2024 IEEE 63rd Conference on Decision and Control (CDC)
December 16-19, 2024. MiCo, Milan, Italy

The Distributionally Robust Infinite-Horizon LQR

Joudi Hajar Taylan Kargin Vikrant Malik Babak Hassibi
Department of Electrical Engineering, Caltech, Pasadena, CA 91125
{jhajar,tkargin,vmalik,hassibi}@caltech.edu

***Abstract*— We explore the infinite-horizon Distributionally Robust (DR) linear-quadratic control. While the probability distribution of disturbances is unknown and potentially correlated over time, it is confined within a Wasserstein-2 ball of a radius r around a known nominal distribution. Our goal is to devise a control policy that minimizes the worst-case expected Linear-Quadratic Regulator (LQR) cost among all probability distributions of disturbances lying in the Wasserstein ambiguity set. We obtain the optimality conditions for the optimal DR controller and show that it is non-rational. Despite lacking a finite-order state-space representation, we introduce a computationally tractable fixed-point iteration algorithm. Our proposed method computes the optimal controller in the frequency domain to any desired fidelity. Moreover, for any given finite order, we use a convex numerical method to compute the best rational approximation (in H_∞-norm) to the optimal non-rational DR controller. This enables efficient time-domain implementation by finite-order state-space controllers and addresses the computational hurdles associated with the finite-horizon approaches to DR-LQR problems, which typically necessitate solving a Semi-Definite Program (SDP) with a dimension scaling with the time horizon. We provide numerical simulations to showcase the effectiveness of our approach.**

I. Introduction

Mitigating uncertainties is a core challenge in decision-making. Control systems inherently encounter various uncertainties, such as external disturbances, measurement errors, model disparities, and temporal variations in dynamics [1], [2]. Neglecting these uncertainties in policy design can result in considerable performance decline and may lead to unsafe and unintended behavior [3].

Traditionally, the challenge of uncertainty mitigation in control systems has been predominantly approached through either the stochastic or robust control frameworks [4], [5], [6]. As exemplified in Linear–Quadratic–Gaussian (LQG) or H_2 control, stochastic control aims to minimize an expected cost, assuming disturbances follow a known probability distribution [4], [7]. However, in many practical scenarios, the true distribution is often estimated from sampled data, introducing vulnerability to inaccurate models. On the other hand, robust control minimizes the worst-case cost across potential disturbance realizations, such as those with bounded energy or power (H_∞ control) [8]. While this ensures robustness, it can be overly conservative.

To tackle the above challenge, a recent approach called *Distributionally Robust (DR) Control* has emerged. In contrast to traditional approaches such as H_2 or H_∞ control that focus on a single probability distribution or a worst-case disturbance realization, the DR framework addresses the uncertainty in system dynamics and disturbances by considering ambiguity sets – sets of plausible probability distributions [9], [10], [11], [12], [13], [14], [15]. This methodology aims to design controllers that exhibit robust performance across all probability distributions within a given ambiguity set. The size of the ambiguity set provides control over the desired robustness against distributional uncertainty, ensuring that the resulting controller is not excessively conservative. Thereby, this approach bridges the domains of stochastic and adversarial uncertainties.

Different measures of distributional mismatch, such as total variation [16], [17] and KL divergence [18], are explored in DR control. However, for computational feasibility and geometric interpretability, ambiguity sets are commonly defined as Wasserstein-2 balls around a nominal distribution [19], [20]. This choice is practical since optimizing quadratic costs over Wasserstein-2 balls leads to a semi-definite program (SDP).

A. Contributions

In this study, we explore the Wasserstein-2 distributionally robust LQR (DR-LQR) control framework. DR-LQR control seeks to design controllers that minimize the worst-case expected cost across distributions chosen adversarially within a Wasserstein-2 ambiguity set. Our contributions are summarized as follows.

a) Stabilizing Infinite-Horizon Controller.: Rather than the finite-horizon setting prevalent in the DR control literature [11], [12], [14], [21], [22], we focus on the infinite-horizon setting. Thus, we provide long-term stability and robustness guarantees.

b) Robustness to Arbitrarily Correlated Disturbances.: Unlike several prior works which assume time-independence of the disturbances [9], [10], [11], [12], [23], [13], [14], we do not impose such assumptions so that the resulting controllers are robust against time-correlated disturbances.

c) Computationally Efficient Controller Synthesis: Leveraging a strong duality result, we obtain the exact Karush-Khun-Tucker (KKT) conditions for the worst-case distribution and the optimal causal controller. We show that, although the resulting controller is non-rational, *i.e.*, it does not admit a finite state-space form, it can still be computed very efficiently. We provide a computationally efficient numerical method to compute the optimal non-rational DR-LQR controller in the frequency domain via fixed-point iterations. We further show how to find the best rational approximation and, thereby, the best finite-dimensional state-space controller for any given degree.

Prior works focus on finite horizon problems (see [21], [22], [12]) and therefore have no stability guarantees. More importantly, they are hampered by the need to solve a semi-definite program (SDP) whose size scales with the time horizon. This prohibits their applicability when the time horizon is large. Our approach enables efficient implementation of the infinite-horizon DR-LQR controller.

Similar earlier work by [24] studied the problem of infinite horizon distributionally robust regret-optimal (DR-RO) control. The DR-RO control is similar to the DR-LQR problem considered in this paper since the cost in both cases is quadratic. However, the regret-optimal controller (originally studied in [25], [26]) is a much simpler problem in the distributionally robust setting of [24] since the cost of the optimal non-causal controller is removed from the LQR cost. In this paper, we show that, despite the more complicated form of the quadratic cost, the main results of [24] extend to the LQR case. While the KKT conditions, the fixed-point iterative algorithm, and the final controller differ from the regret-optimal counterpart in[24], the general methodology follows similarly.

II. Preliminaries and Problem Setup

A. Notations

Going forward, calligraphic letters ($\mathcal{K}$, $\mathcal{M}$, $\mathcal{L}$, etc.) represent infinite-horizon operators, while boldface letters ($\mathbf{K}$, $\mathbf{C}$, $\mathbf{w}$, etc.) denote finite-horizon operators. $\mathcal{I}$ and $\mathbf{I}$ are the identity operators. $\mathcal{M}^*$ is the adjoint of $\mathcal{M}$, and $\succ$ denotes the positive-definite ordering. The trace functions for finite and infinite horizon operators are denoted by $\operatorname{tr}(\cdot)$ and $\operatorname{Tr}(\cdot)$, respectively, where $\operatorname{Tr}(\mathcal{I}) = p$ for a finite horizon. The Euclidean norm is denoted by $\|\cdot\|$, while $\|\cdot\|_\infty$ and $\|\cdot\|_2$ refer to the H_∞ (operator) and H_2 (Frobenius) norms, respectively. $\{\mathcal{M}\}_+$ and $\{\mathcal{M}\}_-$ are the causal and strictly anti-causal parts of $\mathcal{M}$. The notation $\sqrt{\mathcal{M}}$ or $\mathcal{M}^{\frac{1}{2}}$ indicates the symmetric positive square root. $[\cdot]_T$ signifies the finite-horizon restriction of operators. $|z|$ is the magnitude and z^* is the conjugate of a complex number $z \in \mathbb{C}$. The complex unit circle is denoted by $\mathbb{T}$. Finally, $\sigma_{\max}$ denotes the maximum singular value, id denotes the identity map, and $\times$ denotes the cartesian product.

B. Linear-Quadratic Control

Consider the state-space representation of discrete-time linear time-invariant (LTI) dynamical system:

$$x_{t+1} = Ax_t + B_u u_t + B_w w_t, \tag{1}$$

Here, $x_t \in \mathbb{R}^n$ denotes the state, $u_t \in \mathbb{R}^d$ is the control input, and $w_t \in \mathbb{R}^p$ is the disturbance. We posit that both the pairs (A, B_u) and (A, B_w) are stabilizable in the usual sense. For a finite horizon $T > 0$, the system in (1) incurs a quadratic cost as

$$\operatorname{cost}_T(\mathbf{u}, \mathbf{w}) := \sum\nolimits_{t=0}^{T-1} x_t^\intercal Q x_t + u_t^\intercal R u_t, \tag{2}$$

where $Q, R \succ 0$. Without loss of generality, we let $Q = I$, $R = I$ by redefining $x_t \leftarrow Q^{\frac{1}{2}} x_t$ and $u_t \leftarrow R^{\frac{1}{2}} u_t$.

a) System Description in Operator Form: We opt for operator notation for system (1) in the rest of this paper. For a given horizon $T > 0$, we let the sequences $\mathbf{x} := \{x_t\}_{t=0}^{T-1}$, $\mathbf{u} := \{u_t\}_{t=0}^{T-1}$, and $\mathbf{w} := \{w_t\}_{t=0}^{T-1}$, represent the state, control input and exogenous disturbances, respectively. Likewise, we represent their infinite-horizon counterparts using the bi-infinite sequences $\mathsf{x} := \{x_t\}_{t\in\mathbb{Z}}$, $\mathsf{u} := \{u_t\}_{t\in\mathbb{Z}}$, and $\mathsf{w} := \{w_t\}_{t\in\mathbb{Z}}$.

Using the above definitions, we express the system dynamics for both finite and infinite horizons in operator form as

$$\begin{aligned} &\text{Finite-horizon: } \mathbf{x} = \mathbf{F}\mathbf{u} + \mathbf{G}\mathbf{w}, \\ &\text{Infinite-horizon: } \mathsf{x} = \mathcal{F}\mathsf{u} + \mathcal{G}\mathsf{w}. \end{aligned} \tag{3}$$

where $(\mathcal{F}, \mathcal{G})$ denote strictly causal (strictly lower triangular) bi-infinite block Toeplitz operators and $(\mathbf{F}, \mathbf{G})$ represent their finite horizon equivalents, for a horizon $T > 0$. Employing this notation, we succinctly express the LQR cost in Eq. (2) as $\operatorname{cost}_T(\mathbf{u}, \mathbf{w}) := \|\mathbf{x}\|^2 + \|\mathbf{u}\|^2$.

b) Control: We focus on linear disturbance feedback control policies (DFC) which map disturbances to the control input: $\mathbf{u} = \mathbf{K}\mathbf{w}$, for any $\mathbf{K} \in \mathscr{K}_T$, where $\mathscr{K}_T$ is the set of *causal* (online) DFC policies in the finite-horizon of length $T > 0$. The infinite-horizon counterpart of our control policy is $\mathsf{u} = \mathcal{K}\mathsf{w}$ for any $\mathcal{K} \in \mathscr{K}$, with $\mathscr{K}$ the set of *causal and time-invariant* DFC policies in the infinite horizon.

Under a fixed control policy $\mathcal{K}$, the closed-loop transfer operator, $\mathcal{T}_\mathcal{K}$, which maps the disturbances to the state and control input, is defined as

$$\mathcal{T}_\mathcal{K} : \mathsf{w} \mapsto \begin{bmatrix} \mathsf{x} \\ \mathsf{u} \end{bmatrix} := \begin{bmatrix} \mathcal{F}\mathcal{K} + \mathcal{G} \\ \mathcal{K} \end{bmatrix} \mathsf{w}. \tag{4}$$

The finite-horizon counterpart of the closed-loop transfer operator (4) denoted as $\mathbf{T_K}$, is used to rewrite the quadratic cost (2) as

$$\operatorname{cost}_T(\mathbf{K}, \mathbf{w}) = \mathbf{w}^* \mathbf{T}_\mathbf{K}^* \mathbf{T}_\mathbf{K} \mathbf{w}. \tag{5}$$

C. Distributionally Robust LQR Control

We study the distributionally robust LQR control problem and seek to design a causal controller that minimizes the worst-case expected LQR cost when the probability distributions of the disturbances reside within a Wasserstein-2 (W_2) ambiguity set. The *Wasserstein-2 metric* between two distributions $\mathbb{P}_1, \mathbb{P}_2$ is defined as

$$\mathsf{W}_2(\mathbb{P}_1, \mathbb{P}_2) := \left(\inf_{\pi \in \Pi(\mathbb{P}_1, \mathbb{P}_2)} \int \|w_1 - w_2\|^2 \, \pi(dw_1, dw_2) \right)^{\frac{1}{2}},$$

with $\Pi(\mathbb{P}_1, \mathbb{P}_2)$ representing all the joint distributions with marginals $\mathbb{P}_1$ and $\mathbb{P}_2$ [27], [28].

We define the W_2-ambiguity set $\mathscr{W}_T(\mathbb{P}_\circ, r_T)$ for horizon $T > 0$ to be the W_2-ball with radius $r_T := r\sqrt{T} > 0$, centered at a nominal distribution $\mathbb{P}_\circ \in \mathscr{P}(\mathbb{R}^{pT})$:

$$\mathscr{W}_T(\mathbb{P}_\circ, r_T) := \left\{ \mathbb{P} \in \mathscr{P}(\mathbb{R}^{pT}) \,|\, \mathsf{W}_2(\mathbb{P}, \mathbb{P}_\circ) \le r_T \right\}. \tag{6}$$

Remark 1. *The choice of $r_T \propto \sqrt{T}$ aligns with the fact that the W_2-distance between two random vectors of length T, each sampled from two different iid processes, scales proportionally to $\sqrt{T}$.*

Unlike the usual LQR cost, which considers the expected quadratic cost under i.i.d Gaussian disturbances (or disturbances sampled from a single probability distribution), the DR-LQR considers the worst-case expected LQR cost across all disturbance probability distributions which lie in the W_2-ambiguity set.

Definition 1 (**Worst-case expected LQR cost under W_2-ambiguity**). *The worst-case expected LQR cost of a control policy $\mathbf{K} \in \mathscr{K}_T$, in the finite-horizon $T > 0$, is defined as*

$$C_T(\mathbf{K}, r_T) := \sup_{\mathbb{P} \in \mathscr{W}_T(\mathbb{P}_\circ, r_T)} \mathbb{E}_{\mathbb{P}}\left[\mathrm{cost}_T(\mathbf{K}, \mathbf{w})\right]. \tag{7}$$

Likewise, in the infinite-horizon, the worst-case expected LQR cost of a control policy $\mathcal{K} \in \mathscr{K}$ is defined as

$$C(\mathcal{K}, r) := \lim_{T \to \infty} \frac{1}{T} C_T([\mathcal{K}]_T, r_T), \tag{8}$$

where $[\mathcal{K}]_T$ results from restricting $\mathcal{K}$ to a horizon $T > 0$. In these definitions, $\mathbb{E}_{\mathbb{P}}$ is the expectation over the disturbances $\mathbf{w}$ which are sampled from $\mathbb{P}$, i.e., $\mathbf{w} \sim \mathbb{P}$.

We formally state the infinite-horizon Distributionally Robust LQR problem as follows:

Problem 1 (**Distributionally robust LQR control in the infinite-horizon**). *Minimize the time-averaged worst-case expected cost* (8) *as $T \to \infty$, over all causal and time-invariant controllers $\mathcal{K} \in \mathscr{K}$, i.e.,*

$$\inf_{\mathcal{K} \in \mathscr{K}} C(\mathcal{K}, r) = \inf_{\mathcal{K} \in \mathscr{K}} \lim_{T \to \infty} \frac{1}{T} C_T([\mathcal{K}]_T, r_T). \tag{9}$$

In section III, we provide an equivalent formulation of Problem 1 by establishing strong duality for the worst-case expected cost $C(\mathcal{K}, r)$.

III. Main Theoretical Results

In this section, our initial step involves reformulating Problem 1 through the lens of strong duality, breaking it down into a task of addressing a suboptimal problem. We proceed to characterize the controller using the Karush-Kuhn-Tucker (KKT) conditions and present arguments to show that it is stabilizing.

For simplicity, we assume that the nominal disturbance is uncorrelated, *i.e.* $\mathcal{M}_\circ := \mathbb{E}_{\mathbb{P}_\circ}[\mathsf{w}_\circ \mathsf{w}_\circ^*] = \mathcal{I}$.

Theorem 1 (**Strong duality**). *The distributionally robust LQR control problem* (9) *is equivalent to the dual optimization problem:*

$$\inf_{\substack{\mathcal{K} \in \mathscr{K} \\ \gamma \geq 0}} \sup_{\mathcal{M} \succ 0} \mathrm{Tr}(\mathcal{T}_\mathcal{K}^* \mathcal{T}_\mathcal{K} \mathcal{M}) + \gamma\left(r^2 - \mathrm{Tr}\left(\mathcal{M} - 2\sqrt{\mathcal{M}} + \mathcal{I}\right)\right) \tag{10}$$

Additionally, with a fixed $\mathcal{K}$, the worst-case disturbance $\mathsf{w}_\star$ can be given using the nominal disturbance $\mathsf{w}_\circ$ as $\mathsf{w}_\star = (\mathcal{I} - \gamma_\star^{-1} \mathcal{T}_\mathcal{K}^* \mathcal{T}_\mathcal{K})^{-1} \mathsf{w}_\circ$ with $\gamma_\star$ satisfying:

$$\mathrm{Tr}\left((\mathcal{I} - \gamma_\star^{-1} \mathcal{T}_\mathcal{K}^* \mathcal{T}_\mathcal{K})^{-1} - \mathcal{I}\right)^2 = r^2. \tag{11}$$

Proof. This theorem leverages the proof of Theorem 5 and Lemma 8 from [24], adapted here to incorporate our LQR cost, $\mathcal{T}_\mathcal{K}^* \mathcal{T}_\mathcal{K}$. □

This insight reveals that the essence of distributional robustness can be distilled into the analysis of the worst-case power spectrum of the disturbances, $\mathcal{M}$, which deviates by no more than $r > 0$ from the baseline nominal disturbance spectrum $\mathcal{M}_\circ = \mathcal{I}$.

Moreover, for any chosen $r > 0$, a corresponding optimal $\gamma > 0$ exists that facilitates the computation of the optimal controller. Given the feasibility of searching across a singular parameter $\gamma > 0$, our focus shifts towards the γ-optimal problem once γ is fixed.

Let $\Delta^* \Delta = \mathcal{I} + \mathcal{F}^* \mathcal{F}$ be the spectral factorization with causal Δ, and Δ^{-1}. And, let $\mathcal{K}_\circ := -(\mathcal{I} + \mathcal{F}^* \mathcal{F})^{-1} \mathcal{F}^* \mathcal{G}$ be the unique optimal non-causal policy which minimizes the infinite-horizon cost $\lim_{T \to \infty} \frac{1}{T} \mathrm{cost}_T(\mathbf{K}, \mathbf{w})$ [29], with $\mathcal{T}_{\mathcal{K}_\circ}$ its associated closed-loop transfer operator (4).

With the above, we can now give the saddle point conditions for the controller $\mathcal{K}$ and the disturbance covariance $\mathcal{M}$ in Theorem 2, for a fixed $\gamma \geq 0$:

Theorem 2 (**γ-optimal solution via saddle points**). *Let $\gamma > \gamma_{H_\infty} := \inf_{\mathcal{K} \in \mathscr{K}} \|\mathcal{T}_\mathcal{K}^* \mathcal{T}_\mathcal{K}\|_\infty$ be fixed. The γ-optimal LQR control problem in Theorem 1 for fixed γ is equivalent to the following dual problem:*

$$\sup_{\mathcal{M} \succ 0} \inf_{\mathcal{K} \in \mathscr{K}} \mathrm{Tr}(\mathcal{T}_\mathcal{K}^* \mathcal{T}_\mathcal{K} \mathcal{M}) - \gamma\, \mathrm{Tr}\left(\mathcal{M} - 2\sqrt{\mathcal{M}} + \mathcal{I}\right), \tag{12}$$

and the unique saddle point $(\mathcal{K}_\gamma, \mathcal{M}_\gamma)$ satisfies the equations below uniquely:

$$\mathcal{K}_\gamma = \Delta^{-1}\{\Delta \mathcal{K}_\circ \mathcal{L}_\gamma\}_+ \mathcal{L}_\gamma^{-1} \tag{13a}$$

$$\mathcal{L}_\gamma^* \mathcal{L}_\gamma = \frac{1}{4}\left(\mathcal{I} + \sqrt{\mathcal{I} + 4\gamma^{-1}(\mathcal{S}_{\mathcal{L}_\gamma}^* \mathcal{S}_{\mathcal{L}_\gamma} + \mathcal{U}_{\mathcal{L}_\gamma}^* \mathcal{U}_{\mathcal{L}_\gamma})}\right)^2 \tag{13b}$$

where $\mathcal{L}_\gamma \mathcal{L}_\gamma^* = \mathcal{M}_\gamma$ with causal $\mathcal{L}_\gamma$, and $\mathcal{L}_\gamma^{-1}$, and where $\mathcal{S}_{\mathcal{L}_\gamma}$ and $\mathcal{U}_{\mathcal{L}_\gamma}$ are defined as $\mathcal{S}_{\mathcal{L}_\gamma} := \{\Delta \mathcal{K}_\circ \mathcal{L}_\gamma\}_-$, $\mathcal{U}_{\mathcal{L}_\gamma} := \mathcal{T}_{\mathcal{K}_\circ} \mathcal{L}_\gamma$.

Proof. The proof is built upon the KKT conditions for (12) and the Wiener-Hopf technique [30]. Details of the proof are shown in the extended version of the paper [31, Appendix II]. □

Note that fixing γ results in the controller $\mathcal{K}_\gamma$ being *sub-optimal*. Theorem 2 hints at a way to obtain the optimal controller $\mathcal{K}_\gamma$ from a positive operator $\mathcal{M}_\gamma$ (13a). Since the optimality conditions on $\mathcal{M}_\gamma$ can be solely expressed in terms of its spectral factors and system-specific parameters as in (13b), we will shift our focus to obtaining a solution $\mathcal{M}_\gamma$, which is *the worst-case time-invariant covariance operator.*

Remark 2. *Denoting $\mathcal{N}_\gamma := \mathcal{L}_\gamma^* \mathcal{L}_\gamma$, we note that there is a one-to-one correspondence between $\mathcal{M}_\gamma$ and $\mathcal{N}_\gamma$ through* *spectral factorization. As the optimality conditions in Theorem 2 are stated in terms of $\mathcal{L}_\gamma^*\mathcal{L}_\gamma$, we call both $\mathcal{N}_\gamma$ and $\mathcal{M}_\gamma$ as the γ-optimal solution, interchangeably.*

Remark 3. *When r approaches infinity, $\gamma_\star$ approaches the lower bound $\|\mathcal{T}_\mathcal{K}^*\mathcal{T}_\mathcal{K}\|_\infty$, which corresponds to the worst-case cost, or in other words, the H_∞ cost. In this scenario, the optimal DR-LQR controller transitions into the traditional H_∞ controller. On the flip side, as r decreases to zero, $\gamma_\star$ goes to infinity. This shift results in the worst-case expected LQR cost, $C(\mathcal{K}_{\gamma_\star}, r)$, aligning with the anticipated cost under nominal disturbance, $\mathcal{M}_\circ$, thereby allowing the optimal DR-LQR controller to embody the conventional LQR (H_2) controller, particularly when $\mathcal{M}_\circ = \mathcal{I}$. Thus, adjusting r enables the DR-LQR controller to navigate a spectrum between H_∞ and H_2 control paradigms.*

This insight reveals that once γ surpasses $\gamma_{H\infty}$, the expected worst-case LQR cost becomes finite. This observation leads to the formulation of Corollary 3.

Corollary 3. *For any chosen γ value exceeding γ_{H_∞}, the resulting suboptimal controller, $\mathcal{K}_\gamma$, ensures the stabilization of the system's dynamics.*

IV. An Algorithm for Controller Synthesis in the Frequency Domain

In this section, we assert that the sub-optimal DR-LQR controller is *non-rational*, thereby precluding a finite-dimensional state-space realization. Consequently, we introduce a fixed-point iteration scheme aimed at computing the saddle-point solution $(\mathcal{K}_\gamma, \mathcal{N}_\gamma)$ for the γ-optimal problem outlined in Theorem 2, with fixed γ. The optimal $\gamma_\star$ and its associated saddle-point solution $(\mathcal{K}_{\gamma_\star}, \mathcal{N}_{\gamma_\star})$, can be effectively determined by employing the bisection method on equation (11). Finally, we prove the convergence of the fixed-point method in the scalar system case, *i.e.* $p = d = n = 1$, with the proof provided in the extended version of the paper [31, Appendix II].

A. A Fixed-Point Characterization of the Controller

Let the frequency domain counterparts of the operators $\mathcal{S}, \mathcal{U}_L$, and $\mathcal{N}$ be $S_{L_\gamma}(z) := \{\Delta K_\circ L_\gamma\}_{-}(z)$, $U_{L_\gamma}(z) := T_{K_\circ}(z)L_\gamma(z)$ and $N_\gamma(z) := L_\gamma(z)^* L_\gamma(z)$. We use the fact that $\{\mathcal{Y}\}_+ = \mathcal{Y} - \{\mathcal{Y}\}_-$ to express the KKT equations (13) in the frequency domain as:

$$K_\gamma(z) = K_\circ(z) - \Delta^{-1}(z) S_{L_\gamma}(z) L_\gamma^{-1}(z), \tag{14a}$$

$$N_\gamma(z) = \frac{1}{4}\left(I + \sqrt{I + 4\gamma^{-1}(S_{L_\gamma}^*(z) S_{L_\gamma}(z) + U_{L_\gamma}^*(z) U_{L_\gamma}(z))}\right)^2 \tag{14b}$$

The anticausal transfer function $S_{L_\gamma}(z)$ has the following state-space form representation (as shown in [24]): $S_{L_\gamma}(z) := \overline{C}(z^{-1}I - \overline{A})^{-1}\overline{B}_{L_\gamma}$ where we define $\overline{B}_{L_\gamma} := \frac{1}{2\pi}\int_0^{2\pi}(I - e^{j\omega}\overline{A})^{-1}\overline{D}L_\gamma(e^{j\omega})d\omega$, and the state-space parameters $(\overline{A}, \overline{C}, \overline{D})$ depend on the system parameters (A, B, C) (see Appendix I-A for the full definitions).

Given the notation above, we introduce the following theorem, characterizing the γ-optimal solution as a fixed point of a mapping.

Theorem 4 (**$\gamma-$optimal solution is a fixed-point solution**). *For a fixed $\gamma > \gamma_{H_\infty}$, consider the following set of mappings:*

$$F_1 : L(z) \mapsto \overline{B}_L := \frac{1}{2\pi}\int_0^{2\pi}(I - z\overline{A})^{-1}\overline{D}L(e^{j\omega})d\omega,$$
$$F_{2,\gamma} : (\overline{B}_L, L(z)) \mapsto N(z),$$
$$N(z) := \frac{1}{4}\left(I + \sqrt{I + 4\gamma^{-1}(S_L^*(z)S_L(z) + U_L(z)^*U_L(z))}\right),$$
with $S_L(z) = \overline{C}(z^{-1}I - \overline{A})^{-1}\overline{B}_L,\ U_L(z) = T_{K_\circ}(z)L(z),$
$$F_3 : N(z) \mapsto L(z),$$

where F_3 returns a unique spectral factor of $N(z) > 0$. The composition $F_3 \circ F_{2,\gamma} \circ (F_1 \times id) : L(z) \mapsto L(z)$ admits a unique fixed-point $L(z)$, and the positive transfer matrix $N_\gamma(z) := F_{2,\gamma} \circ (F_1 \times id)(L_\gamma(z))$ satisfies the KKT conditions (14).

Proof. The proof is similar to the proof of Theorem 13 in [24]. It utilizes the concavity of the problem in $\mathcal{M}_\gamma$ to argue for the uniqueness of $\mathcal{M}_\gamma$, and thus of its spectral factor $\mathcal{L}_\gamma$ up to a unitary transformation, which leads to a unique fixed-point. □

Subsequently, we argue that $N_\gamma(z)$ is non-rational. Note that $S_{L_\gamma}(z)$ is rational, and assuming $L_\gamma(z)$ is rational implies that $U_{L_\gamma}(z)$ is rational. Hence, $N_\gamma(z)$ involves the square root of a rational term. As square root does not generally preserve rationality, both $N_\gamma(z)$ and its spectral factor $L_\gamma(z)$ are non-rational, leading to Corollary 5.

Corollary 5. *$N_\gamma(z)$ and the suboptimal DR-LQR controller, $K_\gamma(z)$, are non-rational, for any fixed $\gamma > \gamma_{H_\infty}$. Hence, $K_\gamma(z)$ does not have a finite-dimensional state-space representation.*

Despite the fact that $K_\gamma(z)$ does not lend itself to a finite-dimensional state-space form, Theorem 4 affirms that the suboptimal controller $K_\gamma(z)$ (14a) can be derived from $L_\gamma(z)$, by executing a fixed-point iteration on $L_\gamma(z)$ as elucidated in Section IV-B.

B. Algorithm Description

In light of Theorem 4, we propose Algorithm 1 to compute the suboptimal controller $K_\gamma(z)$ at uniformly sampled points on the unit circle, $\mathbb{T}_N := \{e^{j2\pi n/N} \mid n = 0, \ldots, N-1\}$. With an initial estimate $L_\gamma^{(0)}(z)$, `Fixed-Point` iteratively computes the n-th step as $N_\gamma^{(n)}(z) = F_{2,\gamma} \circ (F_1 \times \mathrm{id})(L_\gamma^{(n)}(z))$.

Following this, we compute the spectral factor $L_\gamma^{(n+1)}(z)$ at regularly spaced points along the unit circle using the `SpectralFactor` algorithm. Upon reaching convergence within a predetermined tolerance at the N-th iteration, we determine the suboptimal $N_\gamma^{(N)}(z)$ from which we derive the *suboptimal* controller $K_\gamma^{(N)}(z)$ at each sampled frequency point using (14a).

Algorithm 1 `Fixed-Point`

Input: $\gamma>\gamma_{H_\infty}$, system $(\overline{A},\overline{C},\overline{D})$, discretization N
Initialize $L_\gamma^{(0)}(z), \forall z\in\mathbb{T}_N=\{\mathrm{e}^{j2\pi n/N} \mid n=0,..,N-1\}$
repeat
Compute $\overline{B}_{L_\gamma}^{(n)} = F_1(L_\gamma^{(n)}(z))$ numerically

$$\overline{B}_{L_\gamma}^{(n)} \leftarrow \frac{1}{N}\sum_{z\in\mathbb{T}_N}(I-z\overline{A})^{-1}\overline{D}L_\gamma^{(n)}(z)$$

Compute $N_\gamma^{(n)}(z) \leftarrow F_{2,\gamma}(\overline{B}_{L_\gamma}^{(n)}, L_\gamma^{(n)}(z))$
Get $L_\gamma^{(n+1)}(z) \leftarrow$ `SpectralFactor`$(N_\gamma^{(n)}(z))$
Update $n \leftarrow n+1$
until convergence of $N_\gamma^{(n)}(z)$

Note that this approach can assess $N_\gamma^{(n)}(z)$ for any arbitrary point on the unit circle via Eq. (14b), however, the `SpectralFactor` algorithm computes $L_\gamma^{(n)}(z)$ and $K_\gamma^{(n)}(z)$ only for a discrete number of samples on the unit circle. The details of the `SpectralFactor` algorithm, which is based on discrete Fourier transform (DFT) [32] and tailored for *non-rational* spectra, are in [31, Appendix V].

V. Near-optimal State-Space Realization via Rational Approximation

A. Rational Approximation

Our objective is to find the optimal m^{th} order rational approximation, denoted as $P(z)/Q(z)$, for the positive non-rational function $N(z)$. This rational approximation serves as the basis for obtaining the spectral factor $L(z)$ as noted in Lemma 7 and thus deriving the controller $K(z)$ (14a). We provide results for scalar disturbances, *i.e.* $p=1$, while the states and control inputs can be arbitrarily dimensional vectors. We leave the generalization to vector disturbances (*i.e.*, $p>1$) for future work.

Problem 2 (**Rational approximation using H_∞ norm minimization**). *Given a positive non-rational function $N(z)$, find the best rational approximation of order at most $m\in\mathbb{N}$ with respect to H_∞-norm:*

$$\inf_{p_0,\dots,p_m,q_0,\dots,q_m\in\mathbb{R}} \left\|\frac{P(z)}{Q(z)}-N(z)\right\|_\infty, \tag{15}$$

with $P(z)=\sum_{k=-m}^{m}p_kz^{-k}$, $p_k=p_{-k}\in\mathbb{R}$, and $P(z)>0$, (and similarly for $Q(z)$).

To solve Problem 2 using standard convex optimization tools, we follow the approach in [33] and consider instead the sublevel sets of the objective function (15) and reduce the problem to a convex feasibility problem.

Lemma 6 (**Rational approximation using a convex feasibility problem**). *Fixing a minimum level $\epsilon>0$, Problem 2 can be relaxed to a convex problem:*

$$\text{Find } \mathbf{p}=(p_0,p_1,\dots,p_m),\ \mathbf{q}=(q_0,q_1,\dots,q_m)$$
$$\text{s.t } P(z),Q(z)\geq 0,\ \max_{z\in\mathbb{T}}\left|\frac{P(z)}{Q(z)}-N(z)\right|\leq\epsilon,$$

or equivalently

$$\text{s.t }\begin{cases} P(z)-(N(z)+\epsilon)Q(z)\leq 0, & \forall z\in\mathbb{T}\\ P(z)-(N(z)-\epsilon)Q(z)\geq 0, & \forall z\in\mathbb{T}\\ P(z),Q(z)\geq 0, \quad \forall z\in\mathbb{T}\end{cases}$$

Although the inequalities in Lemma 6 are infinitely many, we can check these inequalities solely for a finite set of frequencies, such as $\mathbb{T}_N=\{\mathrm{e}^{j2\pi n/N}\mid n=0,\dots,N-1\}$ for $N\gg m$. In fact, the finite polynomials $P(z)$ and $Q(z)$ can be fully characterized with $N\geq 2m$ number of uniformly sampled frequencies on the unit circle by Nyquist sampling theorem. By increasing the number of samples, the accuracy of this method can be improved to any desired fidelity.

Once obtained a rational approximation $P(z)/Q(z)$ for $N(z)$, we can find the rational canonical factor $L(z)$ of $P(z)/Q(z)$ from the following lemma.

Lemma 7 (**Canonical factorization [34]**). *Given a polynomial of order m, $R(z)=\sum_{k=-m}^{m}r_kz^{-k}$, where $r_k=r_{-k}\in\mathbb{R}$, and $R(z)>0$, a causal canonical factor, $L(z)=\sum_{k=0}^{m}l_kz^{-k}$, exists and satisfies $R(z)=|L(z)|^2$.*

B. Controller in Time-Domain

Once we have the rational approximation of $L(z)$, we can write it in the form $L(z)=(I+\tilde{C}(zI-\tilde{A})^{-1}\tilde{B})\tilde{D}^{1/2}$, and we can compute the DR-LQR controller in state-space form.

Lemma 8 (**DR-LQR control in state-space form**). *Given a rational spectral factor $L(z)=(I+\tilde{C}(zI-\tilde{A})^{-1}\tilde{B})\tilde{D}^{1/2}$, the near-optimal DR-LQR controller can be realized by the following state-space controller*

$$e(t+1)=\tilde{F}e(t)+\tilde{G}w(t), \tag{16}$$
$$u(t)=\tilde{H}e(t)+\tilde{J}w(t), \tag{17}$$

where $(\tilde{F},\tilde{G},\tilde{H},\tilde{J})$ are functions of the matrices $(A,B,C,\tilde{A},\tilde{B},\tilde{C})$ (see Appendix I-B for the details of the appropriate definitions, and [31, Appendix IV] for the complete proof).

VI. Numerical Simulations

This section presents a comparative evaluation of the DR-LQR controller vis-à-vis H_2 and H_∞ controllers alongside the finite-horizon DR-LQR counterpart. Our evaluation encompasses both frequency-domain and time-domain assessments, which showcase the efficacy of the rational approximation method. Our analysis focuses on benchmark models from [35] such as [AC15], [REA4] and [HE3]. Given the similarity in controller performance across all systems, we opt to present results solely for [AC15], a four-state aircraft model, due to space limitations. We choose our nominal distribution to be Gaussian, with zero mean and identity covariance.

A. Frequency Domain Evaluations

We examine the dynamics of the DR-LQR controller and its rational approximation across varying radii r. The power spectrum $N(e^{j\omega})$ of the worst-case disturbance is illustrated for three distinct r values for the [AC15] system in Figure 1. Notably, for $r = 0.01$, the worst-case disturbance exhibits near-white behavior, consistent with the nominal disturbance. However, as r increases, the temporal correlation of the worst-case disturbance intensifies, leading to a more pronounced peak in the power spectrum.

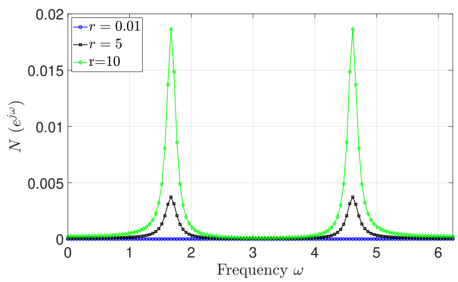

Fig. 1: The power spectrum $N(e^{j\omega})$ of the worst-case disturbance, when $r \in (0.01, 5, 10)$ for system [AC15].

Figures 2 and 3 illustrate the worst-case expected LQR cost for the DR-LQR, H_2, and H_∞ controllers applied to the [AC15] system. As r varies, the performance of the DR-LQR closely mirrors that of the H_2 for smaller values of r. However, with increasing r, the worst-case LQR cost tends to align more closely with that of the H_∞ controller. Across all ranges of r, the DR-LQR consistently outperforms the other controllers and achieves the lowest worst-case expected cost. Additionally, we consider another

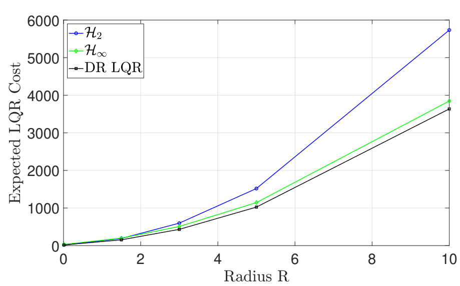

Fig. 2: The worst-case expected LQR cost of the classical controllers H_2, H_∞ compared to the DR-LQR, for the system [AC15], for different r values. The DR-LQR minimizes the cost at all $r's$.

performance metric—the operator norm of $\mathcal{T}_\mathcal{K}$ minimized by the H_∞ controller. This metric, expressed in the frequency domain as $\|\mathcal{T}_\mathcal{K}\|_{op}^2 = \max_{0\le\omega\le 2\pi} \sigma_{\max}(T_K^*(e^{j\omega})T_K(e^{j\omega}))$, is depicted across all frequencies in Figure 4. The results show that the DR controller consistently interpolates the H_2 and H_∞ controllers across all frequencies.

To practically implement the DR-LQR controller, we find the rational controller by employing the method outlined in Section V, from which we obtain the rational approximation

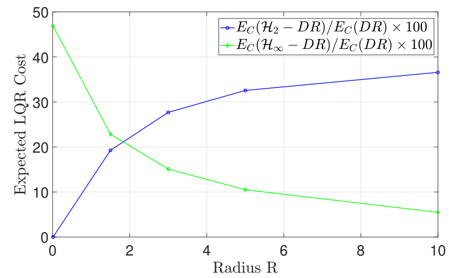

Fig. 3: The percentage difference in the worst-case LQR cost relative to the DR-LQR (see the legend) for the system [AC15], for different values of r. When r is small (large) r, the cost of DR-LQR controller closely aligns with that of H_2 (H_∞). For $r = 1.5$, the cost of the DR-LQR is less than that of H_2 by 22.8%, and that of H_∞ by 19.2% .

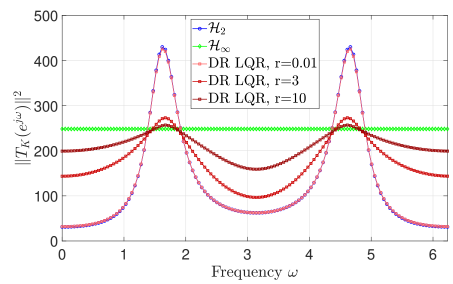

Fig. 4: The operator norm, $\|T_K^*(e^{j\omega})T_K(e^{j\omega})\|$, of different controllers at all frequencies $\omega \in [0, 2\pi]$, for system [AC15]. The DR-LQR cost interpolates between H_2 and H_∞ based on the value of r. When r is small (large), DR closely aligns with H_2 (H_∞) *across all frequencies.*

of $N(e^{j\omega})$ as $\frac{P(e^{j\omega})}{Q(e^{j\omega})}$ with degrees $m = 1, 2$, for the [AC15] system. Table I compares the performance of these resulting rational controllers to the non-rational DR-LQR. Notably, the rational approximation with an order of 2 achieves an expected LQR cost that closely matches that of the non-rational controller with a difference of less than 1% in costs for all r values.

	r=0.01	**r=1.5**	**r=5**	**r=10**
DR-LQR	**17.2**	**153.2**	**1024**	**3635.6**
RA(1)	17.2	5789.1	6214.6	33262
RA(2)	**17.19**	**153.2**	**1024.1**	**3645.8**

TABLE I: The worst-case expected LQR cost of the non-rational DR-LQR controller, compared to the rational controllers RA(1), and RA(2), obtained from degree 1, and 2 rational approximations to $N(e^{j\omega})$.

Finally, Figure 5 shows the convergence ratio defined as:

$$\texttt{Convergence Ratio} := \frac{\mathsf{BW}(\mathcal{M}_1^{i+1}, \mathcal{M}_2^{i+1})}{\mathsf{BW}(\mathcal{M}_1^{i}, \mathcal{M}_2^{i})} \tag{18}$$

where $\mathsf{BW}(\mathcal{M}_1^i, \mathcal{M}_2^i)$ represents the Bures-Wasserstein distance [36] between positive operators $\mathcal{M}_1$ and $\mathcal{M}_2$ at iteration i. The plot shows that the iterates $\{\mathcal{M}_1^i\}_{i\ge 0}$ converge and the rate of convergence of the `Fixed-Point` is exponential.

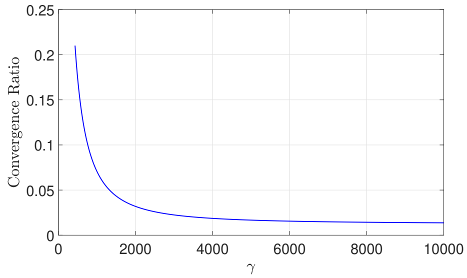

Fig. 5: Convergence ratio (18) for different values of γ. The `Fixed-Point` algorithm converges at an exponential rate.

B. Time Domain Evaluations

We compare the performance of the infinite horizon DR-LQR controller to its finite horizon counterpart, *across time*. The finite-horizon DR-LQR, presented in [21], is computed by solving an SDP whose dimension scales with the time horizon. We graph the mean LQR cost across 210 time steps, consolidating data from 1000 separate trials.The performance of DR controllers under white Gaussian noise and the worst-case disturbances of the infinite-horizon and the finite horizon DR controllers is shown in Figures 6a, 6b and 6c, respectively. For the sake of computational efficacy, the finite horizon controller operates within a constrained time horizon of $s = 30$ steps, being recurrently applied every s steps. Likewise, the worst-case disturbances used in Figures 6b,6c are generated at the same periodicity, resulting in correlated disturbances solely within each s step interval.

Our investigations underscore the unparalleled performance of the infinite-horizon DR-LQR controller across all three scenarios. Note that attempting to extend the horizon of the SDP for prolonged durations to approximate the infinite-horizon performance proves to be excessively computationally intensive. These findings accentuate the inherent advantages of adopting the infinite-horizon DR-LQR controller.

VII. Future Work

Our study proposes several overarching directions for future investigation. Firstly, we aim to prove the convergence of the fixed-point algorithm for the general case of non-scalar systems. Moreover, we plan to generalize the rational approximation method for non-scalar disturbances. Additionally, we seek to expand our findings to partially observed, *i.e.*, measurement feedback systems.

Appendix I: Definitions

A. Parameters Definitions

We define $\overline{A}, \overline{D}$ and $\overline{C}$ as: $\overline{A} := A_K^*$, $\overline{D} := A_K^* P B_w$, and $\overline{C} := -(R + B_u^* P B_u)^{-*/2} B_u^*$ where: i)A_K is the closed loop matrix $A_K := A - B_u K_{\text{lqr}}$, ii) K_{lqr} is the LQR controller $K_{\text{lqr}} := (R + B_u^* P B_u)^{-1} B_u^* P A$ and iii) $P \succ 0$ is the unique stabilizing solution to the LQR Riccati equation $P = Q + A^* P A - A^* P B_u (R + B_u^* P B_u)^{-1} B_u^* P A$.

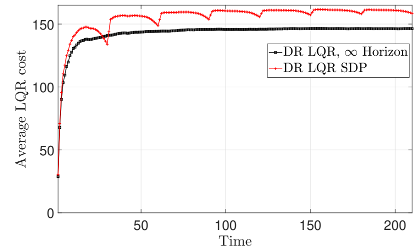

(a) Control costs of DR-LQR controllers in the infinite (I) and finite horizon (II) under white noise.

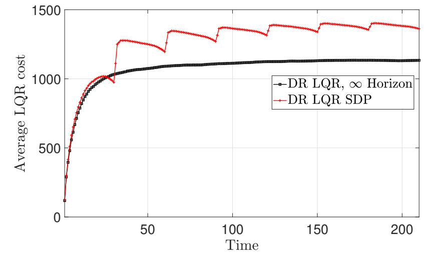

(b) Control costs of DR-LQR controllers in the infinite (I) and finite horizon (II), under the worst-case disturbance for (I).

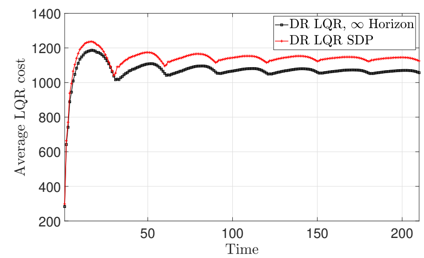

(c) Control costs of DR-LQR controllers in the infinite (I) and finite horizon (II), under the worst-case disturbances for (II).

Fig. 6: Control costs of DR-LQR controllers in the infinite (I) and finite horizon (II) which is solved using SDP. The finite-horizon controller is recurrently applied at intervals of $s = 30$ steps, and the radius of its ball is $r = 1.5\sqrt{s}$, while the radius of (I) is $r = 1.5$. The infinite horizon DR-LQR controller outperforms its finite-horizon counterpart and attains the minimum average cost in all cases, even when the finite horizon DR-LQR is designed to minimize the cost.

B. Equations for the rational controller

We give the equations for $\tilde{F}, \tilde{G}, \tilde{H}$ and $\tilde{J}$ of lemma 8.

$$\tilde{F} = \begin{bmatrix} \tilde{A_K} & 0 \\ B_u \bar{R}^* \bar{R} B_u^* & A_k \end{bmatrix},$$
$$\tilde{G} = \begin{bmatrix} A_K \tilde{B} \\ -B_w + B_u \bar{R}^* \bar{R} B_u^* (P B_w + U \tilde{B}) \end{bmatrix},$$
$$\tilde{H} = -R^{1/2} (\begin{bmatrix} \bar{R}^* \bar{R} B_u^* & -K_{lqr} \end{bmatrix}),$$
$$\tilde{J} = -R^{1/2} (\bar{R}^* \bar{R} B_u^* (P B_w + U \tilde{B}).$$

Here, i) K_{lqr}, A_K and P are as defined in Appendix I-A, ii) $\bar{R} = (R + B_u^* P B_u)^{-*/2}$ iii) $\tilde{A}_k = \tilde{A} - \tilde{B}\tilde{C}$ where $\tilde{A}, \tilde{B}, \tilde{C}$ as in lemma 8 and iv) U satisfies the lyapunov equation $A_k^* P B_w \tilde{C} + A_k^* U A = U$.

References

[1] P. M. E. M. van der Grinten, "Uncertainty in measurement and control," *Statistica Neerlandica*, vol. 22, no. 1, pp. 43–63, 1968.

[2] J. C. Doyle, "Structured uncertainty in control system design," in *1985 24th IEEE Conference on Decision and Control*, 1985, pp. 260–265.

[3] S. Samuelson and I. Yang, "Data-driven distributionally robust control of energy storage to manage wind power fluctuations," in *2017 IEEE Conference on Control Technology and Applications (CCTA)*, 2017, pp. 199–204.

[4] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, Mar. 1960.

[5] G. Zames, "Feedback and optimal sensitivity: Model reference transformations, multiplicative seminorms, and approximate inverses," *IEEE Transactions on Automatic Control*, vol. 26, no. 2, pp. 301–320, Apr. 1981. [Online]. Available: http://ieeexplore.ieee.org/document/1102603/

[6] J. Doyle, K. Glover, P. Khargonekar, and B. Francis, "State-space solutions to standard h_2 and h_∞ control problems," in *1988 American Control Conference*. Atlanta, GA, USA: IEEE, Jun. 1988, pp. 1691–1696. [Online]. Available: https://ieeexplore.ieee.org/document/4789992/

[7] B. Hassibi, A. H. Sayed, and T. Kailath, *Indefinite-Quadratic Estimation and Control*. Society for Industrial and Applied Mathematics, 1999. [Online]. Available: https://epubs.siam.org/doi/abs/10.1137/1.9781611970760

[8] K. Zhou, J. C. Doyle, and K. Glover, *Robust and optimal control*. Upper Saddle River, N.J: Prentice Hall, 1996.

[9] I. Yang, "Wasserstein distributionally robust stochastic control: A data-driven approach," *IEEE Transactions on Automatic Control*, vol. 66, no. 8, pp. 3863–3870, 2020.

[10] K. Kim and I. Yang, "Distributional robustness in minimax linear quadratic control with Wasserstein distance," Feb. 2021, arXiv:2102.12715 [cs, eess, math]. [Online]. Available: http://arxiv.org/abs/2102.12715

[11] A. Hakobyan and I. Yang, "Wasserstein distributionally robust control of partially observable linear systems: Tractable approximation and performance guarantee," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 4800–4807.

[12] B. Taskesen, D. A. Iancu, C. Kocyigit, and D. Kuhn, "Distributionally robust linear quadratic control," *arXiv preprint arXiv:2305.17037*, 2023.

[13] L. Aolaritei, M. Fochesato, J. Lygeros, and F. Dörfler, "Wasserstein Tube MPC with Exact Uncertainty Propagation," Apr. 2023, arXiv:2304.12093 [math]. [Online]. Available: http://arxiv.org/abs/2304.12093

[14] L. Aolaritei, N. Lanzetti, and F. Dörfler, "Capture, Propagate, and Control Distributional Uncertainty," Apr. 2023, arXiv:2304.02235 [math]. [Online]. Available: http://arxiv.org/abs/2304.02235

[15] J.-S. Brouillon, A. Martin, J. Lygeros, F. Dörfler, and G. F. Trecate, "Distributionally robust infinite-horizon control: from a pool of samples to the design of dependable controllers," 2023.

[16] I. Tzortzis, C. D. Charalambous, and T. Charalambous, "Dynamic Programming Subject to Total Variation Distance Ambiguity," Feb. 2014, arXiv:1402.1009 [math]. [Online]. Available: http://arxiv.org/abs/1402.1009

[17] I. Tzortzis, C. D. Charalambous, T. Charalambous, C. K. Kourtellaris, and C. N. Hadjicostis, "Robust Linear Quadratic Regulator for uncertain systems," in *2016 IEEE 55th Conference on Decision and Control (CDC)*. Las Vegas, NV, USA: IEEE, Dec. 2016, pp. 1515–1520. [Online]. Available: http://ieeexplore.ieee.org/document/7798481/

[18] R. Liu, G. Shi, and P. Tokekar, "Data-Driven Distributionally Robust Optimal Control with State-Dependent Noise," Aug. 2023, arXiv:2303.02293 [cs]. [Online]. Available: http://arxiv.org/abs/2303.02293

[19] P. Mohajerin Esfahani and D. Kuhn, "Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations," *Mathematical Programming*, vol. 171, no. 1-2, pp. 115–166, Sep. 2018. [Online]. Available: http://link.springer.com/10.1007/s10107-017-1172-1

[20] R. Gao and A. J. Kleywegt, "Distributionally Robust Stochastic Optimization with Wasserstein Distance," Apr. 2022, arXiv:1604.02199 [math]. [Online]. Available: http://arxiv.org/abs/1604.02199

[21] F. Al Taha, S. Yan, and E. Bitar, "A Distributionally Robust Approach to Regret Optimal Control using the Wasserstein Distance," in *2023 62nd IEEE Conference on Decision and Control (CDC)*, IEEE. IEEE, Dec. 2023, pp. 2768–2775. [Online]. Available: https://ieeexplore.ieee.org/document/10384311/

[22] J. Hajar, T. Kargin, and B. Hassibi, "Wasserstein distributionally robust regret-optimal control under partial observability," in *2023 59th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2023, pp. 1–6.

[23] Z. Zhong and J.-J. Zhu, "Nonlinear wasserstein distributionally robust optimal control," *arXiv preprint arXiv:2304.07415*, 2023.

[24] T. Kargin, J. Hajar, V. Malik, and B. Hassibi, "Wasserstein distributionally robust regret-optimal control over infinite-horizon," in *Proceedings of the 6th Annual Learning for Dynamics & Control Conference*, ser. Proceedings of Machine Learning Research, A. Abate, M. Cannon, K. Margellos, and A. Papachristodoulou, Eds., vol. 242. PMLR, 15–17 Jul 2024, pp. 1688–1701. [Online]. Available: https://proceedings.mlr.press/v242/kargin24a.html

[25] J. Hajar, O. Sabag, and B. Hassibi, "Regret-optimal control under partial observability," 2023. [Online]. Available: https://arxiv.org/abs/2311.06433

[26] O. Sabag, G. Goel, S. Lale, and B. Hassibi, "Regret-optimal lqr control," 2023. [Online]. Available: https://arxiv.org/abs/2105.01244

[27] C. Villani, *Optimal transport: old and new*, ser. Grundlehren der mathematischen Wissenschaften. Berlin: Springer, 2009, no. 338, oCLC: ocn244421231.

[28] F. Santambrogio, "Optimal transport for applied mathematicians," *Birkäuser, NY*, vol. 55, no. 58-63, p. 94, 2015.

[29] O. Sabag, G. Goel, S. Lale, and B. Hassibi, "Regret-optimal controller for the full-information problem," in *2021 American Control Conference (ACC)*. IEEE, 2021, pp. 4777–4782.

[30] T. Kailath, A. H. Sayed, and B. Hassibi, *Linear estimation*, ser. Prentice Hall information and system sciences series. Upper Saddle River, N.J: Prentice Hall, 2000.

[31] J. Hajar, T. Kargin, V. Malik, and B. Hassibi, "The distributionally robust infinite-horizon LQR," 2024, arXiv preprint arXiv:2408.06230.

[32] C. Rino, "Factorization of spectra by discrete Fourier transforms (Corresp.)," *IEEE Transactions on Information Theory*, vol. 16, no. 4, pp. 484–485, Jul. 1970. [Online]. Available: http://ieeexplore.ieee.org/document/1054502/

[33] T. Kargin, J. Hajar, V. Malik, and B. Hassibi, "Infinite-horizon distributionally robust regret-optimal control," in *Proceedings of the 41st International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, R. Salakhutdinov, Z. Kolter, K. Heller, A. Weller, N. Oliver, J. Scarlett, and F. Berkenkamp, Eds., vol. 235. PMLR, 21–27 Jul 2024, pp. 23 187–23 214. [Online]. Available: https://proceedings.mlr.press/v235/kargin24a.html

[34] A. H. Sayed and T. Kailath, "A survey of spectral factorization methods," *Numerical Linear Algebra with Applications*, vol. 8, no. 6-7, pp. 467–496, Sep. 2001. [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1002/nla.250

[35] F. Leibfritz and W. Lipinski, "Description of the benchmark examples in compleib 1.0," *Dept. Math, Univ. Trier, Germany*, vol. 32, 2003.

[36] R. Bhatia, T. Jain, and Y. Lim, "On the Bures-Wasserstein distance between positive definite matrices," Dec. 2017, arXiv:1712.01504 [math]. [Online]. Available: http://arxiv.org/abs/1712.01504