

Robust Linear Quadratic Regulator: Exact Tractable Reformulation

Wouter Jongeneel, Tyler Summers and Peyman Mohajerin Esfahani

Abstract—We consider the problem of controlling an unknown stochastic linear dynamical system subject to an infinite-horizon discounted quadratic cost. Existing approaches for handling the corresponding robust optimal control problem resort to either conservative uncertainty sets or various approximations schemes, and to our best knowledge, the current literature lacks an exact, yet tractable, solution. We propose a class of novel uncertainty sets for the system matrices of the linear system. We show that the resulting robust linear quadratic regulator problem enjoys a closed-form solution described through a generalized algebraic Riccati equation arising from dynamic game theory.

I. INTRODUCTION

A broad variety of problems from engineering, machine learning, and operations research involve optimizing the behaviour of a dynamical system in the face of inherent uncertainties in the system model used for design and decision-making. A vast literature going back several decades has studied various aspects of this robust control problem, including substantial work on system identification; adaptive, robust, and optimal control, *e.g.*, see [1]–[3].

In this work we consider the discrete-time Linear Quadratic Regulator (LQR) problem under parametric uncertainties. Ever since the LQR problem originated, robustness was questioned. It is known that the discrete-time LQR can suffer from the lack of a stability margin [4], or if any, it is typically a noticeably worse margin in comparison with the continuous-time counterpart [5]. Moreover, our understanding of the corresponding perturbation theory is limited [6], [7]. The inherent presence of uncertainties in practice indeed reinforces the need to address these issues. A classical μ -synthesis approach is generally intractable [8], [9] while a tractable LMI approach like proposed in [10] may be conservative. This work investigates to what extend dynamic game theory can be a middle-ground.

A. Related Work

This paper is centered around quantifying the robustness resulting from a dynamic game with quadratic cost and linear dynamics. Early accounts of this viewpoint can be found on for example page 90 of the monograph by Whittle [11]. There, the remark is made that extremizing a *risk-sensitive* multi-stage optimal control cost function can be interpreted as another, yet now *constrained*, optimal control problem.

The authors are with the Delft Center for Systems and Control, TU Delft, The Netherlands (P.MohajerinEsfahani@tudelft.nl, W.Jongeneel@student.tudelft.nl) and the Control, Optimization, and Networks lab, UT Dallas, The United States (tyler.summers@utdallas.edu). The work of T. Summers was sponsored by the Army Research Office and was accomplished under Grant Number: W911NF-17-1-0058.

There is a large body of work in this direction. The celebrated paper [12] provides necessary and sufficient conditions for the continuous-time system $\dot{x}(t) = (A + \Delta_A(t))x(t) + (B + \Delta_B(t))u(t)$, $(\Delta_A \quad \Delta_B) = DF(t)(E_1 \quad E_2)$, $\|F(t)\| \leq 1$ to be stabilizable. This result was later generalized to the discrete-time case in [13]. Although these results are more than 20 years old, describing parametric uncertainties in the pair (A, B) via some matrix-norm-balls is still the prevalent method, however currently driven by measure concentration results, *e.g.*, see [14], [15]. In the stochastic case, distributional uncertainties in the form of *relative entropy* constraints are considered [16], [17].

Although these problems are well understood, the catch within this game theoretic framework is that, the uncertainty set typically depends on the extremizing parameters. Therefore, it is not clear, *a priori*, over which set of models the robust control problem is solved, this is effectively only known *a posteriori*. Moreover, most results do not consider the full uncertainty set their optimization problem can handle, but rather focus on some “*inscribed ball*”, see [17, ch 10] on how to fit an ellipsoid to data. Motivated by renewed interest in tractable reformulations of (Robust) LQR problems (*cf.* [18]–[23]), we investigate which lessons can be drawn from the readily available dynamic game theory.

B. Contribution and Outline

This work focuses on a novel formulation and solution of a robust LQR problem. Our contributions are as follows:

- (i) We propose a novel family of uncertainty sets for the system matrices, and show that the worst-case cost over these sets can be solved efficiently (Proposition III.6).
- (ii) Given the proposed uncertainty sets, we develop an exact, up to an algebraic Riccati equation, solution to the corresponding Robust LQR problem (Theorem III.7).

The article is structured as follows. In Section II, we formally introduce several key definitions along with the robust optimal control problem that will be addressed. The new uncertainty set and the corresponding main results are presented in Section III, which are interpreted from a game theoretic point of view in Section IV. In Section V, we illustrate the presented results through a numerical example.

Due to the lack of space, we refer the interested readers to an extended version of this work [24]. This version contains the technical proofs along with new results and discussions, including an extension to uncertain input matrices and further structural properties of our uncertainty set and worst-case models.

Notation: We use standard notation, but to be clear. Let $\mathbb{R}_{\geq 0}$ denote the set of non-negative real numbers, whereas I_n is the identity element of $\mathbb{R}^{n \times n}$. Let \mathcal{S}_+^n be the cone of symmetric positive semi-definite matrices on which the ordering is denoted by $A \succeq B$. The largest singular-value of a matrix A equals $\|A\|_2$. Let $\text{Tr}(\cdot)$ be the trace operator, then the inner-product between $A \in \mathbb{R}^{m \times n}, B \in \mathbb{R}^{n \times m}$ is defined as $\langle A, B \rangle = \text{Tr}(A^\top B)$ such that $\langle A, A \rangle = \|A\|_F^2$ for $\|\cdot\|_F$ the Frobenius-norm. Similarly, $\|X\|_{F,Q}^2$ is used to denote $\text{Tr}(X^\top Q X)$ for $Q \succ 0$. Furthermore, when A is said to be *exponentially stable* its spectrum is fully contained in the open unit disk. The expectation operator is given by $\mathbb{E}[\cdot]$ and $X \sim \mathcal{P}(\mu, \Sigma)$ is a random variable with mean μ and covariance Σ for a distribution \mathcal{P} . Optimality is indicated with a \star , so x^\star is for example the minimizer of a function $f(x)$ with $f^\star = f(x^\star)$. Also, in the context of an optimization program, *s.t.* stands for *subject to*.

II. PRELIMINARIES

In this section the problem at hand is introduced.

A. Robust LQR problem

Given the matrices $Q \in \mathcal{S}_+^n, R \in \mathcal{S}_{++}^m$, discount factor $\alpha \in (0, 1)$ and $\hat{A} \in \mathbb{R}^{n \times n}, B \in \mathbb{R}^{n \times m}, \Sigma_0, \Sigma_v \in \mathcal{S}_{++}^n$, and $\{v_k\}_{k \in \mathbb{N}}$ being a white noise sequence of independent random variables with zero mean and a time-invariant covariance matrix Σ_v , *i.e.*, $\mathbb{E}[v_i] = 0$ and $\mathbb{E}[v_i v_j^\top] = \delta_{ij} \Sigma_v$ for all $i, j \in \mathbb{N}$. Then we seek an optimal policy $\pi^\star = \{\mu_0^\star, \mu_1^\star, \dots\}$ that solves the discounted *Robust Linear-Quadratic Regulator* (RLQR) problem over the uncertainty set Δ :

$$\begin{aligned} \inf_{\{\mu_k\}_{k=0}^\infty} \sup_{\Delta_A} \mathbb{E}_{x_0, v} \left[\sum_{k=0}^{\infty} \alpha^k (x_k^\top Q x_k + v_k^\top R v_k) \right], \\ \text{s.t. } x_{k+1} = (\hat{A} + \Delta_A) x_k + B u_k + v_k, \\ v_k \stackrel{i.i.d.}{\sim} \mathcal{P}(0, \Sigma_v), \quad x_0 \sim \mathcal{P}(0, \Sigma_0), \\ u_k = \mu_k(x_k), \quad \Delta_A \in \Delta. \end{aligned} \quad (1)$$

In other words, we consider the LQR problem where the system matrix A is not precisely known, but known to be described by $A = \hat{A} + \Delta_A$. Here our prior estimate of A is denoted by \hat{A} , whereas $\Delta_A \in \Delta$ is the uncertainty. A particular example of such a setting naturally emerges in statistics or identification problems where \hat{A} is the current estimate of A and Δ contains Δ_A with high probability.

See [24] for an extension to the case where also B is partially unknown.

Assumption II.1 (Linear time-invariant policy): In problem (1), we restrict the class of control policies μ_k to linear time-invariant (LTI) controllers $\mu_k(x) = Kx$ where $K \in \mathbb{R}^{m \times n}$.

Instead of writing the full program (1) over again, introduce a compact notation:

Definition II.2 (Discounted LQ cost): Consider the dynamical system $x_{k+1} = Ax_k + v_k$ where the noise process and the initial condition follow $v_k \stackrel{i.i.d.}{\sim} \mathcal{P}(0, \Sigma_v)$ and $x_0 \sim \mathcal{P}(0, \Sigma_0)$.

Then we define the linear quadratic (LQ) cost function $\mathcal{J} : \mathbb{R}^{n \times n} \times \mathcal{S}_+^n \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ by

$$\mathcal{J}(A, Q) := \mathbb{E}_{x_0, v} \left[\sum_{k=0}^{\infty} \alpha^k x_k^\top Q x_k \right].$$

Since we consider a *discounted* LQ cost, it is helpful to also introduce a respective notion of stability.

Definition II.3 ($\sqrt{\alpha}$ -stability): Let $\alpha \in (0, 1]$, then the matrix A is $\sqrt{\alpha}$ -stable when its spectrum is fully contained in the open disk with radius $\alpha^{-1/2}$, *i.e.*, $\sqrt{\alpha}A$ is exponentially stable.

One can observe that the classical exponential stability notion in system theory is a sufficient condition, and *not* necessary, for the $\sqrt{\alpha}$ -stability of Definition II.3.

The main objective of this study is to introduce an uncertainty set Δ that facilitates an exact and tractable robust LQR formulation which is meaningful to study. To this end, we first proceed with a brief discussion regarding a desirable property of such an uncertainty set Δ .

B. Convexity in Robust Linear Control

As the next example shows Assumption II.1 restricts possible Δ . There is no time-invariant K which can stabilize all stabilizable pairs (A, B) :

Example II.4 (Lack of universal stabilizing feedback law): Consider for some finite scalar c and $d \in (-1, 1)$ the matrices

$$A_1 = \begin{pmatrix} 1 & c \\ 0 & d \end{pmatrix}, \quad A_2 = \begin{pmatrix} -1 & c \\ 0 & d \end{pmatrix}, \quad B = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

The pairs (A_1, B) and (A_2, B) are stabilizable. However if we let the controller be of the form $K = (K_1 \ K_2)$ then (A_1, B) needs $K_1 \in (-2, 0)$ while (A_2, B) needs $K_1 \in (0, 2)$ to make the closed-loop matrix exponentially stable. Since $(-2, 0) \cap (0, 2) = \{\emptyset\}$ there is no K which can exponentially stabilize both systems.

Example II.4 can be interpreted in the spirit of *switching control*, *i.e.*, once (A_1, B) switches to system (A_2, B) your linear control law should switch as well. As indicated by the discrete-time version of Lemma 3.1 from [25], given a compact subset $\mathcal{K} \subset \mathbb{R}^{n \times n} \times \mathbb{R}^{n \times m}$ of the set of stabilizable pairs (A, B) one can indeed introduce a finite covering where all the elements of each segment can be stabilized via a common feedback gain, *e.g.*, (A_1, B) and (A_2, B) are never members of the same segment while for example $\|A_1\|_2 = \|A_2\|_2$. This simple observation indicates that the *existence* of a stabilizing solution to (1) is not immediately obvious, even for simple norm-balls.

One may wonder how these individual segments look like, and in particular with the desire of a tractable algorithm in mind, whether a set of stabilizable pairs with a common stabilizing feedback is necessarily convex in $\mathbb{R}^{n \times n} \times \mathbb{R}^{n \times m}$. The following example provides a negative answer to this question.

Example II.5 (Non-convex segment): Consider for $a = 2$ and $d = 0.5$ the matrices

$$A_1 = \begin{pmatrix} d & 0 & a \\ 0 & d & 0 \\ 0 & 0 & d \end{pmatrix}, \quad A_2 = A_1^\top, \quad B = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$$

Then (A_1, B) and (A_2, B) are both stabilizable, perhaps by $K = d^2 B^\top$, while for $A = 0.5A_1 + 0.5A_2$ the pair (A, B) is not stabilizable. Moreover, since one can find a path from (A_1, B) to (A_2, B) which can be stabilized by K , there does exist some non-convex segment containing both of the pairs.

These quick examples indicate that convex uncertainty sets for (A, B) in $\mathbb{R}^{n \times n} \times \mathbb{R}^{n \times m}$, which currently dominate the field, can be a restrictive point of view indeed and may be potentially conservative. Now, note that we do not claim that non-convexity is desirable, but merely observe that it should not be ruled out.

III. MAIN RESULTS

The main objective of this section is to provide a closed-form solution to the RLQR problem (3) and study its implications.

A. Introduction of a new uncertainty set

Definition III.1 (Uncertainty set): Given a tuple $(\hat{A}, D, \Sigma_0, \Sigma_v, \alpha)$ and some $\gamma \in \mathbb{R}_{\geq 0}$, let $W_{0,v} := \Sigma_0 + \alpha(1 - \alpha)^{-1}\Sigma_v$ and define a set of models around \hat{A} by the set:

$$\mathcal{A}_\gamma(\hat{A}) := \left\{ A \in \mathbb{R}^{n \times n} : \begin{array}{l} A = \hat{A} + D\Delta_A, \\ \Sigma_x = \alpha A \Sigma_x A^\top + W_{0,v}, \\ \Sigma_x \succ 0, \\ \langle \Delta_A^\top \Delta_A, \Sigma_x \rangle \leq \gamma \end{array} \right\}. \quad (2)$$

For notational convenience, we shall refer to the collection of Δ_A by $\Delta_\gamma(\hat{A})$. Using this notation, we therefore have the following simple relation between these sets: $\mathcal{A}_\gamma(\hat{A}) = \hat{A} + D\Delta_\gamma(\hat{A})^1$.

Remark III.2 (Absence of translation invariance): Let $B_r(x)$ be an Euclidean ball with radius r and center x . Then one can think of $\mathcal{A}_\gamma(\hat{A})$ as a ball with radius γ and center \hat{A} . However, in contrast to an Euclidean ball, our set is not translation invariant and depends on the center \hat{A} . Moreover, since $W_{0,v} \succ 0$, for Δ_A to be in $\Delta_\gamma(\hat{A})$ is the same as being part of the set $\{\Delta_A \in \mathbb{R}^{d \times n} : \|\Delta_A^\top\|_{F, \Sigma_x}^2 \leq \gamma\}$ for Σ_x as in (2). This further explains why γ is referred to as a “radius”.

Remark III.3 (Structural information): The matrix D in Definition III.1 may be used to incorporate a form of prior structural information into the uncertainty set. Without any prior structural information, one should choose $D = I_n$.

Before addressing (1) under (2), we provide, inspired by Lemma 2 from [22], some insights about the set \mathcal{A}_γ , which are especially interesting from an optimization point of view

¹With slight abuse of notation, by $+$ between two sets we mean the Minkowski sum: $A + B = \{a + b : a \in A, b \in B\}$.

Proposition III.4: The set $\mathcal{A}_\gamma(\hat{A})$ as defined in Definition III.1 has the following properties:

- (i) For $n \geq 3$ there are sets $\mathcal{A}_\gamma(\hat{A})$ which are non-convex.
- (ii) For $\gamma > 0$, the set $\mathcal{A}_\gamma(\hat{A})$ is semi-algebraic.

Further extending the tools from [22] to the game theoretic regime, allows for showing that the set is path-connected. The fact that our uncertainty set is semi-algebraic and does not rule out the lack of convexity is nice from a control theoretic point of view as well. See for example [26], an Euclidean ball of $(A, B) \in \mathbb{R}^{n \times n} \times \mathbb{R}^{n \times m}$ intersected with the set of controllable pairs (A, B) is semi-algebraic. In the following we illustrate the general non-convexity of the proposed uncertainty sets through an example.

Example III.5 (Non-convexity of Δ_γ): We consider the case where the uncertainty set $\mathcal{A}_\gamma \subset \mathbb{R}^{3 \times 3}$ from Definition III.1 is constructed using the parameters $\alpha = 0.95$, $D = Q_{cl} = \Sigma_0 = I_3$, and $\Sigma_v = 0.01I_3$. Here we will consider several “levels sets” of Δ_γ . Since the set Δ_γ is essentially a 9-dimensional object, for the sake of illustration we restrict our attention to a 2-dimensional subset. For this purpose, we consider the closed loop matrix A_{cl} , and especially all Δ_A , to be parametrized by

$$A_{cl} = \underbrace{\begin{pmatrix} 0.25 & 1.25 & -0.84 \\ 0 & 0.25 & 0 \\ 0.70 & 1.25 & 0.25 \end{pmatrix}}_{\hat{A} + BK} + \underbrace{\begin{pmatrix} 0 & 0 & \Delta_{A13} \\ 0 & 0 & 0 \\ \Delta_{A31} & 0 & 0 \end{pmatrix}}_{\Delta_A},$$

where $\Delta_{A13} = 4.98\theta_1 - 0.25\theta_2$, $\Delta_{A31} = 0.45\theta_2 - 1.08\theta_1$, and the parameters (θ_1, θ_2) belong to the interval $[-1, 1]^2$. This choice of (θ_1, θ_2) over $(\Delta_{A13}, \Delta_{A31})$ is purely driven by visualization purposes. Figure 1a depicts the 2-dimensional slice of Δ_γ by means of (θ_1, θ_2) for the levels: $\gamma \in \{2^{-4}, 2^1, 2^4, 2^{14}\}$. Interestingly enough, it is non-convex for large values of γ . Figure 1b also illustrates the LQ cost $\mathcal{J}(A_{cl}, Q_{cl})$ from Definition II.2.

At last, using the shorthand notation, the problem (1) over (2) is written as

$$\inf_{K \in \mathbb{R}^{n \times m}} \sup_{A_{cl} \in \mathcal{A}_\gamma(\hat{A} + BK)} \mathcal{J}(A_{cl}, Q + K^\top RK). \quad (3)$$


It is worth noting the dependence on K in the inner maximization step. A solution to (3) is given by $(K^*(\gamma), A_{cl}^*(\gamma))$.

B. Solving a Robust LQR Problem


In the first step, we tackle the worst-case LQ problem over \mathcal{A}_γ , being the inner maximization of the RLQR problem (3). This problem is defined as

$$\sup_{A_{cl} \in \mathcal{A}_\gamma(\hat{A}_{cl})} \mathcal{J}(A_{cl}, Q_{cl}), \quad (4)$$

for some given controller K $\sqrt{\alpha}$ -stabilizing $\hat{A}_{cl} := \hat{A} + BK$ and $Q_{cl} := Q + K^\top RK$ being the closed-loop cost matrix. Denote the solution to (4) by $A_{cl}^*(\gamma) := \hat{A}_{cl} + D\Delta_A^*(\gamma)$.



(a) Uncertainty set Δ_γ as in Definition III.1.



(b) LQ cost \mathcal{J} as in Definition II.2.

Fig. 1: Given the parameters from Example III.5 we show the uncertainty sets and LQ function parametrized by $(\theta_1, \theta_2) \in [-1, 1]^2$ for different levels $\gamma \in \{2^{-4}, 2^1, 2^4, 2^{14}\}$ from darker to lighter gray.

Proposition III.6 (Worst-case LQ cost): Consider problem (4) with nominal closed-loop model \hat{A}_{cl} , structural matrix D , some $\alpha \in (0, 1)$, initial data $\Sigma_0, \Sigma_v \in \mathcal{S}_{++}^n$, and closed-loop cost matrix $Q_{cl} \in \mathcal{S}_+^n$. Given some $\delta \in \mathbb{R}_{\geq 0}$, let us assume that $(\delta^{-1}I_d - \alpha D^\top S D) \succ 0$ is satisfied for the (minimal) positive semi-definite solution S to the algebraic equation

$$S = Q_{cl} + \alpha \hat{A}_{cl}^\top (S + \alpha S D (\delta^{-1} I_d - \alpha D^\top S D)^{-1} D^\top S) \hat{A}_{cl}.$$

Then define

$$\Delta_A^*(\delta) = (\delta^{-1} I_d - \alpha D^\top S D)^{-1} D^\top S \hat{A}_{cl}.$$

Further, define $\tilde{\Sigma}_x$ as the positive-definite solution to the Lyapunov equation

$$\tilde{\Sigma}_x = \alpha (\hat{A}_{cl} + D \Delta_A^*(\delta)) \tilde{\Sigma}_x (\hat{A}_{cl} + D \Delta_A^*(\delta))^\top + W_{0,v} \quad (5)$$

which in its turn defines the function

$$\tilde{h}(\delta) = \left\langle (\Delta_A^*(\delta))^\top \Delta_A^*(\delta), \tilde{\Sigma}_x \right\rangle. \quad (6)$$

Then, $\Delta_A^*(\gamma) = \Delta_A^*(\delta)$ and $\mathcal{J}^* = \langle \tilde{\Sigma}_x, Q_{cl} \rangle$ are the optimizer (worst-case uncertainty) and the optimal value of the problem (4) with the parameter $\gamma = \tilde{h}(\delta)$.

Now we are at the stage to address (3). This is not completely new, see for example [16], [27], where in the former², the pair (γ, δ) is interpreted via multiplier theory. We provide, in line with Definition III.1, a slightly different *system-* instead of *signal-* theoretic interpretation.

Theorem III.7 (Optimal Robust LQ regulator): Consider the RLQR problem (3) with the nominal $\sqrt{\alpha}$ -stabilizable model (\hat{A}, B) , the structural matrix D , $\alpha \in (0, 1)$, the cost matrices $Q \in \mathcal{S}_+^n, R \in \mathcal{S}_{++}^m$ and the covariance matrices $\Sigma_v, \Sigma_0 \in \mathcal{S}_{++}^n$. Given the parameter $\delta \in \mathbb{R}_{\geq 0}$, assume that the algebraic equation

$$P = Q + \alpha \hat{A}^\top P (I_n + \alpha (B R^{-1} B^\top - \delta D D^\top) P)^{-1} \hat{A} \quad (7)$$

in P admits a minimal³ positive semi-definite solution

²Specifically, see sec. 2.4 and ch.7-8 for a discussion.

³See chapter 3 from [28] for the definition and more information.

denoted $P(\delta)$ and define $\Lambda(\delta)$ correspondingly via $\Lambda := I_n + \alpha (B R^{-1} B^\top - \delta D D^\top) P$. Furthermore, define

$$\Delta_A^*(\delta) = \alpha \delta D^\top P(\delta) (\Lambda(\delta))^{-1} \hat{A} \quad (8)$$

and let $\hat{A}_{cl}^*(\gamma) := \hat{A} + D \Delta_A^*(\delta) + B K^*(\gamma)$. Next, consider the expressions for $\tilde{\Sigma}_x$ and $\tilde{h}(\delta)$ as in (5) and (6) respectively, which are now functions of K as well, to emphasize the difference, the tildes are dropped, *i.e.*, define:

$$\Sigma_x = \alpha \hat{A}_{cl}^*(\gamma) \Sigma_x (\hat{A}_{cl}^*(\gamma))^\top + W_{0,v} \quad (9)$$

$$h(\delta) = \left\langle (\Delta_A^*(\delta))^\top \Delta_A^*(\delta), \Sigma_x \right\rangle. \quad (10)$$

Then,

- (i) the controller $u_k = K^*(\gamma) x_k$ defined by

$$K^*(\gamma) = -\alpha R^{-1} B^\top P(\delta) (\Lambda(\delta))^{-1} \hat{A} \quad (11)$$

is (the minimizing part of) the solution to the RLQR problem for $\gamma = h(\delta)$.

- (ii) Furthermore, the maximizing solution is $\hat{A}_{cl}^*(\gamma)$, differently put, the worst-case⁴ system matrix is given by $A^*(\gamma) = \hat{A} + D \Delta_A^*(\delta)$.
- (iii) At last, the map $h(\delta)$ is analytic and non-decreasing over some interval $[0, \bar{\delta}) \subset \mathbb{R}_{\geq 0}$ for $\bar{\delta} < \infty$.

See section IV for a game theoretic interpretation of this “breakdown point” $\bar{\delta}$.

It is also important to remark that although problem (1) is well-defined for all $\gamma \in \mathbb{R}_{\geq 0}$, Theorem III.7 does not simply hold for any $\gamma \in \mathbb{R}_{\geq 0}$ but rather for some range $[0, \bar{\gamma}) \subseteq \mathbb{R}_{\geq 0}$ where $h(\bar{\delta}) = \bar{\gamma}$. See Section 5.2.2 from [24] for a discussion on the properties of this map h , we do not necessarily have $\lim_{\delta \uparrow \bar{\delta}} h(\delta) = \infty$. This explains the implicit formulation of the Theorem.

Additionally, despite this work being about A , we can make a remark regarding B . As shown in [13], [29], when $\det(\hat{A}) \neq 0$ then, theoretically, the extension to case with a partially known B is available by simply extending the state

⁴In the appendix of [24] we affirmatively answer the question if this worst-case model is actually a least-favourable model.

space and applying the aforementioned theory. See [24] for a further discussion and more ideas.

Finally, recall that Theorem III.7 presents us with an explicit expression for the worst-case model. From there we can infer further structural properties which is discussed at length in [24]. These observations are interesting since game theoretic formulations play a prominent role, either explicitly or implicitly, in many control-related fields. Also, it is expected that these results are quite general since they hinge on the symmetries in the cost. So, for the better or worse, even the most basic game theoretic robust control formulation displays a rich structure. It is the authors hope that this inspires further investigations in tractable robust control algorithms while alleviating predominant conservatism.

IV. GAME THEORETIC INTERPRETATION OF ROBUSTNESS

In this section we sketch the proof of the main results via a brief discussion on the connection of the original problem to dynamic game theory setting. We note that we are not the first to spot this link between game- and control theory, see for example [28] and references therein. Given the parameters $Q \in \mathcal{S}_+^n$, $R \in \mathcal{S}_{++}^m$, and $\delta \in [0, \bar{\delta}] \subseteq \mathbb{R}_{\geq 0}$, we define the function $g: \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^d \rightarrow \mathbb{R}$ by

$$g(x, u, w) = (x^\top Q x + u^\top R u - \delta^{-1} w^\top w)$$

and consider for some $\alpha \in (0, 1)$ the *stochastic (discounted) two-player zero-sum dynamic game* defined as:

$$\begin{aligned} \inf_{\{\mu_k\}_{k=0}^\infty} \sup_{\{\nu_k\}_{k=0}^\infty} & \mathbb{E}_{x_0, v} \left[\sum_{k=0}^{\infty} \alpha^k g(x_k, u_k, w_k) \right], \\ \text{s.t.} & \quad x_{k+1} = Ax_k + Bu_k + Dw_k + v_k, \\ & \quad v_k \stackrel{i.i.d.}{\sim} \mathcal{P}(0, \Sigma_v), \quad x_0 \sim \mathcal{P}(0, \Sigma_0), \\ & \quad u_k = \mu_k(x_k), \quad w_k = \nu_k(x_k). \end{aligned} \quad (12)$$

Here, the parameter δ penalizes the input of the ν -player, whose objective it is to destabilize the system, see [28] for conditions under which (12) can be solved. Note that this game is “*diagonal*” in the sense that there are no cross-terms in the cost. This form is chosen to keep the exposition simple, but one can consider more involved adversarial terms, e.g., $w_k^\top S w_k$ for some $S \succeq 0$. Nevertheless, this program heavily relies on the single parameter δ . The parameter δ is constrained to live in the interval $[0, \bar{\delta})$, where $\bar{\delta}$ is referred to as the *breakdown point*, beyond this value, the ν -player has to pay so little that the it can steer the cost to infinity⁵.

To see a relationship between dynamic game theory and parametric uncertainty sets, suppose (12) admits a solution, then consider the following. The policy of the ν -player aims at maximizing the cost. But since the μ -player can handle this *worst-case* policy, it must also be able to handle policies of a *less* powerful adversary. This effectively gives rise to a whole *family* of state feedback policies the μ -player can handle.

⁵See ch.8 [16] for more on the relation between this breakdown point and \mathcal{H}_∞ control.

Using the Lagrangian formulation from constrained optimization, one can take the adversarial part out of the cost and put it into the constraints, *i.e.*, let g be redefined as $g(x, u) = x^\top Q x + u^\top R u$ and add a constraint of the form $\mathbb{E}_{x_0, v} [\alpha^k w_k^\top w_k] \leq \gamma$, for some $\gamma \in \mathbb{R}_{\geq 0}$. Then, Theorem III.7 establishes a link between γ and δ of the form $h(\delta) = \gamma$ such that we can relate their solutions as well.

V. NUMERICAL EXAMPLE

The goal of this section is to compare the actions of a nominal control law to our robust framework. Consider the controllable pair (\hat{A}, B) and the structural matrix D defined as

$$\hat{A} = \begin{pmatrix} 1.2 & 0.5 \\ 0 & 1.2 \end{pmatrix}, \quad B = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad D = \begin{pmatrix} 1 \\ 1 \end{pmatrix}.$$

Also define the covariance matrices $\Sigma_v = 0.1I_2$, $\Sigma_0 = I_2$, the cost matrices $Q = 0.1I_2$, $R = 10$, and the discount factor $\alpha = 0.95$.

Then, set K to the nominal discounted LQ regulator⁶, *i.e.*, $K = K^*(0)$. Now, Figure 2a depicts the level sets of $\Delta_\gamma(\hat{A} + BK^*(0))$ as defined by Definition III.1 for different levels⁷ $\gamma \in \Gamma := \{0.005, 0.03, 0.09, 0.4, 1\}$. We further solve the worst-case model uncertainty problem (4) via Proposition III.6. Let us recall that the mapping h defined in (6) provides the relation $\gamma = \tilde{h}(\delta)$ between the different values of γ . In this example, the corresponding δ are $10^{-3} \cdot \{2, 3.9, 5.5, 7.3, 7.7\}$. The locations of these worst-case models are marked by a star symbol in Figure 2a.

Looking at the cost in Figure 2a anyone could guess where these worst-case uncertainties reside. However, we have only considered this low-dimensional example for visualization purposes. Computationally speaking, nothing prohibits us from doing high dimensional examples (*e.g.*, $n = 1000$), and then Proposition III.6 might help in indicating where your system is sensitive with respect to the cost.

Finally, it is interesting to highlight what a robust controller $K^*(\gamma)$ would do, for say, $\gamma_5 = 1$. See Figure 2b for the corresponding sets $\Delta_{\gamma=1}$ under both types of controller. When compared with Figure 2a, we see that the robust controller anticipates on where the troubles might occur.


Remark V.1 (From radius to feedback): In [24] we provide tools to do the aforementioned computations efficiently, which hinge on Theorem III.7.(iii). For example, let us be given a desired “*radius*” γ and assume it is feasible in the sense of Theorem III.7. Moreover, let the (local) Lipschitz constant of the map h be some $L \in \mathbb{R}_{>0}$ on $[0, \delta)$ and select $\beta \in \mathbb{R}_{>0} : \beta \leq L^{-1}$. Then, the algorithm

$$\delta_{k+1} = \delta_k + \beta(\gamma - h(\delta_k)), \quad \delta_0 = 0,$$

converges to $\delta : h(\delta) = \gamma$ at a linear rate proportional to the estimation error of L . Now, to obtain the feedback $K^*(\gamma)$, given the correct δ , one can solve the Generalized Algebraic Riccati Equation (7) iteratively as proposed in [30].

⁶ $K = -\alpha(R + \alpha B^\top P B)^{-1} B^\top P \hat{A}$ for $P = Q + \alpha \hat{A}^\top P \hat{A} - \alpha^2 \hat{A}^\top P B (R + \alpha B^\top P B)^{-1} B^\top P \hat{A}$.

⁷ $\gamma = 0$ would yield 0 since $\Sigma_x \succ 0$.



(a) For $\gamma \in \Gamma$, the sets $\Delta_\gamma(\hat{A} + BK^*(0))$ with the corresponding worst-case path, including a projection on the cost surface. (b) Comparison of the uncertainty hedged against for the nominal-, $K^*(0)$, and robust controller $K^*(1)$, i.e., $\Delta_{\gamma=1}$ under both K .

Fig. 2: Given the parameters from section V we show the worst-case uncertainties via Proposition III.6 plus how our robust controller anticipates on where the cost increases the sharpest.

BIBLIOGRAPHY

- [1] K. Astrom and B. Wittenmark, "On self tuning regulators," *Automatica*, vol. 9, pp. 185–199, 1973.
- [2] C.-N. Fiechter, "Pac adaptive control of linear systems," in *Proceedings of the Tenth Annual Conference on Computational Learning Theory*, ser. COLT '97, ACM, 1997, pp. 72–80.
- [3] M. C. Campi and P. R. Kumar, "Adaptive linear quadratic gaussian control: The cost-biased approach revisited," *SIAM J. CONTROL OPTIM.*, vol. 36, no. 6, pp. 1890–1907, 1998.
- [4] J. C. Doyle, "Guaranteed margins for lqg regulators," *IEEE Transactions on Automatic Control*, vol. AC-23, no. 4, pp. 756–757, 1978.
- [5] U. Shaked, "Guaranteed stability margins for the discrete-time linear quadratic optimal regulator," *IEEE Transactions on Automatic Control*, vol. 31, no. 2, pp. 162–165, 1986.
- [6] P. Lancaster and L. Rodman, *Algebraic Riccati Equations*, ser. Oxford Science Publications. Oxford University Press, 1995.
- [7] J.-G. Sun, "Perturbation theory for algebraic riccati equations," *SIAM J. MATRIX ANAL. APPL.*, vol. 19, no. 1, pp. 39–65, 1998.
- [8] S. Poljak and J. Rohn, "Checking robust nonsingularity is np-hard," *Mathematics of Control, Signals and Systems*, vol. 6, no. 1, pp. 1–9, 1993.
- [9] R. D. Braatz, P. M. Young, J. C. Doyle, and M. Morari, "Computational complexity of μ calculation," in *1993 American Control Conference*, 1993, pp. 1682–1683.
- [10] M. de Oliveira, J. Bernussou, and J. Geromel, "A new discrete-time robust stability condition," *Systems & Control Letters*, vol. 37, no. 4, pp. 261–265, 1999.
- [11] P. Whittle, *Risk-sensitive Optimal Control*. Wiley, 1990.
- [12] P. P. Khargonekar, I. R. Petersen, and K. Zhou, "Robust stabilization of uncertain linear systems: Quadratic stabilizability and h_∞ control theory," *IEEE Transactions on Automatic Control*, vol. 35, no. 3, pp. 356–361, 1990.
- [13] G. Garcia, J. Bernussou, and D. Arzelier, "Robust stabilization of discrete-time linear systems with norm-bounded time-varying uncertainty," *Systems & Control Letters*, vol. 22, no. 5, pp. 327–339, 1994.
- [14] M. Abeille and A. Lazaric, "Improved regret bounds for thompson sampling in linear quadratic control problems," in *Proceedings of the 35th International Conference on Machine Learning*, vol. 80, PMLR, 2018, pp. 1–9.
- [15] M. Simchowitz, H. Mania, S. Tu, M. I. Jordan, and B. Recht, "Learning without mixing: Towards a sharp analysis of linear system identification," in *Proceedings of the 31st Conference On Learning Theory*, vol. 75, PMLR, 2018, pp. 439–473.
- [16] L. P. Hansen and T. J. Sargent, *Robustness*. Princeton University Press, 2007.
- [17] I. R. Petersen, V. Ugrinovskii, and A. Savkin, *Robust Control Design Using H_∞ Methods*. Springer, 2000.
- [18] Y. Abbasi-Yadkori and C. Szepesvri, "Regret bounds for the adaptive control of linear quadratic systems," *JMLR: Workshop and Conference Proceedings 19, 24th Annual Conference on Learning Theory*, 2011.
- [19] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu, "On the sample complexity of the linear quadratic regulator," *Foundations of Computational Mathematics*, 2019.
- [20] S. Tu, "Sample complexity bounds for the linear quadratic regulator," PhD thesis, EECS Department, University of California, Berkeley, 2019.
- [21] A. Cohen, T. Koren, and Y. Mansour, "Learning linear-quadratic regulators efficiently with only \sqrt{t} regret," in *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, 2019, pp. 1300–1309.
- [22] M. Fazel, R. Ge, S. Kakade, and M. Mesbahi, "Global convergence of policy gradient methods for the linear quadratic regulator," in *Proceedings of the 35th International Conference on Machine Learning*, vol. 80, PMLR, 2018, pp. 1467–1476.
- [23] B. Gravell, P. Mohajerin Esfahani, and T. Summers, "Learning robust control for LQR systems with multiplicative noise via policy gradient," *arXiv e-prints*, 2019.
- [24] W. Jongeneel, T. Summers, and P. Mohajerin Esfahani, "Robust linear quadratic regulator: Exact tractable reformulation (extended version)," 2019. [Online]. Available: {http://www.dsc.tudelft.nl/~mohajerin/Publications/conference/2019/RLQR_extended.pdf}.
- [25] M. Fu and B. Barmish, "Adaptive stabilization of linear systems via switching control," *IEEE Transactions on Automatic Control*, vol. 31, no. 12, pp. 1097–1103, 1986.
- [26] R. Hermann and C. F. Martin, "Applications of algebraic geometry to systems theory - part i," *IEEE Transactions on Automatic Control*, vol. 22, no. 1, pp. 19–25, 1977.
- [27] G. Didinsky and T. Basar, "Design of minimax controllers for linear systems with non-zero initial states under specified information structures," *International Journal of Robust and Nonlinear Control*, vol. 2, no. 1, pp. 1–30, 1992.
- [28] T. Basar and P. Bernhard, *H_∞ -Optimal Control and Related Minimax Design Problems A Dynamic Game Approach*. Birkhauser, 1995.
- [29] B. Barmish, "Stabilization of uncertain systems via linear control," *IEEE Transactions on Automatic Control*, vol. 28, no. 8, pp. 848–850, 1983.
- [30] A. A. Stoorvogel and A. J.T. M. Weeren, "The discrete-time riccati equation related to the h_∞ control problem," *IEEE Transactions on Automatic Control*, vol. 39, no. 3, pp. 686–691, 1994.