stealthy attacks and disturbances. We equivalently reformulated the nonlinear problem of minimizing regret in the presence of stealthy adversaries as a rank-constrained optimization problem, which can be solved using convex rank minimization methods. We also explored extensions such as incorporating system-level parameterization for sparsity and localizability in online learning, highlighting the potential for predictive, data-driven applications in an interpretable framework.

REFERENCES

- [1] O. Sabag, G. Goel, S. Lale, and B. Hassibi, “Regret-optimal controller for the full-information problem,” in *Proc. Amer. Control Conf.*, 2021, pp. 4777–4782.
- [2] —, “Regret-optimal LQR control,” arXiv:2105.01244, 2023.
- [3] G. Goel and B. Hassibi, “Regret-optimal estimation and control,” *IEEE Trans. Autom. Control*, vol. 68, no. 5, pp. 3041–3053, 2023.
- [4] J. C. Doyle, N. Matni, Y.-S. Wang, J. Anderson, and S. Low, “System level synthesis: A tutorial,” in *IEEE CDC*, 2017, pp. 2856–2867.
- [5] Y.-S. Wang, N. Matni, and J. C. Doyle, “A system-level approach to controller synthesis,” *IEEE Trans. Autom. Control*, vol. 64, no. 10, pp. 4079–4093, 2019.
- [6] J. Anderson, J. C. Doyle, S. H. Low, and N. Matni, “System level synthesis,” *Annu. Rev. Control*, vol. 47, pp. 364–393, 2019.
- [7] A. Didier, J. Sieber, and M. N. Zeilinger, “A system level approach to regret optimal control,” *IEEE Control Syst. Lett.*, vol. 6, pp. 2792–2797, 2022.
- [8] S. Boyd and L. Vandenberghe, *Semidefinite Programming Relaxations of Non-Convex Problems in Control and Combinatorial Optimization*. Boston, MA: Springer US, 1997, pp. 279–287.
- [9] S. Ibaraki and M. Tomizuka, “Rank minimization approach for solving BMI problems with random search,” in *Proc. Amer. Control Conf.*, vol. 3, 2001, pp. 1870–1875.
- [10] M. Fazel, H. Hindi, and S. Boyd, “Rank minimization and applications in system theory,” in *Proc. Amer. Control Conf.*, vol. 4, 2004, pp. 3273–3278 vol.4.
- [11] S. Shalev-Shwartz, A. Gonen, and O. Shamir, “Large-scale convex minimization with a low-rank constraint,” in *28th Int. Conf. Mach. Learn.* Omnipress, 2011, p. 329–336.
- [12] C. Sun and R. Dai, “Rank-constrained optimization and its applications,” *Automatica*, vol. 82, pp. 128–136, 2017.
- [13] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu, “Regret bounds for robust adaptive control of the linear quadratic regulator,” in *Adv. Neural Inf. Process. Syst.*, vol. 31, 2018.
- [14] N. Agarwal, B. Bullins, E. Hazan, S. Kakade, and K. Singh, “Online control with adversarial disturbances,” in *36th Int. Conf. Mach. Learn.*, ser. Proc. Mach. Learn. Res., vol. 97, Jun. 2019, pp. 111–119.
- [15] G. Goel and A. Wierman, “An online algorithm for smoothed regression and LQR Control,” in *22nd Int. Conf. Artif. Intell. Statist.*, ser. Proc. Mach. Learn. Res., vol. 89. Apr., 16–18 Apr 2019, pp. 2504–2513.

- [16] E. Hazan, S. Kakade, and K. Singh, "The nonstochastic control problem," in *31st Int. Conf. Algorithmic Learn. Theory*, ser. Proc. Mach. Learn. Res., vol. 117, Feb. 2020, pp. 408–421.
- [17] N. M. Boffi, S. Tu, and J.-J. E. Slotine, "Regret bounds for adaptive nonlinear control," in *LADC*, ser. Proc. Mach. Learn. Res., vol. 144, Jun. 2021, pp. 471–483.
- [18] C. Yu, G. Shi, S.-J. Chung, Y. Yue, and A. Wierman, "The power of predictions in online control," in *Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 1994–2004.
- [19] Y. Li, X. Chen, and N. Li, "Online optimal control with linear dynamics and predictions: Algorithms and regret analysis," in *Adv. Neural Inf. Process. Syst.*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds., vol. 32, 2019.
- [20] T. Basar and P. Bernhard, *H_∞ Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*. Birkhauser, 1995.
- [21] G. Goel, "Regret-optimal control," Ph.D. dissertation, California Institute of Technology, 2022.
- [22] J. Hajar, O. Sabag, and B. Hassibi, "Regret-optimal control under partial observability," arXiv:2311.06433, 2023.
- [23] J. Yu, V. Gupta, and A. Wierman, "Online adversarial stabilization of unknown linear time-varying systems," in *IEEE CDC*, 2023, pp. 8320–8327.
- [24] J. Yu, D. Ho, and A. Wierman, "Online Adversarial Stabilization of Unknown Networked Systems," *ACM Meas. Anal. Comput. Syst.*, vol. 7, no. 1, Mar. 2023.
- [25] A. Teixeira, *Security Metrics for Control Systems*. Springer Int. Publishing, 2021, pp. 99–121.
- [26] A. Teixeira, H. Sandberg, and K. H. Johansson, "Strategic stealthy attacks: The output-to-output ℓ_2 -gain," in *IEEE CDC*, 2015, pp. 2582–2587.
- [27] A. Teixeira, I. Shames, H. Sandberg, and K. H. Johansson, "A secure control framework for resource-limited adversaries," *Automatica*, vol. 51, pp. 135–148, 2015.
- [28] A. Teixeira, D. Pérez, H. Sandberg, and K. H. Johansson, "Attack models and scenarios for networked control systems," in *Int. Conf. High Confidence Netw. Syst.*, 2012, p. 55–64.
- [29] F. Pasqualetti, F. Dörfler, and F. Bullo, "Attack detection and identification in cyber-physical systems," *IEEE Trans. Autom. Control*, vol. 58, no. 11, pp. 2715–2729, 2013.
- [30] A. Xue and N. Matni, "Data-driven system level synthesis," in *Proc. Conf. Learn. Dyn. Control*, ser. Proc. Mach. Learn. Res., vol. 144, PMLR, Jun. 2021, pp. 189–200.
- [31] C. A. Alonso and N. Matni, "Distributed and localized closed loop model predictive control via a system level synthesis," in *IEEE CDC*, 2020, pp. 5598–5605.
- [32] C. A. Alonso, F. Yang, and N. Matni, "Data-driven distributed and localized model predictive control," *IEEE Open J. Control Syst.*, vol. 1, pp. 29–40, 2022.
- [33] Y. Lian and C. N. Jones, "From system level synthesis to robust closed-loop data-enabled predictive control," in *IEEE CDC*, 2021, pp. 1478–1483.
- [34] G. Goel and B. Hassibi, "Regret-optimal measurement-feedback control," in *LADC*, ser. Proc. Mach. Learn. Res., vol. 144, Jun. 2021, pp. 1270–1280.
- [35] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, Mar. 2004.
- [36] A. J. Van der Schaft, " \mathcal{L}_2 -gain analysis of nonlinear systems and nonlinear state-feedback H_∞ control," *IEEE Trans. Autom. Control*, vol. 37, no. 6, pp. 770–784, Jun. 1992.
- [37] N. Z. Shor, "Quadratic optimization problems," *Sov. J. Circuits Syst. Sci.*, vol. 25, no. 6, pp. 1–11, 1987.
- [38] ———, *Algorithms for Constructing Optimal on Volume Ellipsoids and Semidefinite Programming*. Boston, MA: Springer US, 1998, pp. 169–225.
- [39] D. Martinelli, A. Martin, G. Ferrari-Trecate, and L. Furieri, "Closing the gap to quadratic invariance: a regret minimization approach to optimal distributed control," 2023.
- [40] J. Willems, I. Markovskiy, P. Rapisarda, and B. De Moor, "A note on persistency of excitation," in *IEEE CDC*, vol. 3, 2004, pp. 2630–2631.
- [41] H. J. van Waarde, C. De Persis, M. K. Camlibel, and P. Tesi, "Willems' fundamental lemma for state-space systems and its extension to multiple datasets," *IEEE Control Syst. Lett.*, vol. 4, no. 3, pp. 602–607, 2020.
- [42] Y.-S. Wang, N. Matni, and J. C. Doyle, "Localized LQR optimal control," in *IEEE CDC*, 2014, pp. 1661–1668.
- [43] Y.-S. Wang and N. Matni, "Localized LQG optimal control for large-scale systems," in *Proc. Amer. Control Conf.*, 2016, pp. 1954–1961.
- [44] J. L. Wang, G.-H. Yang, and J. Liu, "An LMI approach to \mathcal{H}_- index and mixed $\mathcal{H}_-/\mathcal{H}_\infty$ fault detection observer design," *Automatica*, vol. 43, no. 9, pp. 1656–1665, 2007.
- [45] J. Liu, J. L. Wang, and G.-H. Yang, "An LMI approach to minimum sensitivity analysis with application to fault detection," *Automatica*, vol. 41, no. 11, pp. 1995–2004, 2005.
- [46] H. Tsukamoto, J. Hajar, J. Ibrahim, J. Ragan, A. Rahmani, S.-J. Chung, and F. Y. Hadaegh, "Robust and optimal network topology switching for zero dynamics attack," in *IEEE CDC*, 2024.
- [47] O. Kjellqvist and J. Yu, "On infinite-horizon system level synthesis problems," in *IEEE CDC*, 2022, pp. 5238–5244.