

Closing the Gap to Quadratic Invariance: a Regret Minimization Approach to Optimal Distributed Control

Daniele Martinelli, Andrea Martin, Giancarlo Ferrari-Trecate, and Luca Furieri

Abstract—In this work, we focus on the design of optimal controllers that must comply with an information structure. State-of-the-art approaches do so based on the \mathcal{H}_2 or \mathcal{H}_∞ norm to minimize the expected or worst-case cost in the presence of stochastic or adversarial disturbances. Large-scale systems often experience a combination of stochastic and deterministic disruptions (e.g., sensor failures, environmental fluctuations) that spread across the system and are difficult to model precisely, leading to sub-optimal closed-loop behaviors. Hence, we propose improving performance for these scenarios by minimizing the regret with respect to an ideal policy that complies with less stringent sensor-information constraints. This endows our controller with the ability to approach the improved behavior of a more informed policy, which would detect and counteract heterogeneous and localized disturbances more promptly. Specifically, we derive convex relaxations of the resulting regret minimization problem that are compatible with any desired controller sparsity, while we reveal a renewed role of the Quadratic Invariance (QI) condition in designing informative benchmarks to measure regret. Last, we validate our proposed method through numerical simulations on controlling a multi-agent distributed system, comparing its performance with traditional \mathcal{H}_2 and \mathcal{H}_∞ policies.

I. INTRODUCTION

Control of large-scale systems, such as smart grids [1] or traffic systems [2], requires communication among multiple interacting agents to ensure efficient and safe operation. A significant challenge arises from the incomplete information available to each agent regarding the overall system state. This partial communication can be due to various factors, including privacy concerns, geographic dispersion, and the inherent difficulties of establishing a reliable communication network. Designing optimal control policies complying with specific information structures is a well-known challenge, even in seemingly straightforward scenarios, as highlighted by the classic work [3].

The design phase concerning large-scale systems primarily revolves around addressing two fundamental challenges: first, how to parameterize controllers to comply with a given information sparsity, and second, how to formulate a metric that accounts for disturbances and the effect of their propagation throughout the dynamics of the distributed system.

Regarding the first challenge, for linear dynamical systems, a landmark contribution was given in [4], where the

D. Martinelli, A. Martin, G. Ferrari-Trecate, and L. Furieri are with the Institute of Mechanical Engineering, EPFL, Switzerland. E-mail addresses: {daniele.martinelli, andrea.martin, giancarlo.ferraritrecate, luca.furieri}@epfl.ch.

Daniele Martinelli and Luca Furieri are grateful to the Swiss National Science Foundation (SNSF) for the Ambizione grant PZ00P2.208951. This research is also supported by the SNSF under the NCCR Automation (grant agreement 51NF40.80545).

authors proposed the notion of Quadratic Invariance (QI), a sufficient [4] and necessary [5] condition for enabling an exact convex reformulation of sparse linear controller synthesis. However, the QI condition can be overly restrictive, i.e., not fulfilled by several systems with highly intertwined dynamics. To address this issue, [6] introduced communication channels between controllers to effectively restore the QI condition. The works [7]–[9] presented convex optimization methods for designing sparse closed-loop dynamics and sparse controllers, even when the QI condition does not hold. Within optimal distributed control, [10] addressed how to promote sparsity in closed-loop controllers while minimizing performance degradation compared to a centralized policy. Building on this, works [11] and [12] proposed methods for the explicit parameterization of sparse \mathcal{H}_2 controllers. [11] focused on physically decoupled systems, while [12] derived control policies that closely resemble an optimal centralized one under specific conditions.

Concerning the second challenge of choosing an appropriate performance metric, traditional control techniques, such as \mathcal{H}_2 and \mathcal{H}_∞ , often rely on certain assumptions about the nature of disturbances to achieve optimal control policies [13]. \mathcal{H}_2 treats disturbances as stochastic noise, while \mathcal{H}_∞ considers them as adversarial attacks. However, these assumptions can be unrealistic and conservative in large-scale scenarios, where disturbances are 1) difficult to model due to the complex dynamics of small local mismatches propagating at large, and 2) difficult to localize, as they may affect agents at unpredictable locations and times. This requires the development of novel control strategies capable of accommodating such non-standard disturbances.

Contributions. To improve the ability of the optimal policy to respond quickly to disturbances with unknown locations and nature, our idea is to close the gap to the performance of an *oracle*, that is, a benchmark policy that possesses more sensor measurements. By minimizing the worst-case difference in cost with this oracle, our controller design encourages emulation of its behavior, potentially leading to improved performance. This new metric is inspired by the recent works on regret technique [14]–[16], which are however limited to centralized control scenarios and focus on a temporal notion of regret.

In this work, we first analyze the conditions for the oracle to express an improved performance to be emulated. We term a spatial regret metric satisfying these capabilities as *well-posed*. This analysis is achieved by revealing a renewed role of the QI condition within the well-posedness of the proposed metric. Finally, we provide a convex reformulation

for designing regret-optimal controllers with arbitrary sparsity structures, optimizing them to close the performance gap to an ideal QI subspace that encodes richer information for the control policy. To illustrate the real-world relevance of our approach, we provide numerical examples involving a multi-agent scenario of a multi-mass spring-damper system that requires to be controlled in the presence of non-standard disturbances.

Structure. The paper is organized as follows. In Section II, we define the problem and introduce the novel spatial regret metric. In Section III, we prove the well-posedness of the spatial regret. Moreover, we formulate the spatial regret optimization problem in a convex form. Lastly, in Section IV we report numerical examples to show the performance of spatial regret in the control of a distributed system.

Notation. $\mathbb{R}^{a \times b}$ represents the set of real matrices with dimension $a \times b$. \mathbb{N}^+ denotes the set of non-zero natural numbers. For a matrix $Y \in \mathbb{R}^{a \times b}$, $Y_{i,j}$ indicates its element in row i and column j . Given the vector $x \in \mathbb{R}^a$, x^i is used to indicate the i^{th} component of x , with $i \in \{1, \dots, a\}$. We denote the set of $a \times b$ binary matrices by $\mathbb{B}^{a \times b}$. For $X \in \mathbb{B}^{a \times b}$, we define the set of matrices with the same sparsity pattern of X as

$$\text{Sparse}(X) := \{Y \in \mathbb{R}^{a \times b} \mid Y_{i,j} = 0 \quad \forall i, j \text{ s. t. } X_{i,j} = 0\}.$$

With $\text{card}(X)$, we represent the number of nonzero elements in the binary matrix X . For a matrix $Z \in \mathbb{R}^{a \times b}$, $X := \text{Struct}(Z) \in \mathbb{B}^{a \times b}$ if

$$X_{i,j} = \begin{cases} 0 & \text{if } Z_{i,j} = 0, \\ 1 & \text{otherwise.} \end{cases}$$

For $X, Y \in \mathbb{B}^{a \times b}$, we say $X \leq Y$ if $X_{i,j} \leq Y_{i,j}$, $\forall i, j$. The operators $\|\cdot\|_F$ and $\|\cdot\|_{2 \rightarrow 2}$ represent the Frobenius and induced 2-norm of a matrix, respectively. Given a symmetric matrix $H \in \mathbb{R}^{a \times a}$, we denote the largest eigenvalue of H as $\lambda_{\max}(H)$. For a square matrix X , we use the notation $X \succ 0$ ($X \succeq 0$) to denote positive (semi-) definiteness. The operator \otimes represents the Kronecker product. Finally, $\text{blkdiag}(A, B, C, \dots)$ represents the block-diagonal matrix with matrices A, B, C, \dots on the diagonal.

Remark: For the sake of conciseness, the proofs of the results in this work are deferred to the Appendix of the extended manuscript [17].

II. PROBLEM FORMULATION

We consider discrete-time linear time-varying dynamical systems described by the state-space equations:

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

where $x_t \in \mathbb{R}^n$, $u_t \in \mathbb{R}^m$, and $w_t \in \mathbb{R}^n$ represent the system state, the control input, and an exogenous disturbance, respectively. Motivated by the observation that it is often difficult to characterize the class of disturbances, we make no assumptions regarding the distribution or the nature of w_t over time. For simplicity, we assume $E_t = I$. The more

general case where $E_t \neq I$ can be addressed using the methods described in [18].

We consider the scenario where the system described by (1) is controlled over a finite horizon $T \in \mathbb{N}^+$, starting from an initial condition $x_0 \in \mathbb{R}^n$. For compactness, we define

$$\mathbf{x} := \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_{T-1} \end{bmatrix}, \quad \mathbf{u} := \begin{bmatrix} u_0 \\ u_1 \\ \vdots \\ u_{T-1} \end{bmatrix}, \quad \boldsymbol{\delta} := \begin{bmatrix} x_0 \\ w_0 \\ \vdots \\ w_{T-2} \end{bmatrix} = \begin{bmatrix} x_0 \\ \mathbf{w} \end{bmatrix}.$$

At each time instant t , we denote the input u_t produced by a controller π_t as $u_t = \pi_t(x_0, \dots, x_t)$. For a policy $\pi := [\pi_0^\top(x_0) \dots \pi_{T-1}^\top(x_0, \dots, x_{T-1})]^\top$ and a disturbance $\boldsymbol{\delta}$, the incurred cost is defined as

$$J(\boldsymbol{\delta}, \pi) := [\mathbf{x}^\top \quad \mathbf{u}^\top] \mathbf{C} \begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix}. \quad (2)$$

where the matrix $\mathbf{C} \succeq 0$ assigns different weights to the states and input signals at different time instants. Notice that it is intractable to cast an optimization program over the class of all general policies π . Motivated by their optimality for centralized linear quadratic control and their tractability properties, we focus on linear feedback policies of the form $\mathbf{u} = \mathbf{K}\mathbf{x}$, where $\mathbf{K} \in \mathbb{R}^{mT \times nT}$ is a lower block-triangular matrix due to causality¹. To highlight the dependency of \mathbf{K} in (2), we will use the notation $J(\boldsymbol{\delta}, \mathbf{K})$.

This paper focuses on large-scale systems, where each controller has only access to partial sensor information. We represent this condition using the constraint

$$\mathbf{K} \in \mathcal{S}, \quad \mathcal{S} = \text{Sparse}(\mathbf{S}),$$

where $\mathbf{S} \in \mathbb{B}^{mT \times nT}$ describes the spatio-temporal information of the system, in the sense that \mathbf{S} describes which scalar control input depends on which scalar state, and at which time instant. We will refer to \mathbf{S} as ‘‘sparsity matrix’’. We always assume \mathbf{S} to be lower block-diagonal to ensure the causality of the control policy.

A. Spatial regret

The cost function (2) depends on the realization of the unknown disturbances $\boldsymbol{\delta}$. Hence, finding a sparse controller \mathbf{K} minimizing (2) for any $\boldsymbol{\delta}$ is an ill-posed problem. To remove the explicit dependency of the cost on $\boldsymbol{\delta}$, the \mathcal{H}_2 paradigm optimizes the expected performance under the assumption of stochastic disturbances with zero average and finite second moment, while the \mathcal{H}_∞ setup focuses on optimizing with respect to the worst-case disturbance realization. However, both setups are unlikely to hold in large-scale systems. The presence of unmodeled dynamics often leads to non-stochastic uncertainties, and disturbances are frequently localized to specific subsystems rather than representing a worst-case scenario. In this work, we propose an alternative controller performance metric that is tailored to large-scale scenarios.

¹An affine policy $\mathbf{u} = \mathbf{K}\mathbf{x} + \mathbf{g}$ can be also considered augmenting the state as $\tilde{x}_t := [x_t^\top \quad 1]^\top$. Thus, we will focus on linear feedback policies, without loss of generality.

Suppose, to have an ideal sparsity matrix $\hat{\mathcal{S}}$ that is denser than the real control sparsity structure \mathcal{S} , i.e., $\mathcal{S} \subset \hat{\mathcal{S}}$, with $\mathcal{S} = \text{Sparse}(\mathbf{S})$, and $\hat{\mathcal{S}} = \text{Sparse}(\hat{\mathbf{S}})$. Then, a controller $\hat{\mathbf{K}} \in \hat{\mathcal{S}}$ would have more sensor information to use, and it would be able to detect and counteract localized disturbances more promptly, regardless of the location where they may hit. Our idea is to promote the design of sparse controllers that imitate, in hindsight, the behavior of denser ones. To do so, first define the error between the cost of the controller $\mathbf{K} \in \mathcal{S}$ and $\hat{\mathbf{K}} \in \hat{\mathcal{S}}$ for a given δ as

$$e(\delta, \mathbf{K}, \hat{\mathbf{K}}) := J(\delta, \mathbf{K}) - J(\delta, \hat{\mathbf{K}}). \quad (3)$$

Then, we introduce a new metric that we call *Spatial Regret* quantifying the worst-case scenario of (3) as

$$\text{SpRegret}(\mathbf{K}, \hat{\mathbf{K}}) := \max_{\|\delta\|_2 \leq 1} e(\delta, \mathbf{K}, \hat{\mathbf{K}}). \quad (4)$$

Finally, the minimization problem we want to solve is

$$\min_{\mathbf{K}} \text{SpRegret}(\mathbf{K}, \hat{\mathbf{K}}) \quad (5a)$$

$$\text{s. t. } \mathbf{K} \in \mathcal{S}, \quad \mathcal{S} \subset \hat{\mathcal{S}}. \quad (5b)$$

From (4), note that a non-positive spatial regret value implies that the oracle $\hat{\mathbf{K}}$ cannot improve the performance of the controller \mathbf{K} to be designed for any disturbance $\delta \neq 0$. Therefore, (4) is not *well-posed* in the sense that it is meaningless to match the performance of a universally worse benchmark for any disturbance. An example demonstrating this point is reported in [17][App. A]. Problem (5) presents two key challenges. First, we need to verify when it is well-posed, meaning that $\text{SpRegret}(\cdot, \cdot)$ is lower bounded for any selected controller $\mathbf{K} \in \mathcal{S}$, given the oracle $\hat{\mathbf{K}}$. Second, notice that (5) is composed of multiple nested optimization problems, making its solution intractable in this form. In Section III-B we provide a convex approximation of (5).

Remark 1: Our SpRegret metric is inspired by [14]–[16], [18]–[20], where centralized controllers are designed by minimizing regret with respect to a benchmark policy that has foreknowledge of future realization of δ . Unlike these and related works rooted in online optimization, see e.g., [21], our approach introduces a spatial notion of regret instead of a temporal one. Studying the interplay between spatial and temporal notions of regret is left as an interesting direction for future research.

Remark 2: In the case of a perfect oracle achieving zero cost for every disturbance realization, (4) reduces to the objective of classic \mathcal{H}_∞ controllers. However, the key strength of our metric emerges when considering imperfect oracles: the spatial regret metric allows mimicking the behavior of an unattainable policy that has additional sensor data (which we lack). This approach prioritizes mitigating disturbances where this additional information would be most beneficial, all while complying with the information sparsity of the system.

B. Review of convex design of distributed controllers

We review results on the convex design of distributed controllers that are instrumental in solving the above

challenges. Let \mathbf{Z} denote the block-downshift operator, namely a block-matrix with identity matrices along its first block sub-diagonal and zeros elsewhere. Define the matrices $\mathbf{A} := \text{blkdiag}(A_0, \dots, A_{T-2}, 0_{n \times n})$, and $\mathbf{B} := \text{blkdiag}(B_0, \dots, B_{T-2}, 0_{n \times m})$. Then, the state evolution of (1) can be represented compactly as

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

Considering a control law $\mathbf{u} = \mathbf{K}\mathbf{x}$ and (6), it is straightforward to write the closed-loop maps Φ_x, Φ_u from the disturbances δ to \mathbf{x} and \mathbf{u} as

$$\begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} (I - \mathbf{Z}(\mathbf{A} + \mathbf{B}\mathbf{K}))^{-1} \\ \mathbf{K}(I - \mathbf{Z}(\mathbf{A} + \mathbf{B}\mathbf{K}))^{-1} \end{bmatrix} \delta = \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} \delta = \Phi \delta. \quad (7)$$

Here, Φ is the vertical stacking of Φ_x, Φ_u . The closed-loop responses Φ_x, Φ_u are lower block-diagonal due to causality. It is easy to verify that \mathbf{K} and Φ are linked through the relation $\mathbf{K} = h(\Phi) = \Phi_u \Phi_x^{-1}$, with $h: \mathbb{R}^{(m+n)T \times nT} \rightarrow \mathbb{R}^{mT \times nT}$. Also, as shown in [22], it can be proved that there exists a controller \mathbf{K} such that (7) holds if and only if

$$(I - \mathbf{Z}\mathbf{A})\Phi_x - \mathbf{Z}\mathbf{B}\Phi_u = I. \quad (8)$$

With the introduction of the maps Φ_x, Φ_u , the cost $J(\delta, \mathbf{K})$ can be rewritten as

$$J(\delta, \mathbf{K}) = \delta^\top \Phi^\top \mathbf{C} \Phi \delta. \quad (9)$$

Moreover, as shown in [22], the classic \mathcal{H}_2 and \mathcal{H}_∞ control problems, can be reformulated as

$$\mathcal{H}_2: \quad \mathbb{E}_{\delta \sim \mathcal{D}}[J(\delta, \mathbf{K})] = \left\| \mathbf{C}^{\frac{1}{2}} \Phi \Sigma_\delta \right\|_F^2, \quad (10)$$

$$\mathcal{H}_\infty: \quad \max_{\|\delta\|_2 \leq 1} (J(\delta, \mathbf{K})) = \left\| \mathbf{C}^{\frac{1}{2}} \Phi \right\|_{2 \rightarrow 2}^2, \quad (11)$$

where \mathcal{D} in (10) denotes the probability distribution of δ , with zero mean and covariance $\Sigma_\delta \succeq 0$. Thus, the costs (10) and (11) are convex in Φ . Nonetheless, the sparsity condition (5b) is nonconvex in Φ . However, when the QI condition defined next holds, the sparsity constraint (5b) can be rewritten linearly in Φ [4].

Definition 1: Define $\mathbf{G} := (I - \mathbf{Z}\mathbf{A})^{-1}\mathbf{Z}\mathbf{B}$. A subspace $\mathcal{S} \subseteq \mathbb{R}^{mT \times nT}$ is QI with respect to \mathbf{G} if and only if

$$\mathbf{K}\mathbf{G}\mathbf{K} \in \mathcal{S}, \quad \forall \mathbf{K} \in \mathcal{S}.$$

If \mathcal{S} is QI with respect to \mathbf{G} , then it is was shown [6] that

$$\mathbf{K} \in \mathcal{S} \Leftrightarrow \Phi_u \Gamma \in \mathcal{S}, \quad (12)$$

where $\Gamma := I - \mathbf{Z}\mathbf{A}$. Thus, the QI condition enables us to search over all possible sparse controllers $\mathbf{K} \in \mathcal{S}$, allowing globally optimal minimization of any convex cost, such as (10), (11). However, the QI condition may not be satisfied for every desired dynamical system or sparsity matrix. In [9], a technique was introduced for deriving sparse controllers, even in cases where QI does not hold. Specifically, the work

[9] establishes a method to compute, given \mathcal{S} , a matrix $\mathbf{V}_x \in \mathbb{B}^{nT \times nT}$ such that

$$\Phi_u \Gamma \in \mathcal{S} \quad \text{and} \quad \Phi_x \Gamma \in \text{Sparse}(\mathbf{V}_x), \quad (13a)$$

$$\begin{aligned} &\downarrow \\ \mathbf{K} &= h(\Phi) = \Phi_u \Phi_x^{-1} \in \mathcal{S}, \end{aligned} \quad (13b)$$

where Φ_x, Φ_u satisfy (8). The authors refer to (13) as the Sparsity Invariance (SI) condition. A minimally restrictive choice of the binary matrix \mathbf{V}_x complying with (13) is given by [9, Alg. 1], which we report here as Algorithm 1 for the sake of completeness. In this paper, we denote the set of

Algorithm 1 Generation of \mathbf{V}_x [9, Alg. 1]

Input: \mathbf{S} (System's Sparsity Matrix)

- 1: Initialize $\mathbf{V}_x = \mathbf{1}_{nT \times nT}$
 - 2: **for** each $i = 1, \dots, mT, k = 1, \dots, nT$ **do**
 - 3: **if** $\mathbf{S}_{i,k} == 0$ **then**
 - 4: **for** each $j = 1, \dots, nT$ **do**
 - 5: **if** $\mathbf{S}_{i,j} == 1$ **then** $(\mathbf{V}_x)_{j,k} \leftarrow 0$
-

sparse controllers parametrized by SI as

$$\Theta := \{h(\Phi) : (8), (13a)\}. \quad (14)$$

If \mathbf{V}_x is designed according to Algorithm 1, then we have $\Theta \subseteq \mathcal{S}$.

III. MAIN RESULTS

In this section, we address the two key challenges of problem (5) regarding the well-posedness of the metric and its convex reformulation. First, we establish that if the oracle $\hat{\mathbf{K}} \in \hat{\mathcal{S}}$ is optimal according to \mathcal{H}_2 or \mathcal{H}_∞ criteria, then SpRegret is well-posed in the sense that imitating its behavior is advantageous for some δ . Second, we demonstrate that the QI condition plays a crucial role in the design of the oracle $\hat{\mathbf{K}}$ for enabling a convex synthesis of the spatial regret optimal policy \mathbf{K} .

Proposition 1: Assume an oracle $\hat{\mathbf{K}}$ is derived by solving the optimization problem

$$\min_{\hat{\mathbf{K}}} f(\hat{\mathbf{K}}) \quad \text{s. t.} \quad \hat{\mathbf{K}} \in \hat{\mathcal{S}}, \quad \mathcal{S} \subset \hat{\mathcal{S}} \quad (15)$$

where $f(\hat{\mathbf{K}})$ is $\mathbb{E}_{\delta \sim \mathcal{D}}[J(\delta, \hat{\mathbf{K}})]$ or $\max_{\|\delta\|_2 \leq 1} (J(\delta, \hat{\mathbf{K}}))$. Then, $\text{SpRegret}(\mathbf{K}, \hat{\mathbf{K}}) \geq 0$ for any $\mathbf{K} \in \mathcal{S}$.

Proposition 1 sheds light on two possible design criteria. One could follow to synthesize a denser controller $\hat{\mathbf{K}} \in \hat{\mathcal{S}}$ that is guaranteed to be informative, in the sense that no other controller $\mathbf{K} \in \mathcal{S}$ can outperform it for all possible realizations of δ . Despite these advancements, how to obtain an oracle $\hat{\mathbf{K}}$ with sparsity constraints $\hat{\mathcal{S}}$ by solving (15) remains still a non-trivial challenge. The next section addresses this crucial aspect in detail.

Remark 3: With an unconstrained noncausal benchmark, obtaining an optimal benchmark for any δ is possible, as shown in [13], [16]. Therefore, formulations that minimize worst-case regret with a noncausal oracle are always well-posed. Our result in Proposition 1 reaffirms that spatial

regret preserves this property even when introducing sparsity constraints in the oracle design.

A. Synthesis of the oracle $\hat{\mathbf{K}}$

Thanks to Proposition 1, we have demonstrated there always exist choices for the oracle $\hat{\mathbf{K}}$ that ensure the well-posedness of problem (5). For example, an optimal centralized \mathcal{H}_2 or \mathcal{H}_∞ oracle $\hat{\mathbf{K}}$, is a valid choice. However, a key challenge remains how to effectively synthesize the distributed controllers \mathbf{K} and $\hat{\mathbf{K}}$ while guaranteeing well-posedness of the spatial regret. Since (13) yields linear constraints on Φ that come with tightness guarantees and comply with any desired information structure, a seemingly natural approach would be to exploit Algorithm 1 to synthesize both \mathbf{V}_x and $\hat{\mathbf{V}}_x$, satisfying the SI condition (13) for \mathcal{S} and $\hat{\mathcal{S}}$, respectively. It must be pointed out that these steps might fail to obtain a well-posed spatial regret metric, in general. An example illustrating this point is detailed in [17, App. C]. Nonetheless, in the following theorem, we reveal the crucial role of the QI condition in designing an informative oracle $\hat{\mathbf{K}}$ in conjunction with SI for the synthesis of the spatial regret optimal policy for any \mathcal{S} .

Theorem 1: Let $\hat{\mathcal{S}}$ be QI with respect to \mathbf{G} , with $\mathcal{S} \subset \hat{\mathcal{S}}$. Additionally, denote with \mathbf{V}_x and $\hat{\mathbf{V}}_x$ the binary matrices computed using Algorithm 1 based on \mathcal{S} and $\hat{\mathcal{S}}$, respectively. Assume an oracle $\hat{\mathbf{K}} = h(\hat{\Phi})$ is derived through the optimization problem

$$\min_{\hat{\Phi}} f(\hat{\Phi}) \quad \text{s. t.} \quad (8), \quad \hat{\Phi}_u \Gamma \in \hat{\mathcal{S}}, \quad \hat{\Phi}_x \Gamma \in \text{Sparse}(\hat{\mathbf{V}}_x), \quad (16)$$

where the function $f(\hat{\Phi})$ can be either (10) or (11) and $\Gamma := (I - \mathbf{Z}\mathbf{A})$. Then, $\text{SpRegret}(\mathbf{K}, \hat{\mathbf{K}}) \geq 0$ for any $\mathbf{K} \in \Theta$.

With Theorem 1, we proved that if the oracle structure satisfies the QI condition, then the well-posedness of the spatial regret metric is preserved using the SI condition (13) to enforce both sparsity constraints over the oracle $\hat{\mathbf{K}}$ and then the actual controller \mathbf{K} .

Remark 4: Given the sparsity matrix \mathbf{S} , one can obtain its closest QI superset $\hat{\mathcal{S}}_{QI}$ such that

$$\begin{aligned} \hat{\mathcal{S}}_{QI} &:= \underset{\mathbf{S}^*}{\text{argmin}} \quad \text{card}(\mathbf{S}^* - \mathbf{S}) \\ \text{s. t.} \quad &\mathcal{S} \subseteq \mathbf{S}^*, \quad \mathbf{S}^* = \text{Sparse}(\mathbf{S}^*), \\ &\mathbf{S}^* \text{ is QI w.r.t. } \mathbf{G}. \end{aligned} \quad (17)$$

In [23] the authors proposed a method to obtain efficiently $\hat{\mathcal{S}}_{QI}$ in a finite number of steps. While results hold for any QI subspace $\hat{\mathcal{S}} \supset \mathcal{S}$, the method (17) can always be utilized to generate $\hat{\mathcal{S}}$.

Remark 5: It is important to note that the ‘‘distance’’ in terms of cardinality between the two sparsity matrices \mathbf{S} and $\hat{\mathcal{S}}$ plays an important role in obtaining well-performing controllers. When $\text{card}(\mathbf{S})$ is approximately equal to $\text{card}(\hat{\mathcal{S}})$, the controller \mathbf{K} may struggle to learn adequately, leading it to imitate the oracle $\hat{\mathbf{K}}$. Conversely, when the two structures are significantly different, with $\text{card}(\mathbf{S}) \ll \text{card}(\hat{\mathcal{S}})$, the behavior of the oracle may be too challenging for the actual controller to follow, resulting in poor performance. Choosing

the nearest QI superset has been demonstrated to yield heuristically good performance, particularly when the actual matrix \mathbf{G} is highly sparse.

B. Convex reformulation of the spatial regret problem

Here, we tackle the challenge of obtaining a convex approximation of the SpRegret minimization problem (5).

Proposition 2: Let $\hat{\mathbf{K}} = h(\hat{\Phi}) \in \text{Sparse}(\hat{\mathbf{S}})$ be a solution of (16), where $\hat{\mathbf{S}}$ is QI with respect to \mathbf{G} . Consider the following convex optimization problem

$$\min_{\Phi, \lambda} \lambda \quad (18a)$$

$$\text{s. t.} \quad (8), (13a),$$

$$\begin{bmatrix} I & \mathbf{C}^{\frac{1}{2}} \Phi \\ \Phi^{\top} \mathbf{C}^{\frac{1}{2}} & \lambda I + \hat{\Phi}^{\top} \mathbf{C} \hat{\Phi} \end{bmatrix} \succeq 0, \quad (18b)$$

where \mathbf{V}_x in (13a) is designed according to Algorithm 1. Let the set of parametrized controllers $\Theta \subseteq \mathcal{S}$ be defined as per (14). Finally, let Φ^* denote an optimal solution to (18) and $\mathbf{K}^* = h(\Phi^*)$ the corresponding controller. Then:

- 1) $\text{SpRegret}(\mathbf{K}, \hat{\mathbf{K}}) \geq 0$ for any $\mathbf{K} \in \Theta$,
- 2) $\text{SpRegret}(\mathbf{K}^*, \hat{\mathbf{K}}) \leq \text{SpRegret}(\mathbf{K}, \hat{\mathbf{K}})$ for all $\mathbf{K} \in \Theta$,
- 3) If \mathcal{S} is QI with respect to \mathbf{G} , then \mathbf{K}^* is a globally optimal solution to problem (5).

For the sake of clarity, we summarize the steps of our controller design method in Algorithm 2.

Algorithm 2 Controller synthesis for SpRegret

Input: \mathbf{S} (System's Sparsity Matrix)

- 1: Choose a QI subspace $\hat{\mathcal{S}} \supseteq \mathcal{S}$, for instance, using (17).
 - 2: Compute \mathbf{V}_x and $\hat{\mathbf{V}}_x$ with Algorithm 1, starting from \mathbf{S} and $\hat{\mathbf{S}}$, respectively.
 - 3: Derive the oracle controller $\hat{\mathbf{K}}$ by solving (16).
 - 4: Obtain the SpRegret controller \mathbf{K} through (18).
-

IV. NUMERICAL RESULTS

We conduct a comparative analysis between distributed controllers minimizing spatial regret metrics and traditional $\mathbf{K}_{\mathcal{H}_2}$ and $\mathbf{K}_{\mathcal{H}_\infty}$ ones, all utilizing the same sparsity structure \mathcal{S} and synthesized using the SI condition (13). We use a multi-mass spring-damper system consisting of 10 masses within a time window of $T = 30$ and a sampling time of $T_s = 0.5$ s. For full model details and simulation settings, please refer to [17, App. F].²

We design two SpRegret controllers: $\mathbf{K}_{R_{QI}}$ and \mathbf{K}_{R_C} . The former employs an oracle with sparsity determined by the nearest QI superset of \mathbf{S} . The latter imitates a centralized oracle, which is guaranteed to be QI by definition. We perform two distinct experiments, the results of which are depicted in Fig. 1 and Fig. 2. In both cases, our main focus is to determine the percentage of times in which each controller yields a better (i.e., smaller) cost $J(\cdot)$ compared to the other three.

²The code used in this work is accessible at <https://github.com/DecodEPFL/SpRegret>.

In the first experiment, we aim to replicate the unpredictability and uncertain nature of disturbances by simulating perturbations δ , drawn from a non-centered uniform distribution $\mathcal{U}[-0.5, 1]$. These disturbances are then applied to a random number of masses within the system. The number of affected masses is also drawn uniformly, within the interval $[1, N]$. We gradually increase the value of N up to the total number of interconnected subsystems (in this case 10). The results are presented in Fig. 1. The percentage value of times each policy yields better performance is computed over 1000 realizations of δ and, to obtain the 95% confidence interval of these measurements, we iterated this step 100 times, for a total of 10^5 experiments. As expected, when the number of affected masses is small, the resulting disturbance is zero for most of the agents, aligning with the \mathcal{H}_2 hypothesis on a distribution \mathcal{D} centered around 0. However, as the number of affected masses increases, the cumulative effect of non-standard δ on the overall system becomes more pronounced, thus favoring the performance of the controller $\mathbf{K}_{R_{QI}}$. In the table below, we report the average cost of the controller $\mathbf{K}_{R_{QI}}$ over 10^5 experiments and the relative increase in cost for the other controllers compared to $\mathbf{K}_{R_{QI}}$ for an increasing number of masses.

| N.Masses | Av. cost $\mathbf{K}_{R_{QI}}$ | $\mathbf{K}_{\mathcal{H}_2}$ | $\mathbf{K}_{\mathcal{H}_\infty}$ | \mathbf{K}_{R_C} |
|----------|--------------------------------|------------------------------|-----------------------------------|--------------------|
| 1 | 14.20 | +7.39% | +1.35% | +0.44% |
| 5 | 30.21 | +13.43% | +3.32% | +1.25% |
| 10 | 39.35 | +43.82% | +4.32% | +1.20% |

As we can observe, $\mathbf{K}_{R_{QI}}$ consistently achieves lower average control costs. Notably, the cost advantage of \mathbf{K}_{QI} over other controllers increases with system complexity (number of masses). For instance, with 10 masses, $\mathbf{K}_{R_{QI}}$ outperforms the $\mathbf{K}_{\mathcal{H}_2}$ controller by over 43%, while just improving by 4% and 1% compared to $\mathbf{K}_{\mathcal{H}_\infty}$ and \mathbf{K}_{R_C} , respectively. This demonstrates the effectiveness of our spatial regret-based approach in achieving efficient control, particularly for larger and potentially more challenging systems.

In the second experiment, we evaluate the performance improvement of controllers minimizing spatial regret as the sparsity structure of the large-scale system becomes more and more distributed. For this reason, we choose a QI benchmark as it strikes a favorable balance between imitating a sparse structure and a centralized one. We consider the same multi-mass spring damper system of the first experiment with a progressively increasing number of masses, from 3 to 10. For each new configuration, we synthesize again all the controllers ($\mathbf{K}_{\mathcal{H}_2}$, $\mathbf{K}_{\mathcal{H}_\infty}$, $\mathbf{K}_{R_{QI}}$). Then, we simulate the effects of perturbations δ drawn from a non-centered uniform distribution $\mathcal{U}[-0.5, 1]$ applied to all the masses with the same amount of experiments as in the previous study case. The results are shown in Fig. 2. It is evident that for a small number of systems affected by noise, $\mathbf{K}_{\mathcal{H}_\infty}$ outperforms the other control policies. However, as the size of the system increases, and consequently its sparsity, $\mathbf{K}_{R_{QI}}$ exhibits superior performance, highlighting its capacity to leverage information from the ideal oracle.

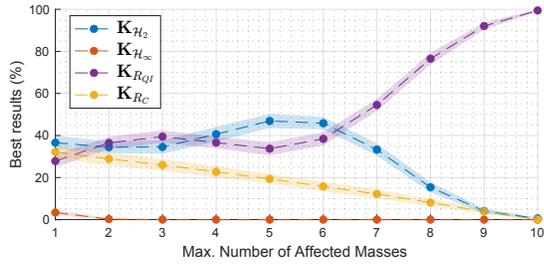


Fig. 1. Percentage of times a control policy outperforms the remaining three as a function of the maximum number of masses affected by uniformly distributed disturbances in the interval $[-0.5, 1]$ in a system of 10 masses. Shaded areas around the dotted lines represent the 95% confidence interval around the corresponding mean values.

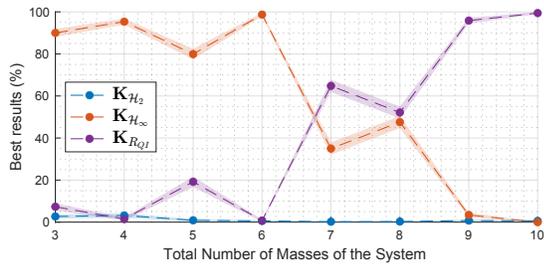


Fig. 2. Percentage of times a control policy outperforms the remaining two as a function of the number of masses constituting the large-scale systems when all the masses are affected by uniformly distributed disturbances in the interval $[-0.5, 1]$. Shaded areas around the dotted lines represent the 95% confidence interval around the corresponding mean values.

V. CONCLUSIONS

In this work, we aimed to design and synthesize distributed controllers for large-scale linear dynamical systems, affected by localized and highly heterogeneous disturbances. To do so, we first introduced the novel metric SpRegret. Then, we demonstrated its well-posedness when the oracle satisfies the QI condition. Finally, we provided a convex formulation for designing regret-optimal controllers with arbitrary sparsity structures, optimizing them to close the performance gap to an ideal QI subspace that encodes richer information for the control policy. To illustrate the real-world relevance of our approach, we provided numerical examples involving a multi-mass spring-damper system that requires to be controlled in the presence of non-standard disturbances. Through comparisons with classic \mathcal{H}_2 and \mathcal{H}_∞ policies, our results showcased the superior performance of SpRegret controllers in handling disturbances that may target large-scale distributed systems.

Future research will explore methods for designing constrained benchmarks that ensure non-negative regret and improved performance over \mathcal{H}_2 and \mathcal{H}_∞ for user-defined disturbance classes. Possible new directions will include investigating how to automatically select the sparsity structure of the oracle tailored to the problem to guarantee better performance, extend our results to the infinite-horizon case, and consider larger and more complex applications to showcase the potentiality of our novel metric.

REFERENCES

- [1] D. K. Molzahn, F. Dörfler, H. Sandberg, S. H. Low, S. Chakrabarti, R. Baldick, and J. Lavaei, "A survey of distributed optimization and control algorithms for electric power systems," *IEEE Transactions on Smart Grid*, vol. 8, no. 6, pp. 2941–2962, 2017.
- [2] Y. Zheng, S. E. Li, K. Li, F. Borrelli, and J. K. Hedrick, "Distributed model predictive control for heterogeneous vehicle platoons under unidirectional topologies," *IEEE Transactions on Control Systems Technology*, vol. 25, no. 3, pp. 899–910, 2016.
- [3] H. S. Witsenhausen, "A counterexample in stochastic optimum control," *SIAM Journal on Control*, vol. 6, no. 1, pp. 131–147, 1968.
- [4] M. Rotkowitz and S. Lall, "A characterization of convex problems in decentralized control," *IEEE Transactions on Automatic Control*, vol. 50, no. 12, pp. 1984–1996, 2005.
- [5] L. Lessard and S. Lall, "Quadratic invariance is necessary and sufficient for convexity," in *Proceedings of the American Control Conference*, pp. 5360–5362, IEEE, 2011.
- [6] L. Furieri and M. Kamgarpour, "Unified approach to convex robust distributed control given arbitrary information structures," *IEEE Transactions on Automatic Control*, vol. 64, no. 12, pp. 5199–5206, 2019.
- [7] Y.-S. Wang, N. Matni, and J. C. Doyle, "A system-level approach to controller synthesis," *IEEE Transactions on Automatic Control*, vol. 64, no. 10, pp. 4079–4093, 2019.
- [8] L. Furieri, Y. Zheng, A. Papachristodoulou, and M. Kamgarpour, "On separable quadratic lyapunov functions for convex design of distributed controllers," in *18th European control conference (ECC)*, pp. 42–49, IEEE, 2019.
- [9] L. Furieri, Y. Zheng, A. Papachristodoulou, and M. Kamgarpour, "Sparsity invariance for convex design of distributed controllers," *IEEE Transactions on Control of Network Systems*, vol. 7, no. 4, pp. 1836–1847, 2020.
- [10] R. Arastoo, M. Bahavarnia, M. V. Kothare, and N. Motee, "Closed-loop feedback sparsification under parametric uncertainties," in *IEEE 55th Conference on Decision and Control (CDC)*, pp. 123–128, IEEE, 2016.
- [11] E. Jensen and B. Bamieh, "An explicit parametrization of closed loops for spatially distributed controllers with sparsity constraints," *IEEE Transactions on Automatic Control*, vol. 67, no. 8, pp. 3790–3805, 2021.
- [12] S. Fattahi, G. Fazelnia, J. Lavaei, and M. Arcak, "Transformation of optimal centralized controllers into near-globally optimal static distributed controllers," *IEEE Transactions on Automatic Control*, vol. 64, no. 1, pp. 66–80, 2018.
- [13] B. Hassibi, A. H. Sayed, and T. Kailath, *Indefinite-Quadratic estimation and control: a unified approach to \mathcal{H}_2 and \mathcal{H}_∞ theories*. SIAM, 1999.
- [14] O. Sabag, G. Goel, S. Lale, and B. Hassibi, "Regret-optimal controller for the full-information problem," in *American Control Conference (ACC)*, pp. 4777–4782, IEEE, 2021.
- [15] G. Goel and B. Hassibi, "Regret-optimal estimation and control," *IEEE Transactions on Automatic Control*, vol. 68, no. 5, pp. 3041–3053, 2023.
- [16] A. Martin, L. Furieri, F. Dörfler, J. Lygeros, and G. Ferrari-Trecate, "Safe control with minimal regret," in *Learning for Dynamics and Control Conference*, pp. 726–738, PMLR, 2022.
- [17] D. Martinelli, A. Martin, G. Ferrari-Trecate, and L. Furieri, "Closing the gap to quadratic invariance: a regret minimization approach to optimal distributed control," *arXiv preprint arXiv:2311.02068*, 2023.
- [18] A. Didier, J. Sieber, and M. N. Zeilinger, "A system level approach to regret optimal control," *IEEE Control Systems Letters*, vol. 6, pp. 2792–2797, 2022.
- [19] A. Martin, L. Furieri, F. Dörfler, J. Lygeros, and G. Ferrari-Trecate, "On the guarantees of minimizing regret in receding horizon," *arXiv preprint arXiv:2306.14561*, 2023.
- [20] A. Martin, L. Furieri, F. Dörfler, J. Lygeros, and G. Ferrari-Trecate, "Follow the clairvoyant: an imitation learning approach to optimal control," *IFAC-PapersOnLine*, vol. 56, no. 2, pp. 2589–2594, 2023.
- [21] U. Ghai, U. Madhushani, N. Leonard, and E. Hazan, "A regret minimization approach to multi-agent control," in *International Conference on Machine Learning*, pp. 7422–7434, PMLR, 2022.
- [22] J. Anderson, J. C. Doyle, S. H. Low, and N. Matni, "System level synthesis," *Annual Reviews in Control*, vol. 47, pp. 364–393, 2019.
- [23] M. C. Rotkowitz and N. C. Martins, "On the nearest quadratically invariant information constraint," *IEEE Transactions on Automatic Control*, vol. 57, no. 5, pp. 1314–1319, 2011.