
Efficient Optimistic Exploration in Linear-Quadratic Regulators via Lagrangian Relaxation

Marc Abeille¹ Alessandro Lazaric²

Abstract

We study the exploration-exploitation dilemma in the linear quadratic regulator (LQR) setting. Inspired by the extended value iteration algorithm used in optimistic algorithms for finite MDPs, we propose to relax the optimistic optimization of OFU-LQ and cast it into a constrained *extended* LQR problem, where an additional control variable implicitly selects the system dynamics within a confidence interval. We then move to the corresponding Lagrangian formulation for which we prove strong duality. As a result, we show that an ϵ -optimistic controller can be computed efficiently by solving at most $O(\log(1/\epsilon))$ Riccati equations. Finally, we prove that relaxing the original OFU problem does not impact the learning performance, thus recovering the $\tilde{O}(\sqrt{T})$ regret of OFU-LQ. To the best of our knowledge, this is the first computationally efficient confidence-based algorithm for LQR with worst-case optimal regret guarantees.

1. Introduction

Exploration-exploitation in Markov decision processes (MDPs) with continuous state-action spaces is a challenging problem: estimating the parameters of a generic MDP may require many samples, and computing the corresponding optimal policy may be computationally prohibitive. The linear quadratic regulator (LQR) model formalizes continuous state-action problems, where the dynamics is linear and the cost is quadratic in state and action variables. Thanks to its specific structure, it is possible to efficiently estimate the parameters of the LQR by least-squares regression and the optimal policy can be computed by solving a Riccati equation. As a result, several exploration strategies have been adapted to the LQR to obtain effective learning algorithms.

¹Criteo AI Lab ²Facebook AI Research. Correspondence to: Marc Abeille <m.abeille@criteo.com>.

Confidence-based exploration. Bittanti et al. (2006) introduced an adaptive control system based on the “bet on best” principle and proved asymptotic performance guarantees showing that their method would eventually converge to the optimal control. Abbasi-Yadkori & Szepesvári (2011) later proved a finite-time $\tilde{O}(\sqrt{T})$ regret bound for OFU-LQ, later generalized to less restrictive stabilization and noise assumptions by Faradonbeh et al. (2017). Unfortunately, neither exploration strategy comes with a computationally efficient algorithm to solve the optimistic LQR, and thus they cannot be directly implemented. On the TS side, Ouyang et al. (2017) proved a $\tilde{O}(\sqrt{T})$ regret for the Bayesian regret, while Abeille & Lazaric (2018) showed that a similar bound holds in the frequentist case but restricted to 1-dimensional problems. While TS-based approaches require solving a single (random) LQR, the theoretical analysis of Abeille & Lazaric (2018) suggests that a new LQR instance should be solved at each time step, thus leading to a computational complexity growing linearly with the total number of steps. On the other hand, OFU-based methods allow for “lazy” updates, which require solving an optimistic LQR only a *logarithmic* number of times w.r.t. the total number of steps. A similar lazy-update scheme is used by Dean et al. (2018), who leveraged robust control theory to devise the first learning algorithm with polynomial complexity and sublinear regret. Nonetheless, the resulting adaptive algorithm suffers from a $\tilde{O}(T^{2/3})$ regret, which is significantly worse than the $\tilde{O}(\sqrt{T})$ achieved by OFU-LQ.

To the best of our knowledge, the only efficient algorithm for confidence-based exploration with $\tilde{O}(\sqrt{T})$ regret has been recently proposed by Cohen et al. (2019). Their method, called OSLO, leverages an SDP formulation of the LQ problem, where an optimistic version of the constraints is used. As such, it translates the original non-convex OFU-LQ optimization problem into a *convex* SDP. While solving an SDP is known to have *polynomial* complexity, no explicit analysis is provided and it is said that the runtime may scale polynomially with LQ-specific parameters and the time horizon T (Cor. 5), suggesting that OSLO may become impractical for moderately large T . Furthermore, OSLO requires an initial system identification phase of length $\tilde{O}(\sqrt{T})$ to properly initialize the method. This strong requirement effectively reduces OSLO to an explore-then-commit strategy, whose

regret is dominated by the length of the initial phase.

Perturbed certainty equivalence exploration. A recent stream of research (Faradonbeh et al., 2018b; Mania et al., 2019; Simchowitz & Foster, 2020) studies variants of the perturbed certainty equivalence (CE) controller (i.e., the optimal controller for the estimated LQR) and showed that this simple exploration strategy is sufficient to reach worst-case optimal regret $\tilde{O}(\sqrt{T})$. Since the CE controller is not recomputed at each step (i.e., lazy updates) and the perturbation is obtained by sampling from a Gaussian distribution, the resulting methods are computationally efficient. Nonetheless, these methods rely on an isotropic perturbation (i.e., all control dimensions are equally perturbed) and they require the variance to be *large enough* so as to *eventually* reduce the uncertainty on the system estimate along the dimensions that are not naturally “excited” by the CE controller and the environment noise. Being agnostic to the uncertainty of the model estimate and its impact on the average cost, may lead this type of approaches to have longer (unnecessary) exploration and larger regret. On the other hand, confidence-based methods relies on exploration controllers that are explicitly designed to excite more the dimensions with higher uncertainty and impact on the performance. As a result, they are able to perform more effective exploration. We further discuss this difference in Sect. 6.

In this paper, we introduce a novel instance of OFU, for which we derive a computationally efficient algorithm to solve the optimistic LQR with explicit computational complexity and $\tilde{O}(\sqrt{T})$ regret guarantees. Our approach is inspired by the extended value iteration (EVI) used to solve a similar optimistic optimization problem in finite state-action MDPs (e.g. Jaksch et al., 2010). Relying on an initial estimate of the system obtained after a *finite* number of system identification steps, we first relax the confidence ellipsoid constraints and we cast the OFU optimization problem into a constrained LQR with extended control. We show that the relaxation of the confidence ellipsoid constraint does not impact the regret and we recover a $\tilde{O}(\sqrt{T})$ bound. We then turn the constrained LQR into a regularized optimization problem via Lagrangian relaxation. We prove strong duality and show that we can compute an ϵ -optimistic and ϵ -feasible solution for the constrained LQR by solving only $O(\log(1/\epsilon))$ algebraic Riccati equations. As a result, we obtain the *first efficient worst-case optimal confidence-based algorithm for LQR*. In deriving these results, we introduce a novel derivation of explicit conditions on the accuracy of the system identification phase leveraging tools from Lyapunov stability theory that may be of independent interest.

2. Preliminaries

We consider the discrete-time linear quadratic regulator (LQR) problem. At any time t , given state $x_t \in \mathbb{R}^n$ and

control $u_t \in \mathbb{R}^d$, the next state and cost are obtained as

$$\begin{aligned} x_{t+1} &= A_* x_t + B_* u_t + \epsilon_{t+1}; \\ c(x_t, u_t) &= x_t^\top Q x_t + u_t^\top R u_t, \end{aligned} \quad (1)$$

where A_* , B_* , Q , R are matrices of appropriate dimension and $\{\epsilon_{t+1}\}_t$ is the process noise. Let $\mathcal{F}_t = \sigma(x_0, u_0, \dots, x_t, u_t)$ be the filtration up to time t , we rely on the following assumption on the noise.¹

Assumption 1. *The noise $\{\epsilon_t\}_t$ is a martingale difference sequence w.r.t. the filtration \mathcal{F}_t and it is componentwise conditionally sub-Gaussian, i.e., there exists $\sigma > 0$ such that $\mathbb{E}(\exp(\gamma \epsilon_{t+1,i}) | \mathcal{F}_t) \leq \exp(\gamma^2 \sigma^2 / 2)$ for all $\gamma \in \mathbb{R}$. Furthermore, we assume that the covariance of ϵ_t is the identity matrix.*

The dynamics parameters are summarized in $\theta_*^\top = (A_*, B_*)$ and the cost function can be written as $c(x_t, u_t) = z_t^\top C z_t$ with $z_t = (x_t, u_t)^\top$ and the cost matrix

$$C = \begin{pmatrix} Q & 0 \\ 0 & R \end{pmatrix}. \quad (2)$$

The solution to an LQ is a stationary deterministic policy $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^d$ mapping states to controls minimizing the infinite-horizon average expected cost

$$J_\pi(\theta_*) = \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=0}^T c(x_t, u_t) \right], \quad (3)$$

with $x_0 = 0$ and $u_t = \pi(x_t)$. We assume that the LQR problem is “well-posed”.

Assumption 2. *The cost matrices Q and R are symmetric p.d. and known, and (A_*, B_*) is stabilizable, i.e., there exists a controller K , such that $\rho(A_* + B_* K) < 1$.²*

In this case, Thm.16.6.4 in (Lancaster & Rodman, 1995) guarantees the existence and uniqueness of an optimal policy $\pi_* = \arg \min_\pi J_\pi(\theta_*)$, which is linear in the state, i.e., $\pi_*(x) = K(\theta_*)x$, where,

$$\begin{aligned} K(\theta_*) &= -(R + B_*^\top P(\theta_*) B_*)^{-1} B_*^\top P(\theta_*) A_*, \\ P(\theta_*) &= Q + A_*^\top P(\theta_*) A_* + A_*^\top P(\theta_*) B_* K(\theta_*). \end{aligned} \quad (4)$$

For convenience, we will denote $P_* = P(\theta_*)$. The optimal average cost is $J_* = J_{\pi_*}(\theta_*) = \text{Tr}(P_*)$. Further, let $L(\theta_*)^\top = (I \ K(\theta_*)^\top)$, then the closed-loop matrix $A^c(\theta_*) = A_* + B_* K(\theta_*) = \theta_*^\top L(\theta_*)$ is asymptotically stable.

¹As shown by Faradonbeh et al. (2017), this can be relaxed to Weibull distributions with known covariance.

² $\rho(A)$ is the spectral radius of the matrix A , i.e., the largest absolute value of the eigenvalues of A .

While Asm. 2 guarantees the existence of an optimal linear controller, its optimal cost J_* may still grow unbounded when θ_* is nearly unstable. A popular solution is to introduce a “strong” stability assumption (i.e., $\rho(A^c(\theta_*)) \leq \bar{\rho} < 1$). Nonetheless, this imposes stability uniformly over all state dimensions, whereas, depending on the cost matrices Q and R , some dimensions may be less sensitive than others in terms of their impact on the cost. Here we prefer imposing an assumption directly on the optimal cost.³

Assumption 3. *There exists $D > 0$ such that $J_* = \text{Tr}(P_*) \leq D$ and D is known.*

Finally, we introduce $\kappa = D/\lambda_{\min}(C)$, a quantity that will characterize the complexity of many aspects of the learning problem in the following. Intuitively, κ measures the cost of controlling w.r.t. the minimum cost incurred if the uncontrolled system was perfectly stable.

The learning problem. We assume that Q and R are known, while θ_* needs to be estimated from data. We consider the online learning problem where at each step t the learner observes the current state x_t , it executes a control u_t and it incurs the associated cost $c(x_t, u_t)$; the system then transitions to the next state x_{t+1} according to Eq. 1. The learning performance is measured by the cumulative regret over T steps defined as $\mathcal{R}_T(\theta_*) = \sum_{t=0}^T (c_t - J_*(\theta_*))$. Exploiting the linearity of the dynamics, the unknown parameter θ_* can be directly estimated from data by regularized least-squares (RLS). For any sequence of controls (u_0, \dots, u_t) and the induced states $(x_0, x_1, \dots, x_{t+1})$, let $z_t = (x_t, u_t)^\top$, the RLS estimator with a regularization bias θ_0 and regularization parameter $\lambda \in \mathbb{R}_+^*$ defined as⁴

$$\begin{aligned} \hat{\theta}_t &= \arg \min_{\theta \in \mathbb{R}^{(n+d) \times n}} \sum_{s=0}^{t-1} \|x_{s+1} - \theta^\top z_s\|^2 + \lambda \|\theta - \theta_0\|_F^2 \\ &= V_t^{-1} \left(\lambda \theta_0 + \sum_{s=0}^{t-1} z_s x_{s+1}^\top \right), \end{aligned} \quad (5)$$

where $V_t = \lambda I + \sum_{s=0}^{t-1} z_s z_s^\top$ is the design matrix. The RLS estimator concentrates as follows (see App. C.3).

Proposition 1 (Thm. 1 in Abbasi-Yadkori & Szepesvári 2011). *For any $\delta \in (0, 1)$ and any \mathcal{F}_t -adapted sequence (z_0, \dots, z_t) , the RLS estimator $\hat{\theta}_t$ is such that*

$$\|\theta_* - \hat{\theta}_t\|_{V_t} \leq \beta_t(\delta) \quad \text{where} \quad (6)$$

$$\beta_t(\delta) = \sigma \sqrt{2n \log \left(\frac{\det(V_t)^{1/2} n}{\det(\lambda I)^{1/2} \delta} \right)} + \lambda^{1/2} \|\theta_0 - \theta_*\|_F,$$

w.p. $1 - \delta$ (w.r.t. the noise $\{\epsilon_{t+1}\}_t$ and any randomization in the choice of the control).

³An alternative assumption may bound the operator norm of P_* , (see e.g., Simchowitz & Foster (2020)).

⁴For $\theta_0 = 0$, this reduces to the standard estimator. The need for a “centered” regularization term is explained in the next section.

Finally, we recall a standard result of RLS.

Proposition 2 (Lem. 10 in Abbasi-Yadkori & Szepesvári 2011). *Let $\lambda \geq 1$, for any arbitrary \mathcal{F}_t -adapted sequence (z_0, z_1, \dots, z_t) , let V_{t+1} be the corresponding design matrix, then*

$$\sum_{s=0}^t \min(\|z_s\|_{V_s^{-1}}^2, 1) \leq 2 \log \frac{\det(V_{t+1})}{\det(\lambda I)}.$$

Moreover when $\|z_t\| \leq Z$ for all $t \geq 0$, then

$$\sum_{s=0}^t \|z_s\|_{V_s^{-1}}^2 \leq \left(1 + \frac{Z^2}{\lambda}\right)(n+d) \log \left(1 + \frac{(t+1)Z^2}{\lambda(n+d)}\right). \quad (7)$$

3. A Sequentially Stable Variant of OFU-LQ

In this section we introduce a variant of the original OFU-LQ of Abbasi-Yadkori & Szepesvári (2011) that we refer to as OFU-LQ++. Similar to (Faradonbeh et al., 2017), we use an initial system identification phase to initialize the system and we provide *explicit* conditions on the accuracy required to guarantee sequential stability thereafter. This result is obtained leveraging tools from Lyapunov stability theory which may be of independent interest.

Faradonbeh et al. (2018a) showed that it is possible to construct a set $\Theta_0 = \{\theta : \|\theta - \theta_0\| \leq \epsilon_0\}$ containing the true parameters θ_* with high probability, through a system identification phase where a randomized sequence of linear controllers is used to accurately estimate the dynamics. In particular, they proved that a set Θ_0 with accuracy ϵ_0 can be obtained by running the system identification phase for as long as $T_0 = \Omega(\epsilon_0^{-2})$ steps.⁵

After the initial phase, OFU-LQ++ uses the estimate θ_0 to regularize the RLS as in (5) and it proceeds through episodes. At the beginning of episode k it computes a parameter

$$\theta_k = \arg \min_{\theta \in \mathcal{C}_k} J(\theta), \quad (8)$$

where t_k is the step at which episode k begins and the constrained set \mathcal{C}_k is defined as

$$\mathcal{C}_k = \mathcal{C}(\beta_{t_k}, V_{t_k}) := \{\theta : \|\theta - \hat{\theta}_{t_k}\|_{V_{t_k}} \leq \beta_{t_k}\}, \quad (9)$$

where $\beta_t = \beta_t(\delta/4)$ is defined in (6). Then the corresponding optimal control $K(\theta_k)$ is computed (4) and the associated policy is executed until $\det(V_t) \geq 2 \det(V_{t_k})$.

Lemma 1. *Let $\Theta_0 = \{\theta : \|\theta - \theta_0\| \leq \epsilon_0\}$ be the output of the initial system identification phase of Faradonbeh*

⁵An alternative scheme for system identification requires access to a stable controller K_0 and to perturb the corresponding controls to returned a set Θ_0 of desired accuracy ϵ_0 (see e.g., Simchowitz & Foster 2020).

et al. (2018a). For all $t \geq 0$, consider the confidence ellipsoid $\mathcal{C}_t := \{\theta : \|\theta - \hat{\theta}_t\|_{V_t} \leq \beta_t\}$, where $\hat{\theta}_t$ and V_t are defined in (5) with regularization bias θ_0 (the center of Θ_0), regularization parameter

$$\lambda = \frac{2n\sigma^2}{\epsilon_0^2} \left(\log(4n/\delta) + (n+d) \log(1 + \kappa X^2 T) \right), \quad (10)$$

and β_t is defined in (6) where $\|\theta - \theta_*\|$ is replaced by its upper-bound ϵ_0 . Let $\{K(\theta_t)\}_{t \geq 1}$ be the sequence of optimistic controllers generated by OFU-LQ++ and let $\{x_t\}_{t \geq 0}$ be the induced state process (Eq. 1). If $\epsilon_0 \leq O(1/\kappa^2)$, then with probability at least $1 - \delta/2$, for all $t \leq T$,

$$\begin{cases} \theta_* \in \mathcal{C}_t \\ \|x_t\| \leq X := 20\sigma \sqrt{\kappa \|P_*\|_2 \log(4T/\delta) / \lambda_{\min}(C)}. \end{cases} \quad (11)$$

OFU-LQ++ has some crucial differences w.r.t. the original algorithm. OFU-LQ receives as input a Θ_0 such that for any $\theta \in \Theta_0$ the condition $\|\theta^T L(\theta)\| < 1$ holds. While this condition ensures that all the LQ systems in Θ_0 are indeed stable, it does not immediately imply that the optimal controllers $K(\theta)$ stabilize the *true* system θ_* . Nonetheless, Abbasi-Yadkori & Szepesvári (2011) proved that the sequence of controllers generated by OFU-LQ naturally defines a state process $\{x_t\}_t$ which remains bounded at any step t with high probability. Unfortunately, their analysis suffers from several drawbacks: **1)** the state bound scales exponentially with the dimensionality, **2)** as θ_* is required to belong to Θ_0 , it should satisfy itself the condition $\|\theta_*^T L(\theta_*)\| < 1$, which significantly restricts the type of LQR systems that can be solved by OFU-LQ, **3)** the existence of Θ_0 is stated by *assumption* and no concrete algorithm to construct it is provided.

Furthermore, OFU-LQ requires solving (8) under the constraint that θ belongs to the intersection $\mathcal{C}_k \cap \Theta_0$, while OFU-LQ++ only uses the confidence set \mathcal{C}_k to guarantee that the controllers $K(\hat{\theta}_k)$ generated through the episodes induces a sequentially stable state process. Although the resulting optimization problem is still non-convex and difficult to solve directly, removing the constraint of Θ_0 enables the relaxation that we introduce in the next section. Finally, we notice that our novel analysis of the sequential stability of OFU-LQ++ leads to a tighter bound on the state, more explicit conditions on ϵ_0 , and lighter assumptions than Faradonbeh et al. (2018a).

As a result, we can refine the analysis of OFU-LQ and obtain a much sharper regret bound for OFU-LQ++.

Lemma 2. For any LQR (A_*, B_*, Q, R) satisfying Asm. 1, 2, and 3, after T steps OFU-LQ++, if properly initialized and tuned as in Lem. 1, suffers a regret

$$\mathcal{R}(T) = \tilde{O} \left(\left(\kappa \|P_*\|_2^2 + \sqrt{\kappa} \|P_*\|_2 (n+d) \sqrt{n} \right) \sqrt{T} \right). \quad (12)$$

4. An Extended Formulation of OFU-LQ++

The optimization in (8) is non-convex and it cannot be solved directly. In this section we introduce a relaxed constrained formulation of (8) and show that its solution is an optimistic controller with similar regret as OFU-LQ++ at the cost of requiring a slightly more accurate initial exploration phase (i.e., smaller ϵ_0).

4.1. The Extended Optimistic LQR with Relaxed Constraints

Our approach is directly inspired by the extended value iteration (EVI) used to solve a similar optimistic optimization problem in finite state-action MDPs (e.g. Jaksch et al., 2010). In EVI, the choice of dynamics θ from \mathcal{C} is added as an additional control variable, thus obtaining an extended policy $\tilde{\pi}$. Exploiting the specific structure of finite MDPs, it is shown that optimizing over policies $\tilde{\pi}$ through value iteration is equivalent to solving a (finite) MDP with the same state space and an extended (compact) control space and the resulting optimal policy, which prescribes both actions and a choice of the model θ , is indeed optimistic w.r.t. the original MDP. Leveraging a similar idea, we “extend” the LQR with estimated parameter $\hat{\theta}_t$ by introducing an additional control variable w corresponding to a specific choice of $\theta \in \mathcal{C}_k$. In the following we remove the dependency of $\hat{\theta}$, β , V , and \mathcal{C} on the learning step t_k and episode k , while we use a generic time s to formulate the extended LQR.

Let $\theta \in \mathcal{C}$ such that $\theta = \hat{\theta} + \delta_\theta = (A, B) = (\hat{A} + \delta_A, \hat{B} + \delta_B)$, then the dynamics of the corresponding LQR is

$$\begin{aligned} x_{s+1} &= Ax_s + Bu_s + \epsilon_{s+1} \\ &= \hat{A}x_s + \hat{B}u_s + \delta_A x_s + \delta_B u_s + \epsilon_{s+1}, \\ &= \hat{A}x_s + \hat{B}u_s + \delta_\theta z_s + \epsilon_{s+1}, \end{aligned} \quad (13)$$

where we isolate the “perturbations” δ_A and δ_B applied to the current estimates. We replace the perturbation associated to θ with a novel control variable w_s , the *perturbation control variable*, which effectively plays the role of “choosing” the parameters of the perturbed LQR, thus obtaining

$$\begin{aligned} x_{s+1} &= \hat{A}x_s + \hat{B}u_s + w_s + \epsilon_{s+1}, \\ &= \hat{A}x_s + \tilde{B}\tilde{u}_s + \epsilon_{s+1}, \end{aligned} \quad (14)$$

where we conveniently introduced $\tilde{B} = [\hat{B}, I]$ and $\tilde{u}_s = [u_s, w_s]$.⁶ This *extended* system has the same state variables as the original LQ, while the number of control variables moves from d to $n + d$. Since perturbations δ_θ are such that $\theta = \hat{\theta} + \delta_\theta \in \mathcal{C}$, we introduce a constraint on the perturbation control such that $\|w_s\| = \|\delta_\theta^T z_s\| \leq \beta \|z_s\|_{V^{-1}}$

⁶In the following we use tilde-notation such as $\tilde{\pi}$ and $w_t B$ to denote quantities in the extended LQR.

(see Prop. 1). We refer to the resulting system as the *extended LQR with hard constraints*. Unfortunately, this constrained system is no longer a “standard” LQR structure, as the constraint should be verified *at each step*. To overcome this difficulty, we relax the previous constraint and define

$$g_{\tilde{\pi}}(\hat{\theta}, \beta, V) = \lim_{S \rightarrow \infty} \frac{1}{S} \mathbb{E} \left(\sum_{s=0}^S \|w_s\|^2 - \beta^2 \|z_s\|_{V^{-1}}^2 \right), \quad (15)$$

where $\tilde{\pi} = (\pi^u, \pi^w)$ is an *extended* policy defining both standard and perturbation controls, so that $u_s = \pi^u(x_s)$ and $w_s = \pi^w(x_s)$, the expectation is w.r.t. the noise ϵ_{s+1} , and the dynamics of x_s follows (14). As a result, we translate the original constraint $\theta \in \mathcal{C}$, which imposed a per-step condition on w_s to $g_{\tilde{\pi}}(\hat{\theta}, \beta, V) \leq 0$, which considers the asymptotic average behavior of w_s . We are now ready to define the *extended LQR with relaxed constraints* as

$$\begin{aligned} \min_{\tilde{\pi}} \mathcal{J}_{\tilde{\pi}}(\hat{\theta}, \beta, V) &:= \limsup_{S \rightarrow \infty} \frac{1}{S} \mathbb{E} \left[\sum_{s=0}^S c(x_s, \pi^u(x_s)) \right] \\ \text{subject to} \quad &x_{s+1} = \hat{A}x_s + \tilde{B}\tilde{u}_s + \epsilon_{s+1} \quad (16) \\ &g_{\tilde{\pi}}(\hat{\theta}, \beta, V) \leq 0, \end{aligned}$$

where $\mathcal{J}_{\tilde{\pi}}$ is the average cost of (14) when controlled with $\tilde{\pi}$ and c is the cost of the original LQ. We also denote by $\mathcal{J}_*(\hat{\theta}, \beta, V)$ the minimum of (16). Once the constrained LQR is solved, the component $\tilde{\pi}^u$ relative to the variable u is used to control the real system for the whole episode until the termination condition is met. When $\tilde{\pi}$ is linear, we denote by \tilde{K} the associated gain (i.e., $\tilde{\pi}(x) = \tilde{K}x$), and we use K_u (resp. K_w) to refer to the block of \tilde{K} corresponding to the control u (resp. the perturbation control w).

4.2. Optimism and Regret

We show that the optimization in (16) preserves the learning guarantees of the original OFU-LQ algorithm at the cost of a slightly stronger requirement on ϵ_0 . This is not obvious as (16) is relaxing the constraints imposed by the confidence set used in (8) and solving the extended LQR might lead to a perturbation control $\tilde{\pi}^w$ that does not actually correspond to any feasible model θ in \mathcal{C} . Intuitively, we need the constraint $g_{\tilde{\pi}}$ to be loose enough so as to guarantee optimism and tight enough to preserve good regret performance. We start by showing that optimizing (16) gives an optimistic solution.

Lemma 3. *Under Asm. 2, and 3, whenever $\theta_* \in \mathcal{C}$, the optimal solution to (16) is optimistic, i.e.,*

$$\mathcal{J}_*(\hat{\theta}, \beta, V) \leq J_*, \quad (17)$$

The lemma above shows that the optimal controller in the extended LQR has an average cost (in the extended LQR) that is smaller than the true optimal average cost, thus certifying

its optimistic nature of (16). This is expected, since (16) is a relaxed version of the original OFU-LQ++ problem, which returns optimistic solutions by definition. Then we show that applying the optimistic extended controllers induce a sequentially stable state process.

Lemma 4. *Given the same system identification phase and RLS estimator of OFU-LQ++ (see Lem. 1), let $\{\tilde{K}_t\}_{t \geq 1}$ be the sequence of extended optimistic controllers generated by solving (16) and $\{x_t\}_{t \geq 0}$ be the state process (Eq. 1) induced when by the sequence of controllers $\{K_{u,t}\}_{t \geq 0}$. If $\epsilon_0 \leq O(1/\kappa^{3/2})$, then with probability at least $1 - \delta/2$, for all $t \leq T$,*

$$\begin{cases} \theta_* \in \mathcal{C}_t \\ \|x_t\| \leq X := 20\sigma\sqrt{\kappa\|P_*\|_2 \log(4T/\delta)/\lambda_{\min}(C)} \end{cases} \quad (18)$$

This lemma is the counterpart of Lem. 1 for the extended LQR and it illustrates that, due to the relaxed constraint, the condition on ϵ_0 is tighter by a factor $1/\sqrt{\kappa}$, while the bound on the state remains the same and this, in turn, leads to the same regret as OFU-LQ++ (Lem. 2) but for problem dependent constants.

Theorem 1. *Let (A_*, B_*, Q, R) be any LQR satisfying Asm. 1, 2, and 3. If the conditions in Lem. 4 are satisfied and the extended LQR with relaxed constrained (16) is solved exactly at each episode, then w.p. at least $1 - \delta$,*

$$\mathcal{R}(T) = \tilde{O}((n+d)\sqrt{n\kappa^{3/2}\|P_*\|_2^2\sqrt{T}}). \quad (19)$$

5. Efficient Solution to the Constrained Extended LQR via Lagrangian Relaxation

We introduce the Lagrangian formulation of (16). Let $\mu \in \mathbb{R}$ be the Lagrangian parameter, we define

$$\mathcal{L}_{\tilde{\pi}}(\hat{\theta}, \beta, V; \mu) := \mathcal{J}_{\tilde{\pi}}(\hat{\theta}, \beta, V) + \mu g_{\tilde{\pi}}(\hat{\theta}, \beta, V). \quad (20)$$

Since both average cost $\mathcal{J}_{\tilde{\pi}}$ and constraint $g_{\tilde{\pi}}$ measure asymptotic average quantities, we can conveniently define the matrices (C_{\dagger} being the bordering of matrix C in (2))

$$C_{\dagger} = \begin{pmatrix} Q & 0 & 0 \\ 0 & R & 0 \\ 0 & 0 & 0 \end{pmatrix}; \quad C_g = \begin{pmatrix} -\beta^2 V^{-1} & 0 \\ 0 & I \end{pmatrix},$$

and write the Lagrangian as

$$\mathcal{L}_{\tilde{\pi}}(\hat{\theta}, \beta, V; \mu) = \lim_{S \rightarrow \infty} \frac{1}{S} \mathbb{E} \left[\sum_{s=0}^{S-1} \begin{pmatrix} x_s^T & \tilde{u}_s^T \end{pmatrix} C_{\mu} \begin{pmatrix} x_s \\ \tilde{u}_s \end{pmatrix} \right],$$

with $C_{\mu} = C_{\dagger} + \mu C_g$. This formulation shows that $\mathcal{L}_{\tilde{\pi}}(\mu)$ can be seen as the average cost of an extended LQR problem with state x_s , control \tilde{u}_s , linear dynamics (\hat{A}, \tilde{B}) and quadratic cost with matrix C_{μ} . As a result, we introduce

the *Lagrangian extended LQR* problem associated to the extended LQR with relaxed constraints of (16) as

$$\begin{aligned} \mathcal{L}_*(\hat{\theta}, \beta, V) &= \sup_{\mu \in \mathcal{M}} \min_{\tilde{\pi}} \mathcal{L}_{\tilde{\pi}}(\hat{\theta}, \beta, V; \mu) \\ \text{subject to} \quad &x_{s+1} = \hat{A}x_s + \tilde{B}\tilde{u}_s + \epsilon_{s+1} \end{aligned}, \quad (21)$$

where $\mathcal{M} = [0, \tilde{\mu}]$ is the domain of the Lagrangian parameter (more details on $\tilde{\mu}$ are reported in App. G.2). We prove the following fundamental result.

Theorem 2. *For any extended LQR parametrized by $\hat{\theta}$, V , β , and psd cost matrices Q, R , there exists a domain $\mathcal{M} = [0, \tilde{\mu}]$ with $\tilde{\mu} \in \mathbb{R}_+$, such that strong duality between the relaxed optimistic optimization in (16) and its Lagrangian formulation (21) holds:*

$$\mathcal{J}_*(\hat{\theta}, \beta, V) = \mathcal{L}_*(\hat{\theta}, \beta, V).$$

Supported by the strong duality above, we provide a more detailed characterization of $\mathcal{L}_{\tilde{\pi}}(\hat{\theta}, \beta, V; \mu)$, which motivates the design of an efficient algorithm to optimize over $\tilde{\pi}$, μ . For ease of notation, in the following we consider $\hat{\theta}$, β , and V as fixed and we drop them from the definition of $\mathcal{L}_{\tilde{\pi}}(\mu)$, which we study as a function of $\tilde{\pi}$ and μ .

5.1. The Lagrangian Dual Function

We introduce the Lagrangian dual function, for any $\mu \in \mathbb{R}_+$,

$$\begin{aligned} \mathcal{D}(\mu) &= \min_{\tilde{\pi}} \mathcal{L}_{\tilde{\pi}}(\mu) \\ \text{s.t.} \quad &x_{s+1} = \hat{A}x_s + \tilde{B}\tilde{u}_s + \epsilon_{s+1} \end{aligned}, \quad (22)$$

and we denote by $\tilde{\pi}_\mu$ the corresponding extended optimal policy. For *small enough* μ , the cost matrix C_μ is p.d., which allows solving (22) using standard Riccati theory. The main technical challenge arises for larger values of μ when the solution of the dual Lagrangian function may not be computable by Riccati equations or it may not even be defined. Fortunately, the following lemma shows that within the domain \mathcal{M} where Thm. 2 holds, there always exists a Riccati solution for (22).

Lemma 5. *For any extended LQR parametrized by $\hat{\theta}$, V , β , and psd cost matrices Q, R , consider the domain $\mathcal{M} = [0, \tilde{\mu}]$ where Thm. 2 holds, then for any $\mu \in \mathcal{M}$*

1. *The extended LQ in (22) is controllable and it admits a unique solution*

$$\tilde{\pi}_\mu = \arg \min_{\tilde{\pi}} \mathcal{L}_{\tilde{\pi}}(\mu), \quad (23)$$

obtained by solving the generalized discrete algebraic Riccati equation (DARE) (Molinari, 1975)⁷ associated with the Lagrange LQR $(\hat{A}, \tilde{B}, C_\mu)$. Let

$C_\mu = (R_\mu \ N_\mu; N_\mu \ Q_\mu)$ be the canonical formulation for the cost matrix, then

$$\begin{aligned} D_\mu &= R_\mu + \tilde{B}^\top P_\mu \tilde{B} \\ P_\mu &= Q_\mu + A^\top P_\mu A \\ &\quad - [A^\top P_\mu \tilde{B} + N_\mu^\top] D_\mu^{-1} [\tilde{B}^\top P_\mu A + N_\mu], \end{aligned} \quad (24)$$

and the optimal control is $\tilde{K}_\mu = -D_\mu^{-1}[\tilde{B}^\top P_\mu A + N_\mu]$, while the dual function is $\mathcal{D}(\mu) = \text{Tr}(P_\mu)$.

2. *$\mathcal{D}(\mu)$ is concave and continuously differentiable.*
3. *The derivative $\mathcal{D}'(\mu) = g_{\tilde{\pi}_\mu}$, i.e., it is equal to the constraint evaluated at the optimal extended policy for μ . As a result, the Lagrangian dual can be written as*

$$\mathcal{L}_{\tilde{\pi}_\mu}(\mu) = \mathcal{D}(\mu) = \mathcal{J}_{\tilde{\pi}_\mu} + \mu \mathcal{D}'(\mu). \quad (25)$$

The previous lemma implies that in order to solve (21) we may need to evaluate the dual function $\mathcal{D}(\mu)$ only where the the optimal control can be computed by solving a DARE.

Since $\mathcal{D}(\mu)$ is concave and smooth, we can envision using a simple dichotomy approach to optimize $\mathcal{D}(\mu)$ over \mathcal{M} and solve (21). Nonetheless, we notice that Thm. 2 only provides strong duality in a sup/min sense, which means that the optimum may not be attainable within \mathcal{M} . Furthermore, even when there exists a maximum, computing an ϵ -optimal solution in terms of the Lagrangian formulation, i.e., finding a pair $\mu, \tilde{\pi}$ such that $|\mathcal{L}_{\tilde{\pi}}(\mu) - \mathcal{L}_*| \leq \epsilon$, may not directly translate in a policy with desirable performance in terms of its average cost $\mathcal{J}_{\tilde{\pi}}$ and feasibility w.r.t. the constraint $g_{\tilde{\pi}}$.

We illustrate this issue in the example in Fig. 1. We display a qualitative plot of the Lagrangian dual $\mathcal{D}(\mu)$ and its derivative $\mathcal{D}'(\mu)$ when (21) admits a maximum at μ^* and the dichotomy search returned values μ_l and μ_r that are ϵ -close and $\mu^* \in [\mu_l, \mu_r]$. We consider the case where the algorithm returns μ_l as the candidate solution. By concavity and the fact that $\mathcal{D}'(0) > 0$, the function $\mathcal{D}(\mu)$ is Lipschitz in the interval $[0, \mu^*]$ with constant bounded by $\mathcal{D}'(0)$. Thus the accuracy $\mu^* - \mu_l \leq \epsilon$ translates into an equivalent ϵ -optimality in \mathcal{D} (i.e., $\mathcal{D}(\mu^*) - \mathcal{D}(\mu_l) = \mathcal{L}^* - \mathcal{L}_{\tilde{\pi}_{\mu_l}}(\mu_l) \leq \mathcal{D}'(0)\epsilon$). Nonetheless, this does not imply a similar guarantee for $\mathcal{D}'(\mu)$. If the second derivative of $\mathcal{D}(\mu)$ (i.e., the curvature of the function) is large close to μ^* , the original error ϵ can be greatly amplified when evaluating $\mathcal{D}'(\mu_l)$. For instance, if $\mathcal{D}''(\mu) \gg 1/\epsilon$, then $\mathcal{D}'(\mu_l) = \Omega(1)$. Given the last point of Lem. 5, this means that despite returning an ϵ -optimal solution in the sense of (21), $\tilde{\pi}_{\mu_l}$ may significantly violate the constraint (as $\mathcal{D}'(\mu_l) = g_{\tilde{\pi}_{\mu_l}} = \Omega(1)$). While Eq. (25) implies that $\tilde{\pi}_{\mu_l}$ is still optimistic (i.e., $\mathcal{J}_{\tilde{\pi}_{\mu_l}} \leq J^*$, as in Lem. 3), the regret accumulated by $\tilde{\pi}_{\mu_l}$ cannot be controlled anymore, since the perturbation control w_s may be arbitrarily outside the confidence interval. Interestingly, the

⁷The need for generalized DARE is due to the fact that for some $\mu \in \mathcal{M}$, the associated cost C_μ may not be p.s.d.

curvature becomes larger and larger as the optimum shifts to the extremum of \mathcal{M} and, in the limit, (21) only admits a supremum. In this case, no matter how close μ_l is to μ^* , the associated policy $\tilde{\pi}_{\mu_l}$ may perform arbitrarily bad.

More formally, we have the following lemma (the explicit value of α is reported in Lem. 14).

Lemma 6. *For any LQR parametrized by $\hat{\theta}$, V , β , and psd cost matrices Q, R , consider the domain $\mathcal{M} = [0, \tilde{\mu}]$ where Thm. 2 holds. Let \mathcal{M}_+ be a subset of \mathcal{M} such that $\mathcal{M}_+ = \{\mu \in \mathcal{M} \text{ s.t. } \mathcal{D}'(\mu) \geq 0\}$. Then, \mathcal{D} has Lipschitz gradient, i.e., there exists a constant α depending on $\hat{\theta}$, V , β , and the cost matrices Q, R , such that for all $(\mu_1, \mu_2) \in \mathcal{M}_+^2$,*

$$|\mathcal{D}'(\mu_1) - \mathcal{D}'(\mu_2)| \leq |\mu_1 - \mu_2| \frac{\alpha}{\lambda_{\min}(D_{\mu_1})},$$

where D_{μ} is defined in (24).

This result shows that even when $|\mu_1 - \mu_2| \leq \epsilon$, the difference in gradients may be arbitrarily large when $\lambda_{\min}(D_{\mu}) \ll \epsilon$ (i.e., large curvature). In the next section, we build on this lemma to craft an adaptive stopping condition for the dichotomy search and to detect that case of large curvature.

5.2. An Efficient Dichotomy Search

The algorithm we propose, DS-OFU, seeks to find a value of μ of zero gradient $\mathcal{D}'(\mu)$ by dichotomy search. While $\mathcal{D}(\mu)$ is a 1-dim function and Lem. 5 guarantees that it is concave in \mathcal{M} , there are three major challenges to address: **1)** Thm. 2 does not provide any explicit value for $\tilde{\mu}$; **2)** The algorithm needs to evaluate $\mathcal{D}'(\mu)$; **3)** For any ϵ , DS-OFU must return a policy $\tilde{\pi}_{\epsilon}$ that is ϵ -optimistic and ϵ -feasible for the extended LQR with relaxed constraints (16).

DS-OFU starts by checking the sign of the gradient $\mathcal{D}'(0)$. If $\mathcal{D}'(0) \leq 0$, the algorithm ends and outputs the optimal policy $\tilde{\pi}_0$ since by concavity 0 is the arg-max of \mathcal{D} and $\tilde{\pi}_0$ is the exact solution to (21). If $\mathcal{D}'(0) > 0$, the dichotomy starts with accuracy ϵ and a valid⁸ search interval $[0, \mu_{\max}]$, where μ_{\max} is defined as follows.

Lemma 7. *Let $\mu_{\max} := \beta^{-2} \lambda_{\max}(C) \lambda_{\max}(V)$, then $\mathcal{D}'(\mu_{\max}) < 0$.*

The previous lemma does not imply that $[0, \mu_{\max}] \supseteq \mathcal{M}$, but it provides an explicit value of μ with negative gradient, thus defining a bounded and valid search interval for the dichotomy process. At each iteration, DS-OFU updates either μ_l or μ_r so that the interval $[\mu_l, \mu_r]$ is always valid.

The second challenge is addressed in the following proposition, which illustrates how the derivative $\mathcal{D}'(\mu)$ (equivalently the constraint $g_{\tilde{\pi}_{\mu}}$) can be efficiently computed.

⁸We say that $[\mu_l, \mu_r]$ is valid if $\mathcal{D}'(\mu_l) \geq 0$ and $\mathcal{D}'(\mu_r) \leq 0$.

Proposition 3. *For any $\mu \in \mathcal{M}$, let $\tilde{\pi}_{\mu}$ (Eq. 23) have an associated controller \tilde{K}_{μ} that induces a closed-loop dynamics $A^c(\tilde{K}_{\mu}) = \hat{A} + \tilde{B}\tilde{K}_{\mu}$ then $\mathcal{D}'(\mu) = g_{\tilde{\pi}_{\mu}} = \text{Tr}(G_{\mu})$, where G_{μ} is the unique solutions of the Lyapunov equation*

$$G_{\mu} = (A^c(\tilde{K}_{\mu}))^{\top} G_{\mu} A^c(\tilde{K}_{\mu}) + \begin{pmatrix} I \\ \tilde{K}_{\mu} \end{pmatrix}^{\top} C_g \begin{pmatrix} I \\ \tilde{K}_{\mu} \end{pmatrix},$$

This directly from the fact that $g_{\tilde{\pi}}$ is an asymptotic average quadratic quantity (as much as the average cost J), and it is thus the solution of a Lyapunov equation of dimension n .

The remaining key challenge is to design an adaptive stopping condition that is able to keep refining the interval $[\mu_l, \mu_r]$ until either an accurate enough solution is returned, or, the curvature is too large (or even infinite). In the latter case, the algorithm switches to a failure mode, for which we design an ad-hoc solution.

Since the objective is to achieve an ϵ -feasible solution (i.e., $g_{\tilde{\pi}_{\epsilon}} \leq \epsilon$), we leverage Lem. 6 and we interrupt the dichotomy process whenever $(\mu_r - \mu_l)\alpha/\lambda_{\min}(D_{\mu_l}) \leq \epsilon$. Nonetheless, when the optimum of (21) is not attainable in \mathcal{M} , the previous stopping condition may never be verified and the algorithm would never stop. As a result, we interrupt the dichotomy process when $\lambda_{\min}(D_{\mu_l}) \leq \lambda_0 \epsilon^2$ for a given constant λ_0 . In this case, the dichotomy fails to return a viable solution and we need to revert to a *backup* strategy, which consists in either modifying the controller found at μ_l or applying a suitable perturbation to the original cost matrix C_{\dagger} . In the latter case, we design a perturbation such that **1)** the optimization problem (21) associated to the system with the perturbed cost C'_{\dagger} admits a maximum and can be efficiently solved by the same dichotomy process illustrated before and **2)** the corresponding solution $\tilde{\pi}'$ is ϵ -optimistic and ϵ -feasible in the *original* system. The explicit backup strategy is reported in App. I.

Theorem 3. *For any LQR parametrized by $\hat{\theta}$, V , β , and psd cost matrices Q, R , and any accuracy $\epsilon \in (0, 1/2)$, there exists values of α and λ_0 and a backup strategy such that*

1. DS-OFU outputs an ϵ -optimistic and ϵ -feasible policy $\tilde{\pi}_{\epsilon}$ given by the linear controller \tilde{K}_{ϵ} such that

$$\mathcal{J}_{\tilde{\pi}_{\epsilon}} \leq \mathcal{J}_{*} + \epsilon \quad \text{and} \quad g_{\tilde{\pi}_{\epsilon}} \leq \epsilon.$$

2. DS-OFU terminates within at most $N = O(\log(\mu_{\max}/\epsilon))$ iterations, each solving one Riccati and one Lyapunov equation for the extended Lagrangian LQR, both with complexity $O(n^3)$.

This result shows that DS-OFU returns a solution to (16) at any level of accuracy ϵ in a total runtime $O(n^3 \log(1/\epsilon))$. We refer to the algorithm resulting by plugging DS-OFU into the OFU-LQ++ learning scheme as LAGLQ (Lagrangian-LQ). By running DS-OFU with $\epsilon = 1/\sqrt{t}$ provides the regret guarantee of Thm. 1.




Figure 1. Qualitative plots of $\mathcal{D}(\mu)$ and its derivative $\mathcal{D}'(\mu)$.

6. Discussion

We investigate the difference between confidence-based and isotropic exploration in term of complexity, bounds and empirical performance. While not conclusive, we believe this discussion sheds light on how confidence-based methods may be better at adapting to the structure of the problem. As a representative for isotropic exploration, we refer to CECCE (Simchowitz & Foster, 2020), which offers the tightest regret guarantee among the CE strategies. For confidence-based exploration, we discuss the guarantee of both OFU-LQ++ and LAGLQ, but limit the computational and experiment comparisons with LAGLQ only.

Computational complexity. Both LAGLQ and CECCE proceeds through episodes of increasing length. LAGLQ relies on the standard determinant-based rule ($\det(V_t) \geq 2 \det(V_{t_k})$) to decide when to stop, while in CECCE the length of each episode is twice longer than the previous one. In both cases, the length of the episodes is increasing exponentially over time, thus leading to $O(\log T)$ updates. Given the complexity analysis in Thm. 3, we notice that DS-OFU and computing the CE controller have the same order of complexity, where CECCE solves *one* Riccati equation, while DS-OFU solves as many as $\log(1/\epsilon)$ Riccati and Lyapunov equations. On systems of moderate side and given the small number of recomputations, the difference between the two approaches is relatively narrow.

Regret. We limit the comparison to the main order term $\tilde{O}(\sqrt{T})$ and the dependencies on dimensions n and d , and problem-dependent constants such as κ and $\|P^*\|_2$,

```

Input:  $\hat{\theta}, \beta, V, \epsilon, \alpha, \lambda_0$ 
1: if  $\mathcal{D}(0) \leq 0$  then
2:   Set  $\bar{\mu} = 0$  and  $\tilde{\pi}_\epsilon = \tilde{\pi}_{\bar{\mu}}$ 
3: else
4:   Set  $\mu_l = 0, \mu_r = \mu_{\max}$  (Lem. 7)
5:   while  $\alpha \frac{\mu_r - \mu_l}{\lambda_{\min}(D_{\mu_l})} \geq \epsilon$  or  $\lambda_{\min}(D_{\mu_l}) \geq \lambda_0 \epsilon^2$  do
6:     Set  $\bar{\mu} = (\mu_l + \mu_r)/2$ 
7:     if  $\mathcal{D}'(\bar{\mu}) > 0$  then
8:        $\mu_l = \bar{\mu}$ 
9:     else
10:       $\mu_r = \bar{\mu}$ 
11:    end if
12:  end while
13: end if
14: if  $\alpha \frac{\mu_r - \mu_l}{\lambda_{\min}(D_{\mu_l})} < \epsilon$  then
15:   Set  $\bar{\mu} = \mu_l$  and  $\tilde{\pi}_\epsilon = \tilde{\pi}_{\bar{\mu}}$ 
16: else
17:   Set  $\tilde{\pi}_\epsilon$  to the control return by the backup procedure
18: end if
19: return Control policy  $\tilde{\pi}_\epsilon$ 
    
```

Figure 2. The DS-OFU algorithm to solve (21).

$$\begin{aligned} \mathcal{R}_{\text{CECCE}} &= \tilde{O}(\|P^*\|_2^{11/2} d \sqrt{nT}), \\ \mathcal{R}_{\text{OFU-LQ++}} &= \tilde{O}(\kappa^{1/2} \|P^*\|_2 (n+d) \sqrt{n} \sqrt{T}), \\ \mathcal{R}_{\text{LAGLQ}} &= \tilde{O}(\kappa^{3/2} \|P^*\|_2^2 (n+d) \sqrt{n} \sqrt{T}). \end{aligned}$$

The first difference is that CECCE has worst-case optimal dependency $d\sqrt{n}$ on the dimension of the problem, while optimistic algorithms OFU-LQ++ and LAGLQ are slightly worse, scaling with $(n+d)\sqrt{n}$. While this shows that OFU-LQ++ and LAGLQ are worst-case optimal when $n \approx d$, it is an open question whether ϵ -greedy is by nature superior to confidence-based method when $d \ll n$ or whether it is due to a loose analysis. In fact, those dependencies are mostly inherited from the confidence intervals in (1) which is treated differently in (Simchowitz & Foster, 2020), thanks to a *refined* bound and a *different* regret decomposition. This suggests that a finer analysis for OFU-LQ++ and LAGLQ may close this gap.

The main difference lies in the dependency on complexity-related quantities κ and $\|P^*\|_2$. While there is no strict ordering between them,⁹ they both measure the cost of controlling the system. In this respect, CECCE suffers from a significantly larger dependency than optimistic algorithms: OFU-LQ++ offers the best performance while LAGLQ is slightly worse than OFU-LQ++, due to the use of a relaxed constraint to obtain tractability. We believe this difference in performance may be intrinsic in the fact that methods

⁹The definition of $\kappa = D/\lambda_{\min}Q$, with $D \geq \text{Tr}(P^*)$ may suggest $\kappa > \|P^*\|_2$, but the smallest eigenvalue of Q may be large enough so that $\kappa \leq \|P^*\|_2$.

based on isotropic perturbations of the CE are less effective in adapting to the actual structure of the problem. As the isotropic perturbation is tuned to guarantee a sufficient estimation in all directions of the state-control space, it leads to over-exploration w.r.t. some directions as soon as there is an asymmetry in the cost sensitivity in the estimation error. On the other hand, OFU-LQ++ and LAGLQ further leverage this asymmetry from the confidence set, and by optimism, do not waste exploration to learn accurately directions which have little to no impact on the performance.

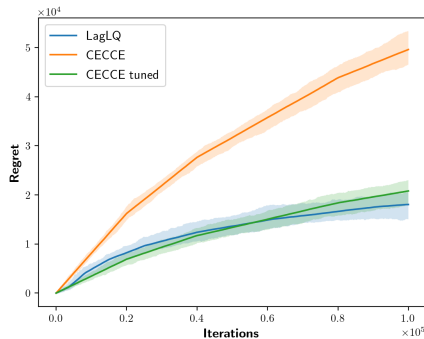


Figure 3. Regret curves for CECCE and LAGLQ.

Empirical comparison. We conclude with a simple numerical simulation (details in App. J). We compare CECCE with the variance parameter (σ_{in}^2) set as suggested in the original paper and a tuned version where we shrink it by a factor $\sqrt{\|P_*\|_2}$, and LAGLQ where the confidence interval is set according to (1). Both algorithms receive the same set Θ_0 obtained from an initial system identification phase. In Fig. 3 we see that LAGLQ performs better than both the original and tuned versions of CECCE. More interestingly, while CECCE is “constrained” to have a $O(\sqrt{T})$ regret by the definition of the perturbation itself, which scales as $1/\sqrt{t}$, it seems LAGLQ’s regret is $o(\sqrt{T})$, suggesting that despite the worst-case lower bound $\Omega(\sqrt{T})$, LAGLQ may be adapt to the structure of the problem and achieve better regret.

References

- Abbasi-Yadkori, Y. and Szepesvári, C. Regret bounds for the adaptive control of linear quadratic systems. In *COLT*, pp. 1–26, 2011.
- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. Online least squares estimation with self-normalized processes: An application to bandit problems. *arXiv preprint arXiv:1102.2670*, 2011.
- Abeille, M. and Lazaric, A. Improved regret bounds for thompson sampling in linear quadratic control problems. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, pp. 1–9, 2018.
- Bittanti, S., Campi, M., et al. Adaptive control of linear time invariant systems: the “bet on the best” principle. *Communications in Information & Systems*, 6(4):299–320, 2006.
- Cohen, A., Koren, T., and Mansour, Y. Learning linear-quadratic regulators efficiently with only \sqrt{T} regret. *CoRR*, abs/1902.06223, 2019. URL <http://arxiv.org/abs/1902.06223>.
- Dean, S., Mania, H., Matni, N., Recht, B., and Tu, S. Regret bounds for robust adaptive control of the linear quadratic regulator. *CoRR*, abs/1805.09388, 2018. URL <http://arxiv.org/abs/1805.09388>.
- Faradonbeh, M. K. S., Tewari, A., and Michailidis, G. Finite time analysis of optimal adaptive policies for linear-quadratic systems. *CoRR*, abs/1711.07230, 2017. URL <http://arxiv.org/abs/1711.07230>.
- Faradonbeh, M. K. S., Tewari, A., and Michailidis, G. Finite time adaptive stabilization of LQ systems. *CoRR*, abs/1807.09120, 2018a. URL <http://arxiv.org/abs/1807.09120>.
- Faradonbeh, M. K. S., Tewari, A., and Michailidis, G. Input perturbations for adaptive regulation and learning. *CoRR*, abs/1811.04258, 2018b. URL <http://arxiv.org/abs/1811.04258>.
- Jaksch, T., Ortner, R., and Auer, P. Near-optimal regret bounds for reinforcement learning. *J. Mach. Learn. Res.*, 11:1563–1600, August 2010.
- Lancaster, P. and Rodman, L. *Algebraic riccati equations*. Oxford University Press, 1995.
- Mania, H., Tu, S., and Recht, B. Certainty Equivalent Control of LQR is Efficient. *arXiv e-prints*, art. arXiv:1902.07826, Feb 2019.
- Molinari, B. P. The stabilizing solution of the discrete algebraic riccati equation. *Automatic Control, IEEE Transactions on*, 20(3):396–399, Jun 1975.
- Ouyang, Y., Gagrani, M., and Jain, R. Learning-based control of unknown linear systems with thompson sampling. *CoRR*, abs/1709.04047, 2017. URL <http://arxiv.org/abs/1709.04047>.
- Rockafellar, R. T. *Convex analysis*. Number 28. Princeton university press, 1970.
- Rugh, W. J. *Linear system theory*, volume 2. prentice hall Upper Saddle River, NJ, 1996.
- Simchowitz, M. and Foster, D. J. Naive Exploration is Optimal for Online LQR. *arXiv e-prints*, art. arXiv:2001.09576, Jan 2020.
- Van Dooren, P. A generalized eigenvalue approach for solving riccati equations. *SIAM Journal on Scientific and Statistical Computing*, 2(2):121–135, 1981.

A. Notation

We recall the basic notation defined in the paper and introduce additional convenient notation.

Linear algebra.

- $\lambda(A)$: the spectrum of the matrix A
- $\rho(A)$: the spectral radius of a matrix A , $\rho(A) = \max |\lambda(A)|$
- $\lambda_{\min}(M)$, $\lambda_{\max}(M)$: the minimum and maximum eigenvalue of a symmetric matrix M
- $Im(A)$, $\ker(A)$: the row and null space of a matrix A
- $|x|$: the absolute value (resp. the modulus) of $x \in \mathbb{R}$ (resp. $x \in \mathbb{C}$)
- $\|x\|$: the euclidian norm of a vector x
- $\|A\|_2$: the 2– norm of a matrix A , $\|A\|_2 = \max_{\|x\|=1} \|Ax\|$
- $\|A\|_F$: the Frobenius norm of a matrix A , $\|A\|_F = \sqrt{\text{Tr}(A^T A)}$
- $\|A\|_M$: the weighted Frobenius norm of A wrt to a psd matrix M , $\|A\|_M = \|M^{1/2} A\|_F$

Original LQR problem.

- n, d : the dimensions of the state and input variables
- x_t, u_t, ϵ_t : the state/input/noise variables at time t
- z_t : aggregated state and input variable at time t , $z_t^T = (x_t^T, u_t^T)$
- A_*, B_* : ground truth for the state dynamics parameters
- Q, R : cost matrices of the LQR
- C : aggregated cost matrix, $C = (Q, 0; 0, R)$
- π, K : an arbitrary linear policy mapping states to inputs, and its matrix representation.
- θ : aggregated state dynamics parametrization, $\theta^T = (A, B)$
- $P(\theta)$: solution of the Riccati equation associated with the LQR system parametrized by θ and C
- $K(\theta)$: optimal controller for the LQR system parametrized by θ and C
- $A^c(\theta, K)$: the closed-loop matrix of a dynamical system parametrized by θ controlled by K , $A^c(\theta, K) = \theta^T \begin{pmatrix} I \\ K \end{pmatrix}$
- $A^c(\theta)$: the closed-loop matrix of a dynamical system parametrized by θ controlled by $K(\theta)$, $A^c(\theta) = A^c(\theta, K(\theta))$
- $\Sigma(\theta, K)$: The steady-state covariance of the state process driven by $A^c(\theta, K)$
- $J_\pi(\theta)$: the infinite horizon cost of an LQR parametrized by θ and controlled by π
- $J(\theta)$: the infinite horizon cost of an LQR parametrized by θ and optimally controlled by $K(\theta)$
- J_* : the optimal infinite horizon cost for the true LQR system, $J_* = J(\theta_*)$

Extended/Lagrangian LQR problem.

- w_t : the input perturbation at time t
- \tilde{u}_t : the extended input at time t , $\tilde{u}_t^\top = (u_t^\top, w_t^\top)$
- C_\dagger, C_g : the cost matrices of the extended LQR objective and constraint
- \tilde{B} : the extended input matrix, $\tilde{B} = (B, I)$
- $\tilde{\pi}, \tilde{K}$: an arbitrary linear extended policy mapping states to extended inputs, and its matrix representation
- $\mathcal{J}_{\tilde{\pi}}(\theta, \beta, V)$: the infinite horizon cost of the constrained extended LQR parametrized by (θ, β, V) and controlled by $\tilde{\pi}$
- \mathcal{J}_* : the optimal infinite horizon cost of the constrained extended LQR.
- $g_{\tilde{\pi}}(\theta, \beta, V)$: the infinite horizon constraint value of the constrained extended LQR parametrized by (θ, β, V) and controlled by $\tilde{\pi}$
- $\mathcal{L}_{\tilde{\pi}}(\theta, \beta, V; \cdot)$: the Lagrangian function associated with the extended constrained LQR
- $\mathcal{D}(\cdot)$: the dual function associated with the Lagrangian relaxation

B. Structure of the Appendix

The appendix is organized over different sections:

- App. C recalls elements of Lyapunov stability theory and concentration inequality (high-probability events) (Prop. 1).
- App. D contains the proofs of OFU-LQ++ (sequential stability Lem. 1 and regret Lem. 2).
- App. E contains the proofs of LAGLQ (optimism Lem. 3, sequential stability Lem. 4, and regret Thm. 1) .
- Appendices F,G, H reports the strong-duality proofs (Thm. 2, Lem. 5, Lem. 6, Lem. 7, Prop. 3).
- App. I details the algorithms and the proof of Thm. 3.
- App. J details the experimental protocol used to obtain Fig. 3.

At the beginning of each part we recall and complement the notation used in that specific part.

C. Preliminaries

C.1. Lyapunov Stability

We recall basic results from Lyapunov stability theory. Most of the definitions below are adapted from [Rugh \(1996\)](#). We consider a general discrete-time deterministic linear dynamical process $\{\bar{x}_t\}_{t \geq 0}$ with dynamics A and initial state ξ defined as

$$\bar{x}_{t+1} = A\bar{x}_t, \quad \bar{x}_0 = \xi. \quad (26)$$

We recall the following definition of stability.

Definition 1. A deterministic process (26) is uniformly exponentially stable (UES) if there exists a constant $\gamma > 0$ and a constant $0 \leq \lambda < 1$ such that for any $t \geq 0$,

$$\|\bar{x}_t\| \leq \gamma \lambda^{t-s} \|\bar{x}_s\|. \quad (27)$$

The stability of the deterministic process (26) can be translated into stability properties of the transition matrix A . The following propositions provide spectral characterizations.

Proposition 4 (Thm. 22.1 ([Rugh, 1996](#))). A time-invariant deterministic process (26) is uniformly exponentially stable if and only if all eigenvalues of A have magnitude strictly smaller than 1, i.e., $\rho(A) < 1$.

Proposition 5 (Thm. 22.7 ([Rugh, 1996](#))). The process (26) is uniformly exponentially stable (UES) if and only if there exists a constant $\gamma > 0$ and a constant $0 \leq \lambda < 1$ such that for all $t \geq s$, $\|A^{t-s}\| \leq \gamma \lambda^{t-s}$.

An alternative characterization is provided by Lyapunov stability theory.

Proposition 6 (Thm. 23.3 ([Rugh, 1996](#))). A deterministic process (26) is uniformly exponentially stable if and only if there exists a symmetric ps.d. matrix Σ such that

$$\Sigma = A\Sigma A^\top + I. \quad (28)$$

Further, let ρ be finite positive constants such that $\Sigma \preceq \rho I$, then for any $t \geq s \geq 0$, we have $\|A^{t-s}\| \leq \gamma \lambda^{t-s}$ with

$$\gamma = \sqrt{\rho}, \quad \lambda = 1 - 1/\rho. \quad (29)$$

Uniform exponential stability for a deterministic process directly implies boundedness for a perturbed version of the process with bounded perturbations.

Proposition 7. Consider the perturbed instance $\{\tilde{x}_t\}_{t \geq 0}$ of the deterministic process (26), defined as

$$\tilde{x}_{t+1} = A\tilde{x}_t + \epsilon_t, \quad \tilde{x}_0 = \xi, \quad \text{for all } t \geq 0. \quad (30)$$

where $\{\epsilon_t\}_{t \geq 0}$ is a martingale difference sequence with zero-mean, σ -subGaussian increments. If the deterministic process is UES associated with a psd matrix Σ given in (28), the perturbed process $\{\tilde{x}_t\}_{t \geq 0}$ is uniformly bounded, that is, with probability at least $1 - \delta$,

$$\text{for all } t \geq 0, \quad \|\tilde{x}_t - A^t \xi\|^2 \leq 8\sigma^2 \text{Tr}(\Sigma) \log(1/\delta). \quad (31)$$

Proof of Prop. 7. First we use (30) to reformulate the perturbed process as

$$\tilde{x}_t = A^t \xi + \sum_{s=1}^t A^{t-s} \epsilon_{s-1}. \quad (32)$$

Then, we follow similar steps as in [Abbasi-Yadkori et al. \(2011\)](#): we begin by constructing an exponential super-martingale. Formally, for all $\alpha \in \mathbb{R}^n$,

$$M_t^\alpha = \exp \left(\alpha^\top (\tilde{x}_t - A^t \xi) - \frac{\sigma^2}{2} \|\alpha\|_{\sum_{s=1}^t A^{t-s} (A^{t-s})^\top}^2 \right),$$

is a super-martingale, and $\mathbb{E}(M_t^\alpha) \leq 1$. By (32), we have for any $s \geq 0$,

$$\mathbb{E}(M_s^\alpha | \mathcal{F}_{s-1}) = M_{s-1}^\alpha \mathbb{E} \left(\alpha^\top A^{t-s} \epsilon_{s-1} - \frac{\sigma^2}{2} \alpha^\top A^{t-s} (A^{t-s})^\top \alpha \middle| \mathcal{F}_{s-1} \right) \leq M_{s-1}^\alpha,$$

which ensures that M_t^α is a super-martingale while $M_0^\alpha = 1$ ensures that $\mathbb{E}(M_t^\alpha) \leq 1$. Further, recursively applying (28) ensures that:

$$\text{for all } t \geq 0, \quad \sum_{s=1}^t A^{t-s} (A^{t-s})^\top \preceq \Sigma.$$

Following Abbasi-Yadkori et al. (2011), we consider the random variable $\mathcal{A} \sim \mathcal{N}(0, \sigma^{-2} \lambda_{\max}(\Sigma)^{-1} I)$, independent of the noise sequence $\{\epsilon_s\}_{s \geq 0}$ and obtain, by the tower rule, that for all $t \geq 0$, $M_t = \mathbb{E}_{\mathcal{A}}(M_t^{\mathcal{A}})$ is such that $\mathbb{E}(M_t) \leq 1$. Further,

$$\begin{aligned} M_t &\geq \sqrt{\frac{\det(\sigma^2 \lambda_{\max}(\Sigma) I)}{(2\pi)^n}} \int_{\alpha \in \mathbb{R}^n} \exp\left(\alpha^\top (\tilde{x}_t - A^t \xi) - \frac{\sigma^2}{2} \|\alpha\|_{\Sigma + \lambda_{\max}(\Sigma) I}^2\right) \\ &= \exp\left(\frac{1}{2\sigma^2} \|\tilde{x}_t - A^t \xi\|_{(\Sigma + \lambda_{\max}(\Sigma) I)^{-1}}^2\right) \sqrt{\frac{\det(\sigma^2 \lambda_{\max}(\Sigma) I)}{\det(\sigma^2 (\Sigma + \lambda_{\max}(\Sigma) I))}} \\ &= \exp\left(\frac{1}{2\sigma^2} \|\tilde{x}_t - A^t \xi\|_{(\Sigma + \lambda_{\max}(\Sigma) I)^{-1}}^2\right) \sqrt{\frac{\det(\lambda_{\max}(\Sigma) I)}{\det(\Sigma + \lambda_{\max}(\Sigma) I)}} \end{aligned}$$

We conclude as in Cor. 1 in (Abbasi-Yadkori et al., 2011) using Markov inequality and a stopping time construction to obtain that with probability at least $1 - \delta$,

$$\text{for all } t \geq 0, \quad \|\tilde{x}_t - A^t \xi\|^2 \leq 2\sigma^2 \lambda_{\max}(\Sigma) \log\left(\frac{\det(\Sigma + \lambda_{\max}(\Sigma) I)}{\det(\lambda_{\max}(\Sigma) I)}\right) + 4\sigma^2 \lambda_{\max}(\Sigma) \log(1/\delta).$$

Finally, from $\log(1+x) \leq x$, we obtain

$$\lambda_{\max}(\Sigma) \log\left(\frac{\det(\Sigma + \lambda_{\max}(\Sigma) I)}{\det(\lambda_{\max}(\Sigma) I)}\right) \leq \text{Tr}(\Sigma), \quad \lambda_{\max}(\Sigma) \leq \text{Tr}(\Sigma),$$

which leads to, with probability at least $1 - \delta$,

$$\text{for all } t \geq 0, \quad \|\tilde{x}_t - A^t \xi\|^2 \leq 8\sigma^2 \text{Tr}(\Sigma) \log(1/\delta).$$

□

C.2. Technical results on Lyapunov stability

We recall useful technical results on Lyapunov manipulations.

Proposition 8 (Thm. 23.7 (Rugh, 1996)). *Given a $n \times n$ matrix A , if there exists symmetric positive definite $n \times n$ matrices M and P satisfying the discrete time Lyapunov equation*

$$A^\top P A - P = -M \tag{33}$$

then all eigenvalues of A have magnitude strictly smaller than 1 (i.e., $\rho(A) < 1$) and A is stable. Conversely, if $\rho(A) < 1$, then for every p.s.d. matrix M , the solution of Eq. 33 exists and is unique.

While the above theorem is stated for some matrix A , in the following we rely extensively on a similar characterization but for A^\top . As A^\top and A have the same spectrum, Prop. 8 can be rephrased in term of A^\top .

Corollary 1. *Given a $n \times n$ matrix A , if there exists a symmetric, positive definite $n \times n$ matrices M and Σ satisfying the discrete time Lyapunov equation*

$$A \Sigma A^\top - \Sigma = -M \tag{34}$$

then all eigenvalues of A have magnitude strictly less than one (i.e., $\rho(A) < 1$) and A is stable. Conversely, if $\rho(A) < 1$, then for every p.s.d. matrix M , the solution of Eq. 33 exists and is unique.

Proposition 9. *Let A be stable and (Σ_1, M_1) , and (Σ_2, M_2) be pairs of p.s.d. matrices satisfying the Lyapunov equations*

$$A \Sigma_1 A^\top - \Sigma_1 = -M_1; \quad A \Sigma_2 A^\top - \Sigma_2 = -M_2.$$

If $M_1 \preceq M_2$, then $\Sigma_1 \preceq \Sigma_2$.

Proposition 10. Let A be stable and (Σ_1, M) , and $(\Sigma_2, \alpha M)$ be pairs of p.s.d. matrices satisfying the Lyapunov equations

$$A\Sigma_1A^\top - \Sigma_1 = -M; \quad A\Sigma_2A^\top - \Sigma_2 = -\alpha M.$$

Then, $\Sigma_2 = \alpha\Sigma_1$.

Proposition 11. Let A be stable and M be a psd matrix. Let P and Σ be solution the Lyapunov equations

$$A^\top PA - P = -M; \quad A\Sigma A^\top - \Sigma = -I.$$

Then, $\text{Tr}(P) = \text{Tr}(\Sigma M)$.

We finally provide a technical result that we will intensively use in appendices D.1 and E.2. This proposition shows that the stability property of some matrix A_0 will be preserved for all matrix A_1 sufficiently close to A_0 , where the distance between A_0 and A_1 is measured according to a Lyapunov metric associated with A_0 .

Proposition 12. Let $A_0 \in \mathbb{R}^{n \times n}$ be a stable matrix, and let Σ_0 be the steady-state covariance associated with A_0 satisfying the Lyapunov:

$$\Sigma_0 = A_0\Sigma_0A_0^\top + I.$$

For any matrix A_1 such that $\|(A_0 - A_1)^\top\|_{\Sigma_0} \leq \eta c_0$ where $c_0^2\eta(\eta + 2) < 1$ and $c_0 \geq \|A_0^\top\|_{\Sigma_0}$, then,

1. A_1 is stable, and we denote by Σ_1 its steady-state covariance matrix defined as the unique solution of

$$\Sigma_1 = A_1\Sigma_1A_1^\top + I,$$

2. $\Sigma_1 \preceq \frac{1}{1 - c_0^2\eta(\eta + 2)}\Sigma_0$.

Proof. 1. We first prove that A_1 is stable by Cor. 1. Algebraic manipulations lead to

$$\begin{aligned} A_1\Sigma_0A_1^\top &= A_0\Sigma_0A_0^\top + (A_1 - A_0)\Sigma_0(A_1 - A_0)^\top + A_0\Sigma_0(A_1 - A_0)^\top + (A_1 - A_0)\Sigma_0A_0^\top \\ &= \Sigma_0 - I + M, \end{aligned}$$

where $M = (A_1 - A_0)\Sigma_0(A_1 - A_0)^\top + A_0\Sigma_0(A_1 - A_0)^\top + (A_1 - A_0)\Sigma_0A_0^\top$. As a result, Cor. 1 guarantees the stability of A_1 whenever $\|M\|_2 < 1$. Further, this is asserted since

$$\begin{aligned} \|M\|_2 &\leq \|\Sigma_0^{1/2}(A_1 - A_0)^\top\|_2^2 + 2\|\Sigma_0^{1/2}(A_1 - A_0)^\top\|_2\|\Sigma_0^{1/2}A_0^\top\|_2 \leq \|(A_1 - A_0)^\top\|_{\Sigma_0}^2 + 2\|(A_1 - A_0)^\top\|_{\Sigma_0}\|A_0^\top\|_{\Sigma_0} \\ &\leq c_0^2(\eta^2 + 2\eta) < 1. \end{aligned}$$

2. We now characterize the distance in term of stability between A_0 and A_1 , showing how Σ_1 relates to Σ_0 (in a psd sense). Let $\Delta = \Sigma_1 - \Sigma_0$. Then, algebraic manipulations leads to

$$\begin{aligned} \Delta &= A_1\Delta A_1^\top + (A_1 - A_0)\Sigma_0(A_1 - A_0)^\top + (A_1 - A_0)\Sigma_0A_0^\top + A_0\Sigma_0(A_1 - A_0)^\top \\ &= A_1\Delta A_1^\top + M. \end{aligned}$$

From $M \preceq \|M\|_2 I$, we obtain that $\Delta \preceq Z$ where $Z = A_1ZA_1^\top + \|M\|_2 I$ by Prop. 9 which can be expressed as $Z = \|M\|_2 \Sigma_1$ by Prop. 10. As a result, we have

$$\Delta \preceq \|M\|_2 \Sigma_1 \quad \Rightarrow \quad \Sigma_1 \preceq \frac{1}{1 - \|M\|_2} \Sigma_0,$$

which concludes the proof since $\|M\|_2 \leq c_0^2(\eta^2 + 2\eta)$.

□

C.3. High-probability events

All the derivations provided in appendices D and E will be conducted under the high-probability events that θ_* belongs to the confidence ellipsoid associated with the RLS procedure. Formally, we consider the event

$$E_t = \left\{ \forall s \leq t, \theta_* \in C_s \right\},$$

where C_t is defined in (9). By Prop. 1 and from $\beta_t = \beta_t(\delta/4)$, E_T holds with probability at least $1 - \delta/4$. This result is standard and we report the proofs here for sake of completeness.

Proof of Prop. 1. The proof of Prop. 1 is essentially the same as the one of Abbasi-Yadkori & Szepesvári (2011) but for the way the RLS is regularized. We provide it for sake of completeness.

Using the expression of $\hat{\theta}_t$, one gets:

$$\hat{\theta}_t = V_t^{-1} \left(\theta_0 + \sum_{s=0}^{t-1} z_s x_{s+1}^\top \right) = V_t^{-1} \left(\lambda \theta_0 + \sum_{s=0}^{t-1} z_s z_s^\top \theta_* + \sum_{s=0}^{t-1} z_s \epsilon_{s+1}^\top \right) = V_t^{-1} \left(\lambda \theta_0 + (V_t - \lambda I) \theta_* + \sum_{s=0}^{t-1} z_s \epsilon_{s+1}^\top \right)$$

which leads to

$$\|\hat{\theta}_t - \theta_*\|_{V_t} \leq \lambda \|\theta_0 - \theta_*\|_{V_t^{-1}} + \|S_t\|_{V_t^{-1}},$$

where $S_t = \sum_{s=0}^{t-1} z_s \epsilon_{s+1}^\top$. We decompose $S_t = (S_t^1, \dots, S_t^n)$ as

$$S_t^i = \sum_{s=0}^{t-1} z_s \epsilon_{s+1}^i, \quad \forall i \in \{1, \dots, n\}$$

where $\{\epsilon_s^i\}_{s \geq 0}$ is a martingale difference sequence, with zero-mean σ -subGaussian increment. From Cor. 1 in (Abbasi-Yadkori et al., 2011), for any $0 < \delta < 1$, we have with probability at least $1 - \delta$:

$$\forall t \geq 0, \quad \|S_t^i\|_{V_t^{-1}} \leq \sigma \sqrt{2 \log \left(\frac{\det(V_t)^{1/2}}{\det(\lambda I)^{1/2} \delta} \right)}.$$

Finally, a union bound argument leads to:

$$\forall t \geq 0, \quad \|S_t\|_{V_t^{-1}} \leq \sigma \sqrt{2n \log \left(\frac{\det(V_t)^{1/2} n}{\det(\lambda I)^{1/2} \delta} \right)}.$$

which concludes the proof. □

D. Sequential Stability and Regret of OFU-LQ++

This section is dedicated to the proofs of Lem. 1 and Lem. 2. We first address the sequential stability of OFU-LQ++ (App. D.1) and then show how it translates in a regret bound that scales polynomially with the dimension and the problem dependent constant κ (App. D.2). All the results in the section are derived under assumptions 1 to 3.

Notations. We recall the LQR notation used in this section. We consider LQR systems of the form (1), whose dynamics is characterized by matrices $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{n \times d}$ and cost matrix characterized by $C = \begin{pmatrix} Q & 0 \\ 0 & R \end{pmatrix}$. We summarize these parameters as $\theta^\top = (A \ B) \in \mathbb{R}^{n \times (n+d)}$.

Given an arbitrary control matrix $K \in \mathbb{R}^{d \times n}$, we denote by $A^c(\theta, K) = A + BK \in \mathbb{R}^{n \times n}$ the associated closed-loop matrix. We often use $L^\top = (I \ K^\top) \in \mathbb{R}^{n \times (n+d)}$, so that the closed-loop matrix can be written as $A^c(\theta, K) = \theta^\top L$.

Whenever, $A^c(\theta, K)$ is stable, we denote as $P(\theta, K)$ the (unique) p.s.d. matrix satisfying the Lyapunov equation (see (33))

$$P(\theta, K) = A^c(\theta, K)^\top P(\theta, K) A^c(\theta, K) + Q + K^\top R K. \quad (35)$$

Notice that $\text{Tr}(P(\theta, K))$ represents the average cost of the LQ problem parametrized with θ under control K . In a similar fashion, we denote as $\Sigma(\theta, K)$ the steady-state covariance satisfying (see (34))

$$\Sigma(\theta, K) = A^c(\theta, K) \Sigma(\theta, K) A^c(\theta, K)^\top + I. \quad (36)$$

Furthermore, whenever θ is stabilizable, we denote by $K(\theta)$ the associated optimal control and $L(\theta)^\top = (I \ K(\theta)^\top)$. When θ is controlled by $K(\theta)$ we often use the shorthand notation $A^c(\theta) = A^c(\theta, K(\theta))$, $\Sigma(\theta) = \Sigma(\theta, K(\theta))$ and $P(\theta) = P(\theta, K(\theta))$.

The parameters of the true system are summarized in $\theta_*^\top = (A_* \ B_*) \in \mathbb{R}^{n \times (n+d)}$ and we denote as K_* , L_* , P_* and Σ_* the optimal quantities $K(\theta_*)$, $L(\theta_*)$, $P(\theta_*, K_*)$ and $\Sigma(\theta_*, K_*)$ respectively.

OFU-LQ++. At each time step $t \geq 1$, OFU-LQ++ maintains an estimates $\hat{\theta}_t$ of θ_* together with a confidence set \mathcal{C}_t as given Eqs. (5) and (9). OFU-LQ++ proceeds in episodes, triggered when $\det(V_t) \geq 2 \det(V_{t_k})$, and at the beginning of each episode k , OFU-LQ++ selects an *optimistic* parameters $\theta_{t_k} \in \mathcal{C}_{t_k}$ and computes the optimal controller $K(\theta_{t_k})$ associated with θ_{t_k} . Then, the *optimistic* controller K_{t_k} is used until the end of the episode.

The optimistic nature of θ_{t_k} and $K(\theta_{t_k})$ ensures that for all $k \geq 1$, $\text{Tr}(P(\theta_{t_k}, K(\theta_{t_k}))) \leq \text{Tr}(P_*) \leq D$ on the high-probability event $\theta_* \in \mathcal{C}_{t_k}$.

D.1. Sequential stability of OFU-LQ++

The proof of the sequential stability of OFU-LQ++ is derived on the high-probability event E_T and is conducted in two steps.

Step 1) We show that OFU-LQ++ outputs a sequence of stable controller $\{K(\theta_{t_k})\}_{k \geq 1}$ (i.e., such that $A^c(\theta_*, K(\theta_{t_k}))$ is UES) as long as $\|\theta_{t_k} - \theta_*\| \leq 2\epsilon_0$ for ϵ_0 satisfying

$$\kappa \eta(\epsilon_0) \leq 1; \quad \eta(\epsilon_0) := 4\kappa \epsilon_0 (1 + \epsilon_0) < 1,$$

where $\kappa = D/\lambda_{\min}(C)$.

Step 2) We prove Lem. 1 showing that $\|\theta_{t_k} - \theta_*\| \leq 2\epsilon_0$ do hold for a suitable initialization phase and choice of regularization parameter λ by induction. This allows us to construct recursively a high-probability event over which the state process remains bounded.

D.1.1. STEP 1

The following proposition summarizes how optimism translates into stability properties.

Proposition 13. *Assume that $\theta_* \in \mathcal{C}_t$ for all $t \geq 1$. Let $\{\theta_{t_k}\}_{k \geq 1}$ and $\{K(\theta_{t_k})\}_{k \geq 1}$ be the sequence of optimistic parameters/controllers generated by OFU-LQ++ over episodes $k \geq 1$. Then, under assumptions 2 and 3, for all $k \geq 1$,*

1. θ_{t_k} is a stabilizable system and $A^c(\theta_{t_k}, K(\theta_{t_k}))$ is stable,
2. $\|L(\theta_{t_k})^\top\|_{\Sigma(\theta_{t_k}, K(\theta_{t_k}))} \leq \sqrt{\kappa}$ and $\|\Sigma(\theta_{t_k}, K(\theta_{t_k}))^{1/2}\|_F \leq \sqrt{\kappa}$, where $\kappa = D/\lambda_{\min}(C)$ and

$$\Sigma(\theta_{t_k}, K(\theta_{t_k})) = A^c(\theta_{t_k}, K(\theta_{t_k})) \Sigma(\theta_{t_k}, K(\theta_{t_k})) A^c(\theta_{t_k}, K(\theta_{t_k}))^\top + I.$$

$$3. \|L(\theta_{t_k})\|_F \leq \sqrt{\kappa} \text{ and } \|A^c(\theta_{t_k}, K(\theta_{t_k}))^\top\|_{\Sigma(\theta_{t_k}, K(\theta_{t_k}))} \leq \sqrt{\kappa}.$$

Proof.

1. By construction of θ_{t_k} and $K(\theta_{t_k})$, optimism guarantees stability of the closed-loop matrix $A^c(\theta_{t_k}, K(\theta_{t_k}))$. Indeed, if $A^c(\theta_{t_k}, K(\theta_{t_k}))$ was unstable, we would have $\text{Tr}(P(\theta_{t_k}, K(\theta_{t_k}))) = +\infty$, which contradicts the optimistic property $\text{Tr}(P(\theta_{t_k}, K(\theta_{t_k}))) \leq \text{Tr}(P_*) \leq D < +\infty$. As a result, $A^c(\theta_{t_k}, K(\theta_{t_k}))$ is stable and θ_{t_k} is a stabilizable pair.
2. By definition, we have

$$\begin{aligned} P(\theta_{t_k}, K(\theta_{t_k})) &= A^c(\theta_{t_k}, K(\theta_{t_k}))^\top P(\theta_{t_k}, K(\theta_{t_k})) A^c(\theta_{t_k}, K(\theta_{t_k})) + L(\theta_{t_k})^\top CL(\theta_{t_k}), \\ \Sigma(\theta_{t_k}, K(\theta_{t_k})) &= A^c(\theta_{t_k}, K(\theta_{t_k})) \Sigma(\theta_{t_k}, K(\theta_{t_k})) A^c(\theta_{t_k}, K(\theta_{t_k}))^\top + I. \end{aligned}$$

Then, applying Prop. 11 with $M = L(\theta_{t_k})^\top CL(\theta_{t_k})$ leads to

$$D \geq \text{Tr}(P(\theta_{t_k}, K(\theta_{t_k}))) = \text{Tr}(\Sigma(\theta_{t_k}, K(\theta_{t_k}))M).$$

Noticing that $M \succcurlyeq \lambda_{\min}(C)L(\theta_{t_k})^\top L(\theta_{t_k})$ we obtain that

$$D \geq \text{Tr}(P(\theta_{t_k}, K(\theta_{t_k}))) \geq \text{Tr}(\Sigma(\theta_{t_k}, K(\theta_{t_k}))L(\theta_{t_k})^\top L(\theta_{t_k})) \lambda_{\min}(C) = \lambda_{\min}(C) \|L(\theta_{t_k})^\top\|_{\Sigma(\theta_{t_k}, K(\theta_{t_k}))}^2.$$

Similarly, from $M \succcurlyeq Q + K(\theta_{t_k})^\top RK(\theta_{t_k}) \succcurlyeq Q \succcurlyeq \lambda_{\min}(C)I$ we obtain

$$D \geq \text{Tr}(P(\theta_{t_k}, K(\theta_{t_k}))) \geq \text{Tr}(\Sigma(\theta_{t_k}, K(\theta_{t_k}))) \lambda_{\min}(C) = \lambda_{\min}(C) \|\Sigma(\theta_{t_k}, K(\theta_{t_k}))\|_F^2.$$

We conclude using that $\kappa = D/\lambda_{\min}(C)$.

3. The last assertions come directly from the previous one, noticing that $\Sigma(\theta_{t_k}, K(\theta_{t_k})) \succcurlyeq I$ and $A^c(\theta_{t_k}, K(\theta_{t_k})) \Sigma(\theta_{t_k}, K(\theta_{t_k})) A^c(\theta_{t_k}, K(\theta_{t_k}))^\top \preccurlyeq \Sigma(\theta_{t_k}, K(\theta_{t_k}))$.

□

Equipped with those properties, we now show that for any $k \geq 1$, the optimistic controller $K(\theta_{t_k})$ stabilizes the true system parameterized by θ_* , as long as θ_{t_k} is *sufficiently close* to θ_* and characterize the stability margin in term of steady-state covariance and cost.

Proposition 14. *Let $k \geq 0$ and assume that $\theta_* \in \mathcal{C}_{t_k}$. Let θ_{t_k} and $K(\theta_{t_k})$ be an optimistic parameter/control pair. Then, if $\|\theta_{t_k} - \theta_*\|_F \leq 2\epsilon_0$ where ϵ_0 is such that $\eta(\epsilon_0) := 4\kappa\epsilon_0(1 + \epsilon_0) < 1$, under assumptions 2 and 3,*

1. $A^c(\theta_*, K(\theta_{t_k}))$ is stable,
2. $\Sigma(\theta_*, K(\theta_{t_k})) \preccurlyeq \frac{1}{1-\eta(\epsilon_0)} \Sigma(\theta_{t_k}, K(\theta_{t_k}))$,
3. $\|P(\theta_*, K(\theta_{t_k}))\|_2 \preccurlyeq \|P(\theta_*, K(\theta_*))\|_2 + \frac{D\eta(\epsilon_0)}{1-\eta(\epsilon_0)}$,
4. $\|P(\theta_{t_k}, K(\theta_{t_k}))\|_2 \leq \|P(\theta_*, K(\theta_*))\|_2(1 + \kappa\eta(\epsilon_0)/2) + \frac{D\eta(\epsilon_0)}{1-\eta(\epsilon_0)}(1 + \kappa\eta(\epsilon_0)/2)$.

Proof. 1. The first two assertions directly follows from propositions 12 and 13. From

$$\left\| (A^c(\theta_*, K(\theta_{t_k})) - A^c(\theta_{t_k}, K(\theta_{t_k})))^\top \right\|_{\Sigma(\theta_{t_k}, K(\theta_{t_k}))} \leq \|\theta_* - \theta_{t_k}\|_F \|L(\theta_{t_k})^\top\|_{\Sigma(\theta_{t_k}, K(\theta_{t_k}))} \leq 2\epsilon_0 \sqrt{\kappa}$$

and

$$\|A^c(\theta_{t_k}, K(\theta_{t_k}))^\top\|_{\Sigma(\theta_{t_k}, K(\theta_{t_k}))} \leq \sqrt{\kappa},$$

we can apply Prop. 12 with $A_0 = A^c(\theta_{t_k}, K(\theta_{t_k}))$, $A_1 = A^c(\theta_*, K(\theta_{t_k}))$, $\eta = 2\epsilon_0$ and $c_0 = \sqrt{\kappa}$, which leads to the desired result.

2. Let $\Delta = P(\theta_*, K(\theta_{t_k})) - P(\theta_*, K(\theta_*))$ and $\eta(\epsilon_0) = 4\kappa\epsilon_0(1 + \epsilon_0)$. We have

$$\begin{aligned} \|\Delta\|_2 &= \lambda_{\max}(\Delta) \leq \text{Tr}\left(P(\theta_*, K(\theta_{t_k}))\right) - \text{Tr}\left(P(\theta_*, K(\theta_*))\right) \\ &= \text{Tr}\left(P(\theta_*, K(\theta_{t_k}))\right) - \text{Tr}\left(P(\theta_{t_k}, K(\theta_{t_k}))\right) + \text{Tr}\left(P(\theta_{t_k}, K(\theta_{t_k}))\right) - \text{Tr}\left(P(\theta_*, K(\theta_*))\right) \\ &\leq \text{Tr}\left(P(\theta_*, K(\theta_{t_k}))\right) - \text{Tr}\left(P(\theta_{t_k}, K(\theta_{t_k}))\right), \end{aligned}$$

where we used that by optimism, $\text{Tr}\left(P(\theta_{t_k}, K(\theta_{t_k}))\right) - \text{Tr}\left(P(\theta_*, K(\theta_*))\right) \leq 0$. Using that

$$\begin{aligned} \text{Tr}\left(P(\theta_*, K(\theta_{t_k}))\right) &= \text{Tr}\left(L(\theta_{t_k})^\top CL(\theta_{t_k})\Sigma(\theta_*, K(\theta_{t_k}))\right) \\ \text{Tr}\left(P(\theta_{t_k}, K(\theta_{t_k}))\right) &= \text{Tr}\left(L(\theta_{t_k})^\top CL(\theta_{t_k})\Sigma(\theta_{t_k}, K(\theta_{t_k}))\right) \end{aligned}$$

we obtain

$$\lambda_{\max}(\Delta) \leq \text{Tr}\left(L(\theta_{t_k})^\top CL(\theta_{t_k})(\Sigma(\theta_*, K(\theta_{t_k})) - \Sigma(\theta_{t_k}, K(\theta_{t_k})))\right).$$

Finally, from $\Sigma(\theta_*, K(\theta_{t_k})) - \Sigma(\theta_{t_k}, K(\theta_{t_k})) \preceq \frac{\eta(\epsilon_0)}{1 - \eta(\epsilon_0)}\Sigma(\theta_{t_k}, K(\theta_{t_k}))$ we have

$$\lambda_{\max}(\Delta) \leq \frac{\eta(\epsilon_0)}{1 - \eta(\epsilon_0)} \text{Tr}\left(L(\theta_{t_k})^\top CL(\theta_{t_k})\Sigma(\theta_{t_k}, K(\theta_{t_k}))\right) = \frac{\eta(\epsilon_0)}{1 - \eta(\epsilon_0)} \text{Tr}\left(P(\theta_{t_k}, K(\theta_{t_k}))\right) \leq \frac{D\eta(\epsilon_0)}{1 - \eta(\epsilon_0)}.$$

3. Let $\Delta = P(\theta_{t_k}, K(\theta_{t_k})) - P(\theta_*, K(\theta_{t_k}))$ and $\eta(\epsilon_0) = 4\kappa\epsilon_0(1 + \epsilon_0)$. We have

$$\begin{aligned} \Delta &= A^c(\theta_{t_k}, K_{t_k})^\top \Delta A^c(\theta_{t_k}, K_{t_k}) + A^c(\theta_{t_k}, K_{t_k})^\top P(\theta_*, K(\theta_{t_k})) A^c(\theta_{t_k}, K_{t_k}) - A^c(\theta_*, K_{t_k})^\top P(\theta_*, K(\theta_{t_k})) A^c(\theta_*, K_{t_k}) \\ &\leq A^c(\theta_{t_k}, K_{t_k})^\top \Delta A^c(\theta_{t_k}, K_{t_k}) + \epsilon_0^2 \|L_t\|_2^2 \|P(\theta_*, K(\theta_{t_k}))\|_2 + 2\epsilon_0 \|L_t\|_2 \|P(\theta_*, K(\theta_{t_k}))\|_2 \\ &\leq A^c(\theta_{t_k}, K_{t_k})^\top \Delta A^c(\theta_{t_k}, K_{t_k}) + (\epsilon_0^2 \kappa + 2\epsilon_0 \sqrt{\kappa}) \|P(\theta_*, K(\theta_{t_k}))\|_2, \\ &\leq A^c(\theta_{t_k}, K_{t_k})^\top \Delta A^c(\theta_{t_k}, K_{t_k}) + \frac{1}{2} \eta(\epsilon_0) \|P(\theta_*, K(\theta_{t_k}))\|_2. \end{aligned}$$

which implies that $\|\Delta\|_2 \leq \kappa\eta(\epsilon_0)/2 \|P(\theta_*, K(\theta_{t_k}))\|_2$. Combined with the previous point, we obtain,

$$\|P(\theta_{t_k}, K(\theta_{t_k}))\|_2 \leq \|P(\theta_*, K(\theta_{t_k}))\|_2 (1 + \kappa\eta(\epsilon_0)/2) \leq \|P(\theta_*, K(\theta_*))\|_2 (1 + \kappa\eta(\epsilon_0)/2) + \frac{D\eta(\epsilon_0)}{1 - \eta(\epsilon_0)} (1 + \kappa\eta(\epsilon_0)/2)$$

□

We conclude **Step 1**) showing that whenever θ_{t_k} is *sufficiently close* to θ_* for all $t \geq 1$, the sequence of optimistic controller $\{K(\theta_{t_k})\}_{k \geq 1}$ induces a state process $\{x_t\}_{t \geq 1}$ for each episode k given by

$$x_{t+1} = A^c(\theta_*, K(\theta_{t_k}))x_t + \epsilon_t, \quad x_0 = \xi, \quad \forall t \in [t_k, t_{k+1} - 1], \quad (37)$$

that is *uniformly exponentially stable*.

Proposition 15. Assume that $\theta_* \in \mathcal{C}_t$ for all $t \geq 1$. Let $\{\theta_{t_k}\}_{k \geq 1}$ and $\{K(\theta_{t_k})\}_{k \geq 1}$ be the sequence of optimistic parameters/controls pair. Then, if for all $k \geq 0$, $\|\theta_{t_k} - \theta_*\| \leq 2\epsilon_0$ where ϵ_0 is such that

$$\kappa\eta(\epsilon_0) \leq 1; \quad \eta(\epsilon_0) := 4\kappa\epsilon_0(1 + \epsilon_0) \leq 1/2, \quad (38)$$

under assumptions 1 to 3,

1. for all k , the state process in (37) is uniformly exponentially stable,
2. for all k , for any $0 < \delta < 1$, with probability at least $1 - \delta$,

$$\begin{aligned} \|x_t\| &\leq 4\sigma\sqrt{\kappa \log(1/\delta)} + \sqrt{\kappa}(1 - 1/\kappa)^t \|\xi\| \\ \text{for all } t \in [t_k, t_{k+1} - 1], \quad \|x_t\| &\leq 4\sigma\sqrt{\kappa \log(1/\delta)} + 2\sqrt{\|P(\theta_*, K(\theta_*))\|_2 / \lambda_{\min}(C)} \|\xi\|. \end{aligned} \quad (39)$$

Proof.

1. By propositions 13 and 14, it exists $\Sigma_k = \Sigma(\theta_*, K(\theta_{t_k}))$ such that

$$\Sigma_k = A^c(\theta_*, K(\theta_{t_k}))\Sigma_k A^c(\theta_*, K(\theta_{t_k}))^\top + I, \quad \text{Tr}(\Sigma_k) \leq \kappa.$$

As a result, Prop. 6 ensures that the process in (37) is *uniformly exponentially stable*.

2. Noticing that $t_{k+1} = \inf\{t \geq t_k \text{ s.t. } \det(V_t) > 2 \det(V_{t_k})\}$ is a stopping time, we obtain the first bound in the second assertion directly by applying propositions 6 and 7. To obtain the second bound, we use Prop. 14

$$\begin{aligned} \|A^c(\theta_*, K(\theta_{t_k}))^t \xi\| &\leq \|A^c(\theta_*, K(\theta_{t_k}))^t\|_2 \|\xi\| \leq \left\| \sqrt{P(\theta_*, K(\theta_{t_k}))} \right\|_2 / \lambda_{\min}(C) \|\xi\| \\ &\leq \sqrt{\|P(\theta_*, K(\theta_*))\|_2 / \lambda_{\min}(C) + 1} \|\xi\| \\ &\leq 2 \sqrt{\|P(\theta_*, K(\theta_*))\|_2 / \lambda_{\min}(C)} \|\xi\|. \end{aligned}$$

□

D.1.2. STEP 2. PROOF OF LEM. 1

Let $X := 20\sigma\sqrt{\kappa \log(4T/\delta)}\sqrt{\|P(\theta_*, K(\theta_*))\|_2 / \lambda_{\min}(C)}$, we are now ready to construct recursively a sequence of high-probability events

$$F_t = \left\{ \forall s \leq t, \|x_s\| \leq X \right\}$$

over which the state process remains bounded. We work on $\{E_t\}_{t \leq T}$ and we proceed by episode on the increasing sequence of probability events $\{E_{t_k} \cap F_{t_k}\}_{k \geq 1}$.

Proof of Lem. 1. We assume that OFU-LQ++ is initialized with a stabilizing set $\Theta_0 = \{\theta : \|\theta - \theta_0\| \leq \epsilon_0\}$ such that ϵ_0 satisfies

$$\kappa\eta(\epsilon'_0) \leq 1; \quad \eta(\epsilon'_0) := 4\kappa\epsilon'_0(1 + \epsilon'_0) \leq 1/2; \quad \epsilon'_0 = 2\epsilon_0.$$

Further, we assume that the RLS estimation in (6) is instantiated with

$$\lambda = \frac{2n\sigma^2}{\epsilon_0^2} \left(\log(4n/\delta) + (n+d) \log(1 + \kappa X^2 T) \right).$$

Setting $\epsilon_0 \leq \frac{1}{32\kappa^2}$ ensures that $\eta(\epsilon'_0) \leq 1/2$, $\kappa\eta(\epsilon'_0) \leq 1$ and that $\lambda \geq 4^3\sigma^2\kappa^4 \log(4T/\delta) / \log(2)$.

We now proceed by induction, for sake of simplicity, we assume that $x_0 = 0$. We show that w.h.p.

- 1) $\|x_t\| \leq X := 20\sigma\sqrt{\kappa \log(4T/\delta)}\sqrt{\|P(\theta_*, K(\theta_*))\|_2 / \lambda_{\min}(C)}$ for all $t \in [t_k, t_{k+1}]$

- 2) $\|x_{t_{k+1}}\| \leq 8\sigma\sqrt{\kappa \log(4T/\delta)}$.

- At $k = 0$, $\theta_{t_0} \in \mathcal{C}_0$ implies that

$$\|\theta_{t_0} - \theta_*\| \leq 2\beta_{t_0} / \sqrt{\lambda_{\min}(V_{t_0})} = 2(\epsilon_0 + \sigma\sqrt{2n \log(4n/\delta) / \sqrt{\lambda}}) \leq 4\epsilon_0 := 2\epsilon'_0.$$

As a result, Prop. 15 guarantees that with probability at least $1 - \delta/4T$,

$$\text{for all } s \leq t_1, \quad \|x_s\| \leq 4\sigma\sqrt{\kappa \log(4T/\delta)} + 2\sqrt{\|P(\theta_*, K(\theta_*))\|_2 / \lambda_{\min}(C)} \|x_0\| \leq X \quad \text{and} \quad \|x_{t_1}\| \leq 8\sigma\sqrt{\kappa \log(4T/\delta)}.$$

Hence, $\mathbb{P}(F_{t_1} | E_{t_0}) \geq 1 - \delta/4T$.

- At $k + 1$, assume that $E_{t_k} \cap F_{t_k}$ holds. From $\|x_s\| \leq X$ for all $s \leq t_{k+1}$, we obtain that

$$\begin{aligned}
 \beta_{t_{k+1}} &= \sigma \sqrt{2n \log \left(\frac{\det(V_{t_{k+1}})^{1/2} 4n}{\det(\lambda I)^{1/2} \delta} \right)} + \lambda^{1/2} \|\theta_0 - \theta_*\|_F \\
 &\leq \sigma \sqrt{2n \log(4n/\delta) + 2n(n+d) \log \left(1 + \sup_{s \leq t_{k+1}-1} \|z_s\|^2 T \right)} + \lambda^{1/2} \epsilon_0 \\
 &\leq \sigma \sqrt{2n \log(4n/\delta) + 2n(n+d) \log \left(1 + \kappa X^2 T \right)} + \lambda^{1/2} \epsilon_0 \\
 &\leq 2\epsilon_0 \sqrt{\lambda},
 \end{aligned}$$

where we used that $z_s = L(\theta_s)x_s$ and $\|L(\theta_s)\|_2 \leq \sqrt{\kappa}$ by Prop. 13, and the definition of λ . As a result, $\|\theta_{t_{k+1}} - \theta_*\| \leq 2\epsilon'_0$ and we can apply Prop. 15. Using that $\|x_{t_k}\| \leq 8\sigma\sqrt{\kappa \log(4T/\delta)}$ we obtain with probability at least $1 - \delta/4T$,

$$\begin{aligned}
 \text{for all } t_k \leq s \leq t_{k+1}, \quad \|x_s\| &\leq 4\sigma\sqrt{\kappa \log(4T/\delta)} + 2\sqrt{\|P(\theta_*, K(\theta_*))\|_2 / \lambda_{\min}(C)} \|x_{t_k}\| \\
 &\leq 4\sigma\sqrt{\kappa \log(4T/\delta)} (1 + 4\sqrt{\|P(\theta_*, K(\theta_*))\|_2 / \lambda_{\min}(C)}) \\
 &\leq X.
 \end{aligned}$$

Hence, $\mathbb{P}(F_{t_{k+1}} | E_{t_k}, F_{t_k}) \geq 1 - \delta/4T$. Further, on this event,

$$\|x_{t_{k+1}}\| \leq 4\sigma\sqrt{\kappa \log(4T/\delta)} (1 + \sqrt{\kappa} (1 - 1/\kappa)^{t_{k+1} - t_k + 1}) \leq 8\sigma\sqrt{\kappa \log(4T/\delta)}$$

which concludes the induction. To obtain the last inequality, we used that $t_{k+1} - t_k + 1 \geq \kappa \log(\kappa)/2$ since otherwise it enters in contradiction with the update rule: suppose that $t_{k+1} - t_k + 1 \leq \kappa \log(\kappa)/2$,

$$\begin{aligned}
 \det(V_{t_{k+1}}) &\leq \det(V_{t_k}) \prod_{t=t_k}^{t_{k+1}-1} (1 + \|z_t\|_{V_t}^2) \leq \det(V_{t_k}) \prod_{t=t_k}^{t_{k+1}-1} (1 + \|x_t\|^2 \kappa / \lambda) \\
 &\leq \det(V_{t_k}) \exp \left(\frac{\kappa}{\lambda} \sum_{t=t_k}^{t_{k+1}-1} \|x_t\|^2 \right) \leq \det(V_{t_k}) \exp \left(32\sigma^2 \kappa^2 \log(4T/\delta) ((t_{k+1} - t_k + 1) + \kappa^2) / \lambda \right) \\
 &\leq \det(V_{t_k}) \exp \left(4^3 \sigma^2 \kappa^4 \log(4T/\delta) / \lambda \right) \leq 2 \det(V_{t_k}),
 \end{aligned}$$

where we used that $\lambda \geq 4^3 \sigma^2 \kappa^4 \log(4T/\delta) / \log(2)$.

Overall, we showed that the sequence of event F_t is such that $\mathbb{P}(F_{t_{k+1}} | F_{t_k}, E_{t_k}) \geq 1 - \delta/4T$ for all $k \geq 0$. As a result, we have

$$\mathbb{P}(F_T \cap E_T) \geq \mathbb{P}(E_T \cap \bigcap_{k \geq 0} F_{t_k} | F_{t_{k-1}}, E_{t_{k-1}}) \geq 1 - \mathbb{P}(E_T^c) - \sum_{k \geq 0} \mathbb{P}((F_{t_k} | F_{t_{k-1}}, E_{t_{k-1}})^c) \geq 1 - \delta/2.$$

□

D.2. Regret bound for OFU-LQ++

The proof of Lem. 2 follows from carefully using Lem. 1 in the analysis of Abbasi-Yadkori & Szepesvári (2011). We provide it for sake of completeness. In line with Abbasi-Yadkori & Szepesvári (2011), we perform a regret decomposition

on the sequence of events $\{E_t \cap F_t\}_{t \geq 0}$ and obtain

$$\begin{aligned}
 R(T) &= \sum_{t=0}^T x_t^\top Q x_t + u_t^\top R u_t - J_* = \sum_{t=0}^T J(\theta_t) - J_* \\
 &\quad + \sum_{t=0}^T x_t^\top P_t x_t - \mathbb{E}(x_{t+1}^\top P_{t+1} x_{t+1} | \mathcal{F}_t) \\
 &\quad + \sum_{t=0}^T \mathbb{E}(x_{t+1}^\top (P_{t+1} - P_t) x_{t+1} | \mathcal{F}_t) \\
 &\quad + \sum_{t=0}^T (\theta_t^\top z_t)^\top P_t (\theta_t^\top z_t) - (\theta_*^\top z_t)^\top P_t (\theta_*^\top z_t).
 \end{aligned} \tag{40}$$

where $\theta_t = \theta_{t_k}$ and $P_t = P(\theta_{t_k})$ for all $[t_k, t_{k+1} - 1]$. θ_{t_k} is the optimistic parameter chosen by OFU-LQ++ at the beginning of episode k . Each regret terms are then bounded separately as in [Abbasi-Yadkori & Szepesvári \(2011\)](#) but for the last term.

Bounding $R^{opt}(T) = \sum_{t=0}^T J(\theta_t) - J_*$. Since θ_{t_k} is optimistic, on E_t , $J(\theta_t) \leq J_*$. As a result, $R^{opt}(T) \leq 0$. Further, by Prop. 14 and the tuning of ϵ_0 , $\|P_t\|_2 = O(\|P(\theta_*)\|_2)$.

Bounding $R^{mart}(T) = \sum_{t=0}^T x_t^\top P_t x_t - \mathbb{E}(x_{t+1}^\top P_{t+1} x_{t+1} | \mathcal{F}_t)$. On $E_t \cap F_t$, Lem. 1 guarantees that $\|x_t\| \leq X$ where $X = 20\sigma\sqrt{\kappa \log(4T/\delta)}\sqrt{\|P(\theta_*, K(\theta_*))\|_2 / \lambda_{\min}(C)}$. Thus, $R^{mart}(T)$ is a bounded martingale sequence and applying Azuma's inequality leads to

$$R^{mart}(T) = \tilde{O}(X^2 \|P(\theta_*)\|_2 \sqrt{T}) = \tilde{O}(\kappa \|P(\theta_*)\|_2^2 \sqrt{T}).$$

Bounding $R^{lazy}(T) = \sum_{t=0}^T \mathbb{E}(x_{t+1}^\top (P_{t+1} - P_t) x_{t+1} | \mathcal{F}_t)$. The lazy update scheme used to re-evaluate the controllers $K_{t,s}$, ensures that the number of updates scales logarithmically. Hence, $R^{lazy}(T) = \tilde{O}(X^2 \|P(\theta_*)\|_2 (n + d))$.

Bounding $R^{pred}(T) = \sum_{t=0}^T (\theta_t^\top z_t)^\top P_t (\theta_t^\top z_t) - (\theta_*^\top z_t)^\top P_t (\theta_*^\top z_t)$. From the triangular inequality, we have:

$$R^{pred}(T) = \sum_{t=0}^T \|\theta_t^\top z_t\|_{P_t}^2 - \|\theta_*^\top z_t\|_{P_t}^2 = \sum_{t=0}^T (\|\theta_t^\top z_t\|_{P_t} - \|\theta_*^\top z_t\|_{P_t}) (\|\theta_t^\top z_t\|_{P_t} + \|\theta_*^\top z_t\|_{P_t}).$$

Further, $\|\theta_t^\top z_t\|_{P_t} = \|A^c(\theta_t, K(\theta_t))x_t\|_{P_t} \leq \|x_t\|_{P_t} \leq \|P(\theta_*)\|_2^{1/2} X$. Similarly, on $E_t \cap F_t$,

$$\begin{aligned}
 \|\theta_*^\top z_t\|_{P_t} &\leq \|\theta_t^\top z_t\|_{P_t} + \|A^c(\theta_*, K(\theta_t))x_t - A^c(\theta_t, K(\theta_t))x_t\|_{P_t} \\
 &\leq \|P(\theta_*)\|_2^{1/2} X + 4\|P(\theta_*)\|_2^{1/2} X \sqrt{\kappa} \epsilon_0 \\
 &\leq 2X \|P(\theta_*)\|_2^{1/2}.
 \end{aligned}$$

where we used $\|\theta_t - \theta_*\| \leq 2\epsilon_0$ from line 1 to 2, and the condition on ϵ_0 from line 2 to 3. Thus,

$$R^{pred}(T) \leq 3X \|P(\theta_*)\|_2 \sum_{t=0}^T \|(\theta_t - \theta_*)^\top z_t\| \leq 6X \|P(\theta_*)\|_2 \beta_T \sum_{t=1}^T \|z_t\|_{V_t^{-1}}.$$

Finally, using (7), we have that $\sum_{t=1}^T \|z_t\|_{V_t^{-1}} = \tilde{O}(\sqrt{(n+d)T})$. Noticing that $\kappa X^2 / \lambda = O(1)$, $X = \tilde{O}(\sqrt{\kappa \|P(\theta_*)\|_2})$ and $\beta_T = \tilde{O}(\sqrt{n(n+d)})$, we obtain $R^{pred}(T) = \tilde{O}(\sqrt{\kappa} \|P(\theta_*)\|_2 (n+d) \sqrt{n} \sqrt{T})$.

E. Optimism, Sequential Stability and Regret of LAGLQ

This section is dedicated to the proofs of LAGLQ and concerns the *Extended Optimistic LQR with Relaxed Constraints* introduced in Sect. 4. We first prove Lem. 3 showing that LAGLQ does use an *optimistic* controller (App. E.1). We then prove Lem. 4 showing that it preserves the sequential stability properties of OFU-LQ++ (App. E.2). We finally prove Thm. 1 ensuring that LAGLQ suffers a similar regret compared to OFU-LQ++ (App. E.3). All the results in the section are derived under assumptions 1 to 3.

Notations. We first recall the notation used in the definition of the extended LQR. The original LQR is associated with the dynamics and cost matrix parametrized respectively by $\theta = (A \ B)$ and $C = \begin{pmatrix} Q & 0 \\ 0 & R \end{pmatrix}$. The extended LQR is obtained by adding an additional perturbation control w , which induces the extended LQR associated with the dynamics and cost matrix parametrized respectively by $\tilde{\theta}^\top = (A \ \tilde{B}) \in \mathbb{R}^{n \times (2n+d)}$, where $\tilde{B} = (B \ I) \in \mathbb{R}^{n \times (n+d)}$ and $C_{\dagger} = \begin{pmatrix} C & 0 \\ 0 & 0 \end{pmatrix}$. The extended control matrix $\tilde{K} \in \mathbb{R}^{(n+d) \times n}$ can be decomposed as $\tilde{K}^\top = (K^{u^\top} \ K^{w^\top})$ with the ‘‘standard’’ control u and the perturbation control w . Following these definitions, it is possible to redefine the notation used for ‘‘standard’’ LQRs as

$$\begin{aligned} A^c(\tilde{\theta}, \tilde{K}) &= A + \tilde{B}\tilde{K} = A + BK^u + K^w \\ P(\tilde{\theta}, \tilde{K}) &= A^c(\tilde{\theta}, \tilde{K})^\top P(\tilde{\theta}, \tilde{K}) A^c(\tilde{\theta}, \tilde{K}) + Q + K^{u^\top} R K^u, \\ \Sigma(\tilde{\theta}, \tilde{K}) &= A^c(\tilde{\theta}, \tilde{K}) \Sigma(\tilde{\theta}, \tilde{K}) A^c(\tilde{\theta}, \tilde{K})^\top + I. \end{aligned} \quad (41)$$

From the last equation we also obtain that for any control policy $\tilde{\pi}(x) = \tilde{K}x$, the corresponding average expected reward is $\mathcal{J}_{\tilde{\pi}}(\tilde{\theta}) = \text{Tr}(P(\tilde{\theta}, \tilde{K}))$. Finally, notice that the extended dynamics $\tilde{\theta}$ is controllable even if the original one θ is not¹⁰.

LAGLQ. LAGLQ follows the same procedure as OFU-LQ++ and operates through episodes. A new episode starts as soon as $\det(V_t) \geq 2 \det(V_{t_k})$, where t_k is the initial starting time of episode k . At each time step $t \geq 1$, as for OFU-LQ++, LAGLQ maintains an estimates $\hat{\theta}_t$ of θ_* together with a confidence set \mathcal{C}_t as given Eqs. (5) and (9). At the beginning of each episode, LAGLQ computes an *optimistic* policy solving the Extended LQR with Relaxed Constraint in (16) using the DS-OFU procedure described in Sect. 5.2. Further, Lem. 5 ensures that the *extended optimistic* policy in linear and we denote by \tilde{K}_t the associated extended linear controller. Finally, the effective controller K_t^u is obtained by extracting the control part of the extended controller \tilde{K}_t and is used until the end of the episode.

E.1. Optimism (Proof of Lemma 3)

In order to prove the lemma, it is sufficient to display a control policy $\bar{\pi}$ that satisfy the constraint, and such that $\mathcal{J}_{\bar{\pi}}(\hat{\theta}, \beta, V) \leq J_*$.

Step 1 (optimism). Consider a controller $\bar{\pi} = (\bar{\pi}^u, \bar{\pi}^w)$ defined as

$$\bar{\pi}^u(x) = K_* x \quad \text{and} \quad \bar{\pi}^w(x) = (A_* + B_* K_*)x - (\hat{A} + \hat{B} K_*)x = (A_* - \hat{A})x + (B_* - \hat{B})\bar{\pi}^u(x).$$

Then the dynamics of the corresponding extended LQR reduces to

$$\begin{aligned} x_{s+1} &= \hat{A}x_s + \hat{B}u_s + w_s + \epsilon_{s+1} \\ &= \hat{A}x_s + \hat{B}\bar{\pi}^u(x_s) + (A_* - \hat{A})x_s + (B_* - \hat{B})\bar{\pi}^u(x_s) \\ &= (A_* + B_* K_*)x_s + \epsilon_{s+1}, \end{aligned} \quad (42)$$

which coincides with the dynamics induced by the optimal controller K_* in the original LQR. Notice that Asm. 2 guarantees $A_* + B_* K_*$ to be stable, and hence $\bar{\pi}$ is a stabilizing policy for the extended system parametrized with $\hat{\theta}$. As a result, the expected average cost of $\bar{\pi}$ is $J^* = \mathcal{J}_{\bar{\pi}}(\hat{\theta})$.

Step 2 (constraint $g_{\bar{\pi}}$). We now need to verify that $\bar{\pi}$ satisfies the constraints $g_{\bar{\pi}}(\hat{\theta}, \beta, V)$. By definition of $\bar{\pi}$ we have that for any $s \geq 0$,

$$\|w_s\| = \|(A_* - \hat{A})x_s + (B_* - \hat{B})u_s\| = \left\| (\theta_* - \hat{\theta})^\top \begin{pmatrix} x_s \\ u_s \end{pmatrix} \right\| \leq \beta \left\| \begin{pmatrix} x_s \\ u_s \end{pmatrix} \right\|_{V^{-1}} = \beta \|z_s\|, \quad (43)$$

¹⁰This directly comes from the fact that the extended matrix \tilde{B} is full column rank.

where the inequality is obtained from Prop. 1 under the event $\theta_* \in \mathcal{C}$. As a result, the constraint $g_{\bar{\pi}}(\hat{\theta}, \beta, V) = \lim_{S \rightarrow \infty} \frac{1}{S} \mathbb{E}(\sum_{s=0}^S \|w_s\|^2 - \beta \|z_s\|_{V^{-1}}^2) \leq 0$ is met.

E.2. Sequential Stability (Proof of Lem. 4)

The proof follows a similar path as the one of Lem. 1. Notice that now, the effective closed-loop matrix which drives the system over the episode k is given by

$$A^c(\theta_*, K_{t_k}^u) = A_* + B_* K_{t_k}^u = \theta_*^\top L_{t_k}, \quad \text{where} \quad L_{t_k} = \begin{pmatrix} I \\ K_{t_k}^u \end{pmatrix}, \quad (44)$$

while the closed-loop matrix which drives the optimistic extended LQR is given by

$$A^c(\tilde{\theta}_{t_k}, \tilde{K}_{t_k}) = \hat{A}_{t_k} + \hat{B}_{t_k} K_{t_k}^u + K_{t_k}^w = \theta_{t_k}^\top L_{t_k} = A^c(\theta_{t_k}, K_{t_k}^u), \quad \text{where} \quad \theta_{t_k}^\top = \tilde{\theta}_{t_k}^\top + K_{t_k}^w.$$

As for the proof of Lem. 1, the proof of the sequential stability of LAGLQ is derived on the high-probability event E_T and is conducted in two steps.

Step 1) we show that LAGLQ outputs a sequence of stable controller $\{K_{t_k}^u\}_{k \geq 1}$ (i.e., such that $A^c(\theta_*, K_{t_k}^u)$ is UES) as long as $\|\theta_{t_k} - \theta_*\| \leq 2\epsilon_0$ for ϵ_0 satisfying

$$\eta(\epsilon_0) := 4\kappa\epsilon_0(1 + \epsilon_0) < 1,$$

where $\kappa = D/\lambda_{\min}(C)$.

Step 2) we prove Lem. 4 showing that $\|\theta_{t_k} - \theta_*\| \leq 2\epsilon_0$ do hold for a suitable initialization phase and choice of regularization parameter λ by induction. This allows us to construct recursively a high-probability event over which the state process remains bounded.

The following proposition summarizes how optimism (of the Extended LQR with Relaxed Constraint) translates into stability properties.

Proposition 16. *Assume that $\theta_* \in \mathcal{C}_t$ for all $t \geq 1$. Let $\{\tilde{K}_{t_k}\}_{k \geq 1}$ be the sequence of extended optimistic controllers generated by LAGLQ over episodes $k \geq 1$. Let $\{K_{t_k}^u\}_{k \geq 1}$ and $\{K_{t_k}^w\}_{k \geq 1}$ be the sequence of control and perturbation controllers associated with $\{\tilde{K}_{t_k}\}_{k \geq 1}$ and let $\theta_{t_k} = \tilde{\theta}_{t_k} + K_{t_k}^w$. Then, under assumptions 2 and 3, for all $k \geq 1$,*

1. θ_{t_k} is a stabilizable pair and $A^c(\theta_{t_k}, K_{t_k}^u)$ is stable,
2. $\|L_{t_k}^\top\|_{\Sigma(\theta_{t_k}, K_{t_k}^u)} \leq \sqrt{\kappa}$ and $\|\Sigma(\theta_{t_k}, K_{t_k}^u)^{1/2}\|_F \leq \sqrt{\kappa}$, where $\kappa = D/\lambda_{\min}(C)$, $L_{t_k} = \begin{pmatrix} I \\ K_{t_k}^u \end{pmatrix}$ and

$$\Sigma(\theta_{t_k}, K_{t_k}^u) = A^c(\theta_{t_k}, K_{t_k}^u) \Sigma(\theta_{t_k}, K_{t_k}^u) A^c(\theta_{t_k}, K_{t_k}^u)^\top + I.$$
3. $\|L_{t_k}\|_F \leq \sqrt{\kappa}$ and $\|A^c(\theta_{t_k}, K_{t_k}^u)^\top\|_{\Sigma(\theta_{t_k}, K_{t_k}^u)} \leq \sqrt{\kappa}$.

Proof.

1. By construction, optimism guarantees that $\text{Tr}(P(\tilde{\theta}_{t_k}, \tilde{K}_{t_k})) \leq D < \infty$. As a result, $A^c(\tilde{\theta}_{t_k}, \tilde{K}_{t_k})$ is stable. Noticing that $A^c(\tilde{\theta}_{t_k}, \tilde{K}_{t_k}) = \hat{A}_{t_k} + \hat{B}_{t_k} K_{t_k}^u + K_{t_k}^w = A^c(\theta_{t_k}, K_{t_k}^u)$ we obtain that $A^c(\theta_{t_k}, K_{t_k}^u)$ is stable and hence that θ_{t_k} is a stabilizable pair.
2. From $A^c(\tilde{\theta}_{t_k}, \tilde{K}_{t_k}) = A^c(\theta_{t_k}, K_{t_k}^u)$ and $\begin{pmatrix} I \\ \tilde{K}_{t_k} \end{pmatrix}^\top C_\dagger \begin{pmatrix} I \\ \tilde{K}_{t_k} \end{pmatrix} = L_{t_k}^\top C L_{t_k}$, we have that

$$\begin{aligned} P(\tilde{\theta}_{t_k}, \tilde{K}_{t_k}) &= A^c(\theta_{t_k}, K_{t_k}^u)^\top P(\tilde{\theta}_{t_k}, \tilde{K}_{t_k}) A^c(\theta_{t_k}, K_{t_k}^u) + L(\theta_{t_k})^\top C L(\theta_{t_k}), \\ \Sigma(\theta_{t_k}, K_{t_k}^u) &= A^c(\theta_{t_k}, K_{t_k}^u) \Sigma(\theta_{t_k}, K_{t_k}^u) A^c(\theta_{t_k}, K_{t_k}^u)^\top + I. \end{aligned}$$

Hence, Prop. 11 ensures that

$$\begin{aligned} D &\geq \text{Tr}(P(\theta_{t_k}, K_{t_k}^u) \geq \lambda_{\min}(C) \|L_{t_k}^\top\|_{\Sigma(\theta_{t_k}, K_{t_k}^u)}^2, \\ D &\geq \text{Tr}(P(\theta_{t_k}, K_{t_k}^u) \geq \lambda_{\min}(C) \|\Sigma(\theta_{t_k}, K_{t_k}^u)^{1/2}\|_F^2, \end{aligned}$$

and we conclude using that $\kappa = D/\lambda_{\min}(C)$.

3. The last assertions come directly from the previous one, noticing that $\Sigma(\theta_{t_k}, K_{t_k}^u) \succcurlyeq I$ and $A^c(\theta_{t_k}, K_{t_k}^u) \Sigma(\theta_{t_k}, K_{t_k}^u) A^c(\theta_{t_k}, K_{t_k}^u)^\top \preccurlyeq \Sigma(\theta_{t_k}, K_{t_k}^u)$.

□

In line with Lem. 1, we now show that for any $k \geq 1$, the controller $K_{t_k}^u$ stabilizes the true system parameterized by θ_* , as long as θ_{t_k} is *sufficiently close* to θ_* and characterize the stability margin both in term of steady-state covariance and cost (Prop. 17). Further, this translates in uniform exponential stability of the induced deterministic state process (Prop. 18). The proofs are identical to propositions 14 and 15 so we only report the formal statements.

Proposition 17. *Let $k \geq 1$ and assume that $\theta_* \in \mathcal{C}_{t_k}$. Let θ_{t_k} and $K_{t_k}^u$ be an optimistic parameter/controller pair as given in Prop. 16. Then, if $\|\theta_{t_k} - \theta_*\|_F \leq 2\epsilon_0$ where ϵ_0 is such that $\eta(\epsilon_0) := 4\kappa\epsilon_0(1 + \epsilon_0) < 1$, under assumptions 2 and 3,*

1. $A^c(\theta_*, K_{t_k}^u)$ is stable,
2. $\Sigma(\theta_*, K_{t_k}^u) \preccurlyeq \frac{1}{1-\eta(\epsilon_0)} \Sigma(\theta_{t_k}, K_{t_k}^u)$,
3. $\|P(\theta_*, K_{t_k}^u)\|_2 \preccurlyeq \|P(\theta_*, K(\theta_*))\|_2 + \frac{D\eta(\epsilon_0)}{1-\eta(\epsilon_0)}$,
4. $\|P(\theta_{t_k}, K_{t_k}^u)\|_2 \leq \|P(\theta_*, K(\theta_*))\|_2(1 + \kappa\eta(\epsilon_0)/2) + \frac{D\eta(\epsilon_0)}{1-\eta(\epsilon_0)}(1 + \kappa\eta(\epsilon_0)/2)$.

Proposition 18. *Assume that $\theta_* \in \mathcal{C}_t$ for all $t \geq 1$. Let $\{\theta_{t_k}\}_{k \geq 1}$ and $\{K_{t_k}^u\}_{k \geq 1}$ be the sequences optimistic parameter/controller pair as given in Prop. 16. Then, if for all $k \geq 0$, $\|\theta_{t_k} - \theta_*\|_F \leq 2\epsilon_0$ where ϵ_0 is such that*

$$\kappa\eta(\epsilon_0) \leq 1; \quad \eta(\epsilon_0) := 4\kappa\epsilon_0(1 + \epsilon_0) \leq 1/2, \quad (45)$$

under assumptions 1 to 3,

1. for all k , the state process driven by $A^c(\theta_*, K_{t_k}^u)$ is uniformly exponentially stable,
2. for all k , for any $0 < \delta < 1$, with probability at least $1 - \delta$,

$$\begin{aligned} \text{for all } t \in [t_k, t_{k+1} - 1], \quad & \|x_t\| \leq 4\sigma\sqrt{\kappa \log(1/\delta)} + \sqrt{\kappa}(1 - 1/\kappa)^t \|\xi\| \\ & \|x_t\| \leq 4\sigma\sqrt{\kappa \log(1/\delta)} + 2\sqrt{\|P(\theta_*, K(\theta_*))\|_2/\lambda_{\min}(C)} \|\xi\|. \end{aligned} \quad (46)$$

E.2.1. STEP 2. PROOF OF LEM. 4

Let $X := 20\sigma\sqrt{\kappa \log(4T/\delta)}\sqrt{\|P(\theta_*, K(\theta_*))\|_2/\lambda_{\min}(C)}$, we are now ready to construct recursively a sequence of high-probability events

$$F_t = \left\{ \forall s \leq t, \|x_s\| \leq X \right\}$$

over which the state process remains bounded. As we did for Lem. 1 in App. D.2, we work on $\{E_t\}_{t \leq T}$ and we proceed by episode on the increasing sequence of probability events $\{E_{t_k} \cap F_{t_k}\}_{k \geq 0}$.

Proof of Lem. 4. We assume that LAGLQ is initialized with a stabilizing set $\Theta_0 = \{\theta : \|\theta - \theta_0\| \leq \epsilon_0\}$ such that ϵ_0 satisfies

$$\kappa\eta(\epsilon_0) \leq 1; \quad \eta(\epsilon'_0) := 4\kappa\epsilon'_0(1 + \epsilon'_0) \leq 1/2; \quad \epsilon'_0 = 2\sqrt{\kappa}\epsilon_0.$$

Further, we assume that the RLS estimation in (6) is instantiated with

$$\lambda = \frac{2n\sigma^2}{\epsilon_0^2} \left(\log(4n/\delta) + (n+d) \log(1 + \kappa X^2 T) \right).$$

Setting $\epsilon_0 \leq \frac{1}{64\kappa^{3/2}}$ ensures that $\eta(\epsilon_0') \leq 1/2$, $\kappa\eta(\epsilon_0) \leq 1$ and that $\lambda \geq 4^3 \sigma^2 \kappa^4 \log(4T/\delta) / \log(2)$.

As opposed to the proof of Lem. 1, controlling $\theta_{t_k} - \theta_*$ where $\theta_{t_k} = \widehat{\theta}_{t_k} + K_{t_k}^w$ requires not only to guarantee that $\widehat{\theta}_{t_k}$ is close to θ_* but also that $\|K_{t_k}^w\|$ is small. While addressing the former can be done in a similar fashion, addressing the later is specific to the Extended LQR with Relaxed Constraint resolution. Let $\widetilde{\pi}_{t_k}$ be the extended policy corresponding to the optimal solution of (16). By construction, we have that $g_{\widetilde{\pi}_{t_k}}(\widehat{\theta}_{t_k}, \beta_{t_k}, V_{t_k}) \leq 0$. Expanding $\widetilde{\pi}_{t_k} = (\pi_{t_k}^u, \pi_{t_k}^w)$ where $\pi_{t_k}^u$ and $\pi_{t_k}^w$ are the linear policy respectively associated with $K_{t_k}^u$ and $K_{t_k}^w$, the constraint $g_{\widetilde{\pi}_{t_k}} \leq 0$ leads to:

$$\lim_{S \rightarrow \infty} \frac{1}{S} \mathbb{E} \left(\sum_{s=0}^S \|w_s\|^2 - \beta_t \|z_s\|_{V_t^{-1}}^2 \right) = \text{Tr}(K_{t_k}^{w\top} K_{t_k}^w \Sigma(\theta_{t_k}, K_{t_k}^u)) - \beta_{t_k}^2 \text{Tr}(L_{t_k}^\top V_{t_k}^{-1} L_{t_k} \Sigma(\theta_{t_k}, K_{t_k}^u)) \leq 0,$$

where $L_{t_k}^\top = (I, K_{t_k}^{u\top})$. As a result, Prop. 16 ensures that

$$\|K_{t_k}^w\|_F^2 \leq \kappa \beta_{t_k}^2 / \lambda_{\min}(V_{t_k}) \leq \kappa \beta_{t_k}^2 / \lambda. \quad (47)$$

We now proceed by induction, for sake of simplicity, we assume that $x_0 = 0$. We show that w.h.p.

- 1) $\|x_t\| \leq X := 20\sigma \sqrt{\kappa \log(4T/\delta)} \sqrt{\|P(\theta_*, K(\theta_*))\|_2 / \lambda_{\min}(C)}$
- 2) $\|x_{t_{k+1}}\| \leq 8\sigma \sqrt{\kappa \log(4T/\delta)}$.

- At $k = 0$, we have

$$\begin{aligned} \|\theta_{t_0} - \theta_*\| &\leq \|\widehat{\theta}_{t_0} - \theta_* + K_{t_0}^w\| \leq (1 + \sqrt{\kappa}) \beta_{t_0} / \sqrt{\lambda_{\min}(V_{t_0})} \\ &\leq (1 + \sqrt{\kappa}) (\epsilon_0 + \sigma \sqrt{2n \log(4n/\delta)} / \sqrt{\lambda}) \\ &\leq 4\sqrt{\kappa} \epsilon_0 := 2\epsilon_0'. \end{aligned}$$

As a result, Prop. 18 guarantees that with probability at least $1 - \delta/4T$,

$$\text{for all } s \leq t_1, \quad \|x_s\| \leq 4\sigma \sqrt{\kappa \log(4T/\delta)} \leq X \quad \text{and} \quad \|x_{t_1}\| \leq 8\sigma \sqrt{\kappa \log(4T/\delta)}.$$

Hence, $\mathbb{P}(F_{t_1} | E_{t_0}) \geq 1 - \delta/4T$.

- At $k + 1$, assume that $E_{t_k} \cap F_{t_k}$ holds. From $\|x_s\| \leq X$ for all $s \leq t_{k+1}$, we obtain that

$$\begin{aligned} \beta_{t_{k+1}} &= \sigma \sqrt{2n \log \left(\frac{\det(V_{t_{k+1}})^{1/2} 4n}{\det(\lambda I)^{1/2} \delta} \right)} + \lambda^{1/2} \|\theta_0 - \theta_*\|_F \\ &\leq \sigma \sqrt{2n \log(4n/\delta) + 2n(n+d) \log \left(1 + \sup_{s \leq t_{k+1}-1} \|z_s\|^2 T \right)} + \lambda^{1/2} \epsilon_0 \\ &\leq \sigma \sqrt{2n \log(4n/\delta) + 2n(n+d) \log(1 + \kappa X^2 T)} + \lambda^{1/2} \epsilon_0 \\ &\leq 2\epsilon_0 \sqrt{\lambda}, \end{aligned}$$

where we used that $z_s = L_s x_s$ and $\|L_s\|_2 \leq \sqrt{\kappa}$ by Prop. 16, and the definition of λ . As a result,

$$\begin{aligned} \|\theta_{t_{k+1}} - \theta_*\| &\leq \|\widehat{\theta}_{t_{k+1}} - \theta_* + K_{t_{k+1}}^w\| \leq (1 + \sqrt{\kappa}) \beta_{t_{k+1}} / \sqrt{\lambda_{\min}(V_{t_{k+1}})} \\ &\leq 2(1 + \sqrt{\kappa}) \epsilon_0 \leq 4\sqrt{\kappa} \epsilon_0 := 2\epsilon_0' \end{aligned}$$

and we can apply Prop. 18. Using that $\|x_{t_k}\| \leq 8\sigma\sqrt{\kappa\log(4T/\delta)}$ we obtain with probability at least $1 - \delta/4T$,

$$\begin{aligned} \text{for all } t_k \leq s \leq t_{k+1}, \quad \|x_s\| &\leq 4\sigma\sqrt{\kappa\log(4T/\delta)} + 2\sqrt{\|P(\theta_*, K(\theta_*))\|_2/\lambda_{\min}(C)}\|x_{t_k}\| \\ &\leq 4\sigma\sqrt{\kappa\log(4T/\delta)}(1 + 4\sqrt{\|P(\theta_*, K(\theta_*))\|_2/\lambda_{\min}(C)}) \\ &\leq X. \end{aligned}$$

Hence, $\mathbb{P}(F_{t_{k+1}}|E_{t_k}, F_{t_k}) \geq 1 - \delta/4T$. Further, on this event,

$$\|x_{t_{k+1}}\| \leq 4\sigma\sqrt{\kappa\log(4T/\delta)}(1 + \sqrt{\kappa}(1 - 1/\kappa)^{t_{k+1}-t_k+1}) \leq 8\sigma\sqrt{\kappa\log(4T/\delta)}$$

since $t_{k+1} - t_k + 1 \geq \kappa\log(\kappa)/2$ otherwise it enters in contradiction with the update rule and concludes the induction. We are now left to prove that $t_{k+1} - t_k + 1 \geq \kappa\log(\kappa)/2$. Suppose that $t_{k+1} - t_k + 1 \leq \kappa\log(\kappa)/2$,

$$\begin{aligned} \det(V_{t_{k+1}}) &\leq \det(V_{t_k}) \prod_{t=t_k}^{t_{k+1}-1} (1 + \|z_t\|_{V_t}^2) \leq \det(V_{t_k}) \prod_{t=t_k}^{t_{k+1}-1} (1 + \|x_t\|^2\kappa/\lambda) \\ &\leq \det(V_{t_k}) \exp\left(\frac{\kappa}{\lambda} \sum_{t=t_k}^{t_{k+1}-1} \|x_t\|^2\right) \leq \det(V_{t_k}) \exp\left(32\sigma^2\kappa^2\log(4T/\delta)((t_{k+1} - t_k + 1) + \kappa^2)/\lambda\right) \\ &\leq \det(V_{t_k}) \exp\left(4^3\sigma^2\kappa^4\log(4T/\delta)/\lambda\right) \leq 2\det(V_{t_k}), \end{aligned}$$

where we used that $\lambda \geq 4^3\sigma^2\kappa^4\log(4T/\delta)/\log(2)$.

Overall, we showed that the sequence of event F_t is such that $\mathbb{P}(F_{t_{k+1}}|F_{t_k}, E_{t_k}) \geq 1 - \delta/4T$ for all $k \geq 0$. As a result, we have

$$\mathbb{P}(F_T \cap E_T) \geq \mathbb{P}(E_T \cap \bigcap_{k \geq 0} F_{t_k} | F_{t_{k-1}}, E_{t_{k-1}}) \geq 1 - \mathbb{P}(E_T^c) - \sum_{k \geq 0} \mathbb{P}\left((F_{t_k} | F_{t_{k-1}}, E_{t_{k-1}})^c\right) \geq 1 - \delta/2.$$

□

E.3. Regret bound (Proof of Thm. 1)

In line with Abbasi-Yadkori & Szepesvári (2011), we conduct the analysis on the sequence of events $\{E_t \cap F_t\}_{t \geq 1}$, where

$$E_t = \{\forall s \leq t, \theta_* \in \mathcal{C}_s\}, \quad F_t = \{\forall s \leq t, \|x_s\| \leq X\}, \quad X := 20\sigma\sqrt{\kappa\log(4T/\delta)}\sqrt{\|P(\theta_*)\|_2/\lambda_{\min}(C)},$$

which holds with probability at least $1 - \delta/2$ by Lem. 4. We further assume for sake of simplicity that at each update, the controller exactly solves the extended problem in (16) since, as discussed in Sect. 5.2, one can solve it with arbitrarily accuracy.

E.3.1. REGRET DECOMPOSITION

For any $t \leq T$, let $t_k \leq t$ be the last time step at which has been updated. Let $\tilde{K}_t = \tilde{K}_{t_k}$ be the extended controller and let $\tilde{K}_t = (K_t^{u\top} \quad K_t^{w\top})^\top$ be the blocks corresponding to control u and perturbation w respectively. Then,

$$\begin{aligned} \mathcal{J}_*(\hat{\theta}_{t_k}, \beta_{t_k}, V_{t_k}) &= \text{Tr}(P_t) \quad \text{where} \\ P_t &= Q + K_t^{u\top} R K_t^u + (\hat{A}_{t_k} + \hat{B}_{t_k} K_t^u + K_t^w)^\top P_t (\hat{A}_{t_k} + \hat{B}_{t_k} K_t^u + K_t^w) \end{aligned} \quad (48)$$

which leads to the perturbed Bellman equation

$$\mathcal{J}_*(\hat{\theta}_t, \beta_t, V_t) + x_t^\top P_t x_t = x_t^\top Q x_t + u_t^\top R u_t + \mathbb{E}(\tilde{x}_{t+1}^\top P_t \tilde{x}_{t+1} | \mathcal{F}_t), \quad \text{where } \tilde{x}_{t+1} = \hat{\theta}_t^\top z_t + w_t + \epsilon_{t+1}. \quad (49)$$

Introducing, $x_{t+1} = \theta_*^\top x_t + \epsilon_{t+1} = (\theta_* - \hat{\theta}_{t_k})^\top z_t - w_t + \tilde{x}_{t+1}$ leads to:

$$\begin{aligned} \mathcal{J}_*(\hat{\theta}_{t_k}, \beta_{t_k}, V_{t_k}) + x_t^\top P_t x_t &= x_t^\top Q x_t + u_t^\top R u_t + \mathbb{E}(x_{t+1}^\top P_t x_{t+1} | \mathcal{F}_t) \\ &\quad + (\hat{\theta}_{t_k}^\top z_t + w_t)^\top P_t (\hat{\theta}_{t_k}^\top z_t + w_t) - (\theta_*^\top z_t)^\top P_t (\theta_*^\top z_t). \end{aligned} \quad (50)$$

Summing over $[0, T]$, and performing similar manipulation as in (Abbasi-Yadkori & Szepesvári, 2011) leads to the regret decomposition:

$$\begin{aligned}
 R(T) &= \sum_{t=0}^T x_t^\top Q x_t + u_t^\top R u_t - J_* = \sum_{t=0}^T \mathcal{J}_*(\hat{\theta}_{t_k}, \beta_{t_k}, V_{t_k}) - J_* \\
 &\quad + \sum_{t=0}^T x_t^\top P_t x_t - \mathbb{E}(x_{t+1}^\top P_{t+1} x_{t+1} | \mathcal{F}_t) \\
 &\quad + \sum_{t=0}^T \mathbb{E}(x_{t+1}^\top (P_{t+1} - P_t) x_{t+1} | \mathcal{F}_t) \\
 &\quad + \sum_{t=0}^T (\hat{\theta}_{t_k}^\top z_t + w_t)^\top P_t (\hat{\theta}_{t_k}^\top z_t + w_t) - (\theta_*^\top z_t)^\top P_t (\theta_*^\top z_t).
 \end{aligned} \tag{51}$$

We now can bound each term separately. While the analysis is similar to the one of Abbasi-Yadkori & Szepesvári (2011) but for the last term, we provide here a full analysis, for sake of completeness.

Bounding $R^{opt}(T) = \sum_{t=0}^T \mathcal{J}_*(\hat{\theta}_{t_k}, \beta_{t_k}, V_{t_k}) - J_*$. Lem. 3 guarantees that on E_t , $\mathcal{J}_*(\hat{\theta}_{t_k}, \beta_{t_k}, V_{t_k}) \leq J_*$. As a result, $R^{opt}(T) \leq 0$. Further, by Prop. 17 and the tuning of ϵ_0 , $\|P_t\|_2 = O(\|P(\theta_*)\|_2)$.

Bounding $R^{mart}(T) = \sum_{t=0}^T x_t^\top P_t x_t - \mathbb{E}(x_{t+1}^\top P_{t+1} x_{t+1} | \mathcal{F}_t)$. On $E_t \cap F_t$, Lem. 4 guarantees that $\|x_t\| \leq X$ where $X = 20\sigma\sqrt{\kappa \log(4T/\delta)}\sqrt{\|P(\theta_*)\|_2/\lambda_{\min}(C)}$. Thus, $R^{mart}(T)$ is a bounded martingale sequence and applying Azuma's inequality leads to, with probability at least $1 - \delta/4$,

$$R^{mart}(T) = \tilde{O}(X^2 \|P(\theta_*)\|_2 \sqrt{T}) = \tilde{O}(\kappa \|P(\theta_*)\|_2^2 \sqrt{T}).$$

Bounding $R^{lazy}(T) = \sum_{t=0}^T \mathbb{E}(x_{t+1}^\top (P_{t+1} - P_t) x_{t+1} | \mathcal{F}_t)$. LAGLQ uses, as in (Abbasi-Yadkori & Szepesvári, 2011) a lazy update scheme to re-evaluate the controllers K_{t_s} , which translates into $P_t = P_{t+1}$ for all $t \notin \{t_k\}_{k=1}^m$ where $t_k = \min\{t \geq t_{k-1} \text{ s.t. } \det(V_t) \leq 2 \det(V_{t_{k-1}})\}$. Thus, on $E_t \cap F_t$, one has $R^{lazy}(T) \leq 2X^2 Dm$. Finally, Lem. 8 in (Abbasi-Yadkori & Szepesvári, 2011) ensures that the number of updates scales logarithmically,

$$m \leq (n+d) \log_2(1 + TX^2 \kappa / \lambda). \tag{52}$$

which implies that $R^{lazy}(T) = \tilde{O}(X^2 \|P(\theta_*)\|_2 (n+d)) = \tilde{O}(\kappa \|P(\theta_*)\|_2^2 (n+d))$.

Bounding $R^{pred}(T) = \sum_{t=0}^T (\hat{\theta}_{t_k}^\top z_t + w_t)^\top P_t (\hat{\theta}_{t_k}^\top z_t + w_t) - (\theta_*^\top z_t)^\top P_t (\theta_*^\top z_t)$. Using $\|P_t\| \leq \|P(\theta_*)\|_2$ together with the triangular inequality leads to:

$$\begin{aligned}
 R^{pred}(T) &= \sum_{t=0}^T \|\hat{\theta}_{t_k}^\top z_t + w_t\|_{P_t}^2 - \|\theta_*^\top z_t\|_{P_t}^2 = \sum_{t=0}^T (\|\hat{\theta}_{t_k}^\top z_t + w_t\|_{P_t} - \|\theta_*^\top z_t\|_{P_t}) (\|\hat{\theta}_{t_k}^\top z_t + w_t\|_{P_t} + \|\theta_*^\top z_t\|_{P_t}) \\
 &\leq \|P(\theta_*)\|_2^{1/2} \sum_{t=0}^T \|(\hat{\theta}_{t_k} - \theta_*)^\top z_t + w_t\| (\|\hat{\theta}_{t_k}^\top z_t + w_t\|_{P_t} + \|\theta_*^\top z_t\|_{P_t})
 \end{aligned}$$

Further, from (48), $\|\hat{\theta}_{t_k}^\top z_t + w_t\|_{P_t} = \|(\hat{A}_{t_k} + \hat{B}_{t_k} K_{u,t} + K_{w,t}) x_t\|_{P_t} \leq \|x_t\|_{P_t} \leq \|P(\theta_*)\|_2^{1/2} X$. Similarly, on $E_t \cap F_t$,

$$\begin{aligned}
 \|\theta_*^\top z_t\|_{P_t} &\leq \|\hat{\theta}_{t_k}^\top z_t + w_t\|_{P_t} + \|(A_* + B_* K_{u,t}) x_t - (\hat{A}_{t_k} + \hat{B}_{t_k} K_{u,t} + K_{w,t}) x_t\|_{P_t} \\
 &\leq \|P(\theta_*)\|_2^{1/2} X + 4\|P(\theta_*)\|_2^{1/2} X \kappa \epsilon_0 \\
 &\leq 2X \|P(\theta_*)\|_2^{1/2}.
 \end{aligned}$$

where we used $\|\theta_{t_k} - \theta_*\| \leq 4\sqrt{\kappa}\epsilon_0$ from line 1 to 2, and the condition on ϵ_0 from line 2 to 3. Thus,

$$R^{pred}(T) \leq 3X\|P(\theta_*)\|_2 \left(\underbrace{\sum_{t=0}^T \|(\hat{\theta}_{t_k} - \theta_*)^\top z_t\|}_{:=R_u^{pred}(T)} + \underbrace{\sum_{t=0}^T \|w_t\|}_{:=R_w^{pred}(T)} \right).$$

E.3.2. BOUNDING THE CUMULATIVE PERTURBATIONS

Bounding $R_u^{pred}(T)$. In line with (Abbasi-Yadkori & Szepesvári, 2011) (Lem. 12), from Prop. 1, one has on E_t ,

$$\|(\hat{\theta}_{t_k} - \theta_*)^\top z_t\| = \|(\hat{\theta}_{t_k} - \theta_*)^\top V_{t_k}^{1/2} V_{t_k}^{-1/2} z_t\| \leq \beta_{t_k} \|z_t\|_{V_{t_k}^{-1}} \leq 2\beta_{t_k} \|z_t\|_{V_t^{-1}}$$

which implies, using Cauchy-Schwarz and Prop. 2, $R_u^{pred}(T) \leq 2\beta_T \sqrt{T} \left(\sum_{t=0}^T \|z_t\|_{V_t^{-1}}^2 \right)^{1/2}$.

$$R_u^{pred}(T) \leq 2\beta_T \sqrt{T} \left(\sum_{t=0}^T \|z_t\|_{V_t^{-1}}^2 \right)^{1/2} \leq 2\beta_T \gamma_u \sqrt{T}$$

$$\text{where } \gamma_u = (2(n+d) \log(1+T))^{1/2}.$$

Bounding $R_w^{pred}(T)$. So far the regret analysis, that relied mostly on optimism and Prop. 1, was very similar than the one of Abbasi-Yadkori & Szepesvári (2011). The following Lemma which ensures the perturbation w_t to be cumulatively bounded is specific to the new perturbed approach presented in LAGLQ.

Lemma 8. *On event $E_t \cap F_t$, with probability at least $1 - \delta/4$, $R_w^{pred}(T) = \tilde{O}(\sqrt{\kappa}X\beta_T\gamma_u\sqrt{T})$.*

Proof. As discussed in Sec. 4, if the perturbation w_t were derived from an optimistic $\tilde{\theta}_{t_k} \in C_{t_k}$, it would translate into the constraint $\|w_t\| \leq \beta_{t_k} \|z_t\|_{V_{t_k}^{-1}}$ and Lem. 8 would trivially hold. Unfortunately, this constraint cannot be enforced as it does not lead to a *feasible* extended LQR problem. To overcome this issue, we used in problem (16) the relaxed constraint in Eq. 15. Formally, we only guarantee that

$$g_{\tilde{\pi}}(\hat{\theta}_{t_k}, \beta_{t_k}, V_{t_k}) = \lim_{S \rightarrow \infty} \frac{1}{S} \mathbb{E} \left(\sum_{s=0}^S \|w_s\|^2 - \beta_{t_k}^2 \|z_s\|_{V_{t_k}^{-1}}^2 \right) \leq 0,$$

where $\{x_s\}_{s \geq 0}$ follows the extended dynamic in Eq. 14 parametrized by $\hat{\theta}_{t_k}$. This equivalently translates in a constraint on $K_{t_k}^u$ and $K_{t_k}^w$, i.e.,

$$\text{Tr}(K_{t_k}^{w\top} K_{t_k}^w \Sigma(\theta_{t_k}, K_{t_k}^u)) \leq \beta_{t_k}^2 \text{Tr}(L_{t_k}^\top V_{t_k}^{-1} L_{t_k} \Sigma(\theta_{t_k}, K_{t_k}^u)), \quad (53)$$

where $L_{t_k} = \begin{pmatrix} I \\ K_{t_k}^u \end{pmatrix}$ and $\Sigma(\theta_{t_k}, K_{t_k}^u)$ is the steady-state variance of the process $\{x_s\}_{s \geq 0}$. Further, Prop. 16 guarantees that $I \preceq \Sigma(\theta_{t_k}, K_{t_k}^u)$ and $\text{Tr}(\Sigma(\theta_{t_k}, K_{t_k}^u)) \leq \kappa$. As a result,

$$\text{Tr}(K_{t_k}^{w\top} K_{t_k}^w) \leq \kappa \beta_{t_k}^2 \lambda_{\max}(L_{t_k}^\top V_{t_k}^{-1} L_{t_k}) \Rightarrow \|K_{t_k}^w\|_2 = \|K_{t_k}^w\|_2 \leq \sqrt{\kappa} \beta_{t_k} \|V_{t_k}^{-1/2} L_{t_k}\|_2$$

As a result, one gets:

$$R_w^{pred}(T) \leq X \sum_{t=0}^T \|K_t^w\|_2 \leq \sqrt{\kappa} \beta_T X \sum_{t=0}^T \|V_{t_k}^{-1/2} L_{t_k}\|_2 \mathbf{1}_{t \in [t_k, t_{k+1}-1]}.$$

At a high-level, proving Lem. 8 turns into proving that the cumulative sum of weighted *controllers* $\sum_{t=0}^T \|L_t\|_{V_t^{-1}}$ is bounded by \sqrt{T} , whereas Prop. 2 only provides a bound for the cumulative sum of weighted *controls* $\sum_{t=0}^T \|z_t\|_{V_t^{-1}} = \sum_{t=0}^T \|L_t x_t\|_{V_t^{-1}}$. However, as hinted in (Abeille & Lazaric, 2018), those latter quantity is related to the exploration

performed by covariates $z_t = L_t x_t$, which are driven by L_t . Intuitively, since, on F_{t-1} , $\mathbb{V}(x_t | \mathcal{F}_{t-1}) \geq I$, x_t covers on average all directions, turning into an accurate exploration on every directions of L_t . Formally, for any $t > t_k$ (i.e., for every time steps strictly within an episode), V_{t_k}, L_{t_k} are \mathcal{F}_{t-1} measurable and so are $\lambda_{\max}(L_{t_k}^\top V_{t_k}^{-1} L_{t_k})$ and $v_{t_k}^{\max}$. As a result, we have

$$\|L_{t_k} x_t\|_{V_{t_k}^{-1}} \geq \|V_{t_k}^{-1/2} L_{t_k} v_{t_k}^{\max} x_t^\top v_{t_k}^{\max}\| \geq \|V_{t_k}^{-1/2} L_{t_k}\|_2 |x_t^\top v_{t_k}^{\max}|. \quad (54)$$

Without loss of generality, we can assume that $x_{t-1}^\top A^c(\theta_*, K_{t_k}^u)^\top v_{t_k} \geq 0$ and hence,

$$|x_t^\top v_{t_k}^{\max}| \geq |x_{t-1}^\top v_{t_k}^{\max}| \mathbb{1}_{\epsilon_t^\top v_{t_k}^{\max} \geq 0} \geq \epsilon_t \mathbb{1}_{\epsilon_t^\top v_{t_k}^{\max} \geq 0}.$$

Applying property 1 to $\epsilon_t^\top v_{t_k}^{\max}$ and taking the expectation in (54) leads to

$$\|L_{t_k} x_t\|_{V_{t_k}^{-1}} \leq 64\sigma^3 \mathbb{E}(\|L_{t_k} x_t\|_{V_{t_k}^{-1}} | \mathcal{F}_{t-1}, F_{t-1}).$$

Further, we have to guarantee that the state x_t remains bounded, and hence work under $F_t = \{s \leq t \text{ s.t. } \|x_s\| \leq X\}$ (notice that the conditioning only ensures that we work under F_{t-1}). To do so, we use property 2 and obtain:

$$\begin{aligned} \|L_{t_k} x_t\|_{V_{t_k}^{-1}} &\leq 64\sigma^3 \mathbb{E}(\|L_{t_k} x_t\|_{V_{t_k}^{-1}} \mathbb{1}_{\|x_t\| \leq X} | \mathcal{F}_{t-1}, F_{t-1}) + 64\sigma^3 \sqrt{\kappa/\lambda} \mathbb{E}(\|x_t\| \mathbb{1}_{\|x_t\| \geq X} | \mathcal{F}_{t-1}, F_{t-1}) \\ &\leq 64\sigma^3 \mathbb{E}(\|L_{t_k} x_t\|_{V_{t_k}^{-1}} \mathbb{1}_{\|x_t\| \leq X} | \mathcal{F}_{t-1}, F_{t-1}) + \sigma^3 \delta / T. \end{aligned}$$

Notice that $\{t = t_k\}$ is a condition on V_t and hence is \mathcal{F}_{t-1} measurable. Thus,

$$\begin{aligned} \sum_{t=0}^T \|L_{t_k} x_t\|_{V_{t_k}^{-1}} \mathbb{1}_{t \in [t_k, t_{k+1}-1]} &= \sum_{t=0}^T \|L_{t_k} x_t\|_{V_{t_k}^{-1}} \mathbb{1}_{t \in [t_k+1, t_{k+1}-1]} + \sum_{t=0}^T \|L_{t_k} x_t\|_{V_{t_k}^{-1}} \mathbb{1}_{t=t_k} \\ &= \sum_{t=0}^T \|L_{t_k} x_t\|_{V_{t_k}^{-1}} \mathbb{1}_{t \in [t_k+1, t_{k+1}-1]} + \sqrt{\kappa/\lambda} m \\ &\leq 64\sigma^3 \sum_{t=0}^T \mathbb{E}(\|L_{t_k} x_t\|_{V_{t_k}^{-1}} \mathbb{1}_{\|x_t\| \leq X} | \mathcal{F}_{t-1}, F_{t-1}) \mathbb{1}_{t \in [t_k+1, t_{k+1}-1]} + \sigma^3 \delta + \sqrt{\kappa/\lambda} m \\ &\leq 128\sigma^3 \sum_{t=0}^T \mathbb{E}(\|z_t\|_{V_t^{-1}} \mathbb{1}_{\|x_t\| \leq X} | \mathcal{F}_{t-1}, F_{t-1}) \mathbb{1}_{t \in [t_k+1, t_{k+1}-1]} + \sigma^3 \delta + \sqrt{\kappa/\lambda} m \end{aligned}$$

Finally, using Azuma's inequality, with probability at least $1 - \delta/4$

$$\sum_{t=0}^T \mathbb{E}(\|z_t\|_{V_t^{-1}} \mathbb{1}_{\|x_t\| \leq X} | \mathcal{F}_{t-1}) \leq \sum_{t=0}^T \|z_t\|_{V_t^{-1}} + \sqrt{\log(4/\delta)T},$$

We conclude using Cauchy-Schwarz and Prop. 2,

$$\sum_{t=0}^T \mathbb{E}(\|z_t\|_{V_t^{-1}} \mathbb{1}_{\|x_t\| \leq X} | \mathcal{F}_{t-1}) \leq \gamma_u \sqrt{T} + \sqrt{\log(4/\delta)T}.$$

Finally, we have

$$R_w^{\text{pred}}(T) = \tilde{O}(\sqrt{\kappa} X \beta_T \gamma_u \sqrt{T}).$$

□

Bounding $R^{\text{pred}}(T)$. Summarizing the bounds on $R_u^{\text{pred}}(T)$ and $R_w^{\text{pred}}(T)$ we obtain that with probability at least $1 - \delta/4$,

$$R^{\text{pred}}(T) \leq 3X \|P(\theta_*)\|_2 (R_u^{\text{pred}}(T) + R_w^{\text{pred}}(T)) = \tilde{O}\left((1 + \sqrt{\kappa} X) X \|P(\theta_*)\|_2 \beta_T \gamma_u \sqrt{T}\right).$$

From $\beta_T = \tilde{O}(\sqrt{n(n+d)})$, $X = \tilde{O}(\sqrt{\kappa} \|P(\theta_*)\|_2)$ and $\gamma_u = \tilde{O}(\sqrt{(n+d)})$ we obtain,

$$R^{\text{pred}}(T) = \tilde{O}((n+d) \sqrt{n} \kappa^{3/2} \|P(\theta_*)\|_2^2 \sqrt{T}).$$

E.3.3. PUTTING EVERYTHING TOGETHER

Since the regret decomposition and the analysis were derived on $\{E_t \cap F_t\}_{t \geq 1}$, we have that with probability at least $1 - \delta/2$,

$$R(T) \leq R^{opt}(T) + R^{mart}(T) + R^{lazy}(T) + R^{pred}(T).$$

Further, on $\{E_t \cap F_t\}_{t \geq 1}$, $R^{opt}(T) \leq 0$ and $R^{lazy}(T) = \tilde{O}(\kappa \|P(\theta_*)\|_2^2(n+d))$ while $R^{mart}(T) = \tilde{O}(\kappa \|P(\theta_*)\|_2^2 \sqrt{T})$ with probability at least $1 - \delta/4$ and $R^{pred}(T) = \tilde{O}((n+d)\sqrt{n}\kappa^{3/2} \|P(\theta_*)\|_2^2 \sqrt{T})$ with probability at least $1 - \delta/4$. As a result, a union bound argument ensures that we have, with probability at least $1 - \delta$,

$$R(T) = \tilde{O}((n+d)\sqrt{n}\kappa^{3/2} \|P(\theta_*)\|_2^2 \sqrt{T}).$$

E.3.4. TECHNICAL RESULTS

Property 1. Let ϵ be a zero-mean, σ -subGaussian random variable such that $\mathbb{V}(\epsilon) = 1$. Then,

$$\mathbb{E}(\epsilon \mathbf{1}_{\epsilon \geq 0}) = -\mathbb{E}(-\epsilon \mathbf{1}_{\epsilon \leq 0}) \geq 1/(64\sigma^3).$$

Proof of property 1. From $\mathbb{E}(\epsilon) = 0$, we obtain that

$$0 = \mathbb{E}(\epsilon \mathbf{1}_{\epsilon \geq 0}) + \mathbb{E}(-\epsilon \mathbf{1}_{\epsilon \leq 0}),$$

which provides the l.h.s equality. In particular, this implies that $\mathbb{E}(|\epsilon|) = 2\mathbb{E}(\epsilon \mathbf{1}_{\epsilon \geq 0})$. Further, using Holder's inequality one has:

$$1 = \mathbb{E}(\epsilon^2)^2 \leq \mathbb{E}(|\epsilon|^3) \mathbb{E}(|\epsilon|) \leq \mathbb{E}(|\epsilon|^4)^{3/4} \mathbb{E}(|\epsilon|) \leq (32\sigma^4)^{3/4} \mathbb{E}(|\epsilon|),$$

where we used the properties of subGaussian r.v. to obtain the last inequality. As a result,

$$\mathbb{E}(\epsilon \mathbf{1}_{\epsilon \geq 0}) \geq \frac{1}{64\sigma^3}.$$

□

Property 2. On F_{t-1} , one has that

$$\mathbb{E}(\|x_t\| \mathbf{1}_{\|x_t\| \geq X} | \mathcal{F}_{t-1}) \leq \frac{\delta X}{2T}$$

Proof of property 2. The proof borrows many steps in the proof of Prop. 7, that we omit for sake of readability. Let $f_{\|x_t\|}$ and $\bar{F}_{\|x_t\|}$ be respectively the pdf and complementary cdf of $\|x_t\|$ conditionally to \mathcal{F}_{t-1} , F_{t-1} . Then,

$$\mathbb{E}(\|x_t\| \mathbf{1}_{\|x_t\| \geq X} | \mathcal{F}_{t-1}) = \int_X^{+\infty} u f_{\|x_t\|}(u) du = X \bar{F}_{\|x_t\|}(X) + \int_X^{+\infty} \bar{F}_{\|x_t\|}(u) du.$$

From Prop. 7, we have that for all $u \geq X$, $\bar{F}_{\|x_t\|}(u) \leq e^{-(u/12\sigma\kappa)^2}$. Thus, $X \bar{F}_{\|x_t\|}(X) \leq \frac{\delta X}{4T}$ and

$$\int_X^{+\infty} \bar{F}_{\|x_t\|}(u) du \leq \frac{(12\sigma\kappa)^2}{2X} \int_X^{+\infty} -[e^{-(u/12\sigma\kappa)^2}]' du \leq \frac{(12\sigma\kappa)^2}{2X} e^{-(X/12\sigma\kappa)^2} \leq \frac{\delta X}{4T}.$$

□

F. Lagrangian LQR and Strong Duality

In this section, we study the constrained and Lagrangian LQR problem and in particular focus on proving strong duality.

Set of constrained LQR system. We summarize the parametrization of constrained LQR systems in $\mathfrak{s} = (A, \tilde{B}, C_\dagger, C_g)$ where A and \tilde{B} are the parameters of the dynamic, C_\dagger is the matrix parametrizing the original cost function and C_g the matrix parametrizing the constraints. We focus on LQR system which exhibit a certain structure, and denote as \mathfrak{S} such set of system. Formally,

Definition 2. \mathfrak{S} is the set of system $\mathfrak{s} = (A, \tilde{B}, C_\dagger, C_g)$ which has the following structure:

$$\begin{aligned} A &\in \mathbb{R}^{n \times n}, \tilde{B} = (B; I_n), B \in \mathbb{R}^{n \times d}, \\ C_\dagger &= \begin{pmatrix} C_\dagger^1 & 0 \\ 0 & 0 \end{pmatrix} + C_\dagger^2, \quad C_\dagger^1 \in \mathbb{S}_{++}^{n+d}; \quad C_\dagger^2 \in \mathbb{S}_+^{2n+d}; \\ C_g &= \begin{pmatrix} -V^{-1} & 0 \\ 0 & I_n \end{pmatrix}, V^{-1} \in \mathbb{S}_{++}^{n+d}, \end{aligned}$$

and for which it exists at least a feasible stable linear policy $\tilde{\pi}$ characterized by a linear controller \tilde{K} , i.e., such that $\rho(A^c(\tilde{K})) < 1$ and $g_{\tilde{\pi}}(\mathfrak{s}) \leq 0$, where $A^c(\tilde{K}) = A + \tilde{B}\tilde{K}$ and $g_{\tilde{\pi}}(\mathfrak{s})$ is given in (57).

The reason why we study strong-duality for a larger set of system \mathfrak{S} is motivated by the fact that we may have to slightly modify the original constrained LQR problem in (55) to enforce numerical stability of the algorithmic resolution.

Notice that for all $t \geq 0$, the constrained LQR problem LAGLQ has to solve is indeed contained in \mathfrak{S} whenever $\theta_* \in \mathcal{C}_t$. At each time $t \geq 0$, it corresponds to the parametrization

$$C_\dagger^1 = \begin{pmatrix} Q & 0 \\ 0 & R \end{pmatrix}; \quad C_\dagger^2 = 0; \quad C_g = \begin{pmatrix} -\beta_t^2 V_t^{-1} & 0 \\ 0 & I \end{pmatrix}; \quad A = \hat{A}_t; \quad \tilde{B} = (\hat{B}_t; I), \quad (55)$$

which belongs to \mathfrak{S} since $\beta_t^2 V_t^{-1} \in \mathbb{S}_{++}^{n+d}$ and $\begin{pmatrix} Q & 0 \\ 0 & R \end{pmatrix} \in \mathbb{S}_{++}^{n+d}$ under Asm. 2. Further, it exists a feasible stable linear policy from Lem. 3¹¹. As a consequence, the statements we prove in appendices G and H are slightly more general than the one displayed in the main text, which follow as corollaries. More in details, Thm. 2 is implied by Thm. 4, Lem. 5 and Prop. 3 follow from Lem. 9, lemmas 6 and 7 are respectively proven in Lem. 14 and Prop. 23.

From constrained to Lagrangian problem. For any $\mathfrak{s} \in \mathfrak{S}$, we consider the constrained problem

$$\begin{aligned} \mathcal{J}_*(\mathfrak{s}) &= \min_{\tilde{\pi}} \mathcal{J}_{\tilde{\pi}}(\mathfrak{s}) := \limsup_{S \rightarrow \infty} \frac{1}{S} \mathbb{E} \left[\sum_{s=0}^S \begin{pmatrix} x_s \\ \tilde{u}_s \end{pmatrix}^\top C_\dagger \begin{pmatrix} x_s \\ \tilde{u}_s \end{pmatrix} \right], \\ \text{subject to} & \quad x_{s+1} = Ax_s + \tilde{B}\tilde{u}_s + \epsilon_{s+1} \\ & \quad g_{\tilde{\pi}}(\mathfrak{s}) \leq 0 \end{aligned}, \quad (56)$$

where

$$g_{\tilde{\pi}}(\mathfrak{s}) = \lim_{S \rightarrow \infty} \frac{1}{S} \mathbb{E} \left(\sum_{s=0}^S \begin{pmatrix} x_s \\ \tilde{u}_s \end{pmatrix}^\top C_g \begin{pmatrix} x_s \\ \tilde{u}_s \end{pmatrix} \right). \quad (57)$$

Notice that (56) and (57) exactly correspond to the constrained LQR problem LAGLQ has to solve at time step t in (16) and (15) whenever \mathfrak{s} is defined by (55).

For any $\mathfrak{s} \in \mathfrak{S}$, the Lagrangian formulation of the extended LQR share the same dynamics as in (56), while the cost is defined as the combination of the original average cost \mathcal{J} , and the constraint $g_{\tilde{\pi}}$ as

$$\mathcal{L}_{\tilde{\pi}}(\mu; \mathfrak{s}) = \mathcal{J}_{\tilde{\pi}}(\mathfrak{s}) + \mu g_{\tilde{\pi}}(\mathfrak{s}). \quad (58)$$

¹¹The feasibility condition actually holds for almost every parametrization. Indeed, a system \mathfrak{s} with the appropriate structure is feasible as soon as (A, B) is stabilizable. Since the set of controllable pairs (and hence stabilizable pairs) is open and dense, this is the case for almost every systems.

Furthermore, it can be conveniently written as a quadratic cost function

$$\begin{aligned} \mathcal{L}_{\tilde{\pi}}(\mu; \mathfrak{s}) &= \lim_{S \rightarrow \infty} \frac{1}{S} \mathbb{E} \left[\sum_{s=0}^{S-1} (x_s^\top \quad \tilde{u}_s^\top) C_\mu \begin{pmatrix} x_s \\ \tilde{u}_s \end{pmatrix} \right] \\ \text{subject to } x_{s+1} &= Ax_s + \tilde{B}\tilde{u}_s + \epsilon_{s+1}, \quad \tilde{u}_s = \tilde{\pi}(x_s) \quad \forall s \geq 0 \end{aligned} \quad (59)$$

where $C_{\mu, \nu} = C_\dagger + \mu C_g$. Finally, we will denote as $\tilde{\pi}_\mu(\mathfrak{s})$ the optimal policy achieving $\min_{\tilde{\pi}} \mathcal{L}_{\tilde{\pi}}(\mu; \mathfrak{s})$ for the system \mathfrak{s} and Lagrangian parameter μ whenever it exists.

Strong duality. We are concerned in this section in proving that for any $\mathfrak{s} \in \mathfrak{S}$, strong duality holds i.e. $\mathcal{J}_*(\mathfrak{s}) = \sup_{\mu} \min_{\tilde{\pi}} \mathcal{L}_{\tilde{\pi}}(\mu; \mathfrak{s})$. Further, we derive an algorithm that can find efficiently a feasible linear policy $\tilde{\pi}_*$ such that $\mathcal{J}_{\tilde{\pi}_*}(\mathfrak{s})$ is arbitrarily close to $\mathcal{J}_*(\mathfrak{s})$. The proofs are structured as follow:

- In App. G, we study the inner minimization problem $\tilde{\pi}_\mu(\mathfrak{s}) = \arg \min_{\tilde{\pi}} \mathcal{L}_{\tilde{\pi}}(\mu; \mathfrak{s})$ and characterize a set \mathcal{M} such that for all $\mu \in \mathcal{M}$, the optimal policy is linear in the state, i.e., $\tilde{\pi}_\mu(\mathfrak{s})(x) = \tilde{K}_\mu x$. Those results are summarized in Lem. 9.

Lemma 9. For any $\mathfrak{s} \in \mathfrak{S}$, it exists $\tilde{\mu} > 0$ such that on $\mathcal{M} = [0, \tilde{\mu}]$ the following properties hold

1. The extended Lagrangian LQ in (59) admits a unique solution $\tilde{\pi}_\mu(\mathfrak{s}) = \arg \min_{\tilde{\pi}} \mathcal{L}_{\tilde{\pi}}(\mu; \mathfrak{s})$ obtained by solving a discrete algebraic Riccati equation. As a result, $\tilde{\pi}_\mu(\mathfrak{s})(x) = K_\mu x$ and $\mathcal{D}(\mu, \mathfrak{s}) = \text{Tr}(P_\mu)$ where P_μ (resp. K_μ) are solution of the Riccati equation (resp. the optimal control) associated with (A, \tilde{B}, C_μ) .
 2. $\mathcal{D}(\cdot, \mathfrak{s})$ is concave, $\mathcal{D}(\cdot, \mathfrak{s}) \in \mathcal{C}^1(\mathcal{M})$ and the derivative $\mathcal{D}'(\mu, \mathfrak{s}) = g_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s})$, i.e., it coincides with the constraint g evaluated at the optimal control $\tilde{\pi}_\mu(\mathfrak{s})$.
- In App. H, we leverage the structure of the inner minimization solution to show that strong duality holds (see Thm. 4)

Theorem 4. For any $\mathfrak{s} \in \mathfrak{S}$, let \mathcal{M} be the admissible Riccati set associated with \mathfrak{s} as defined in (72). Then,

$$\mathcal{J}_*(\mathfrak{s}) = \sup_{\mu \in \mathcal{M}} \min_{\tilde{\pi}} \mathcal{L}_{\tilde{\pi}}(\mu; \mathfrak{s}).$$

Further, for any $\epsilon > 0$, it exists $\mu^\epsilon \in \mathcal{M}$ and a linear policy $\tilde{\pi}^\epsilon$ such that

$$1. g_{\tilde{\pi}^\epsilon}(\mathfrak{s}) \leq 0, \quad 2. \mu^\epsilon g_{\tilde{\pi}^\epsilon}(\mathfrak{s}) = 0, \quad 3. J_*(\mathfrak{s}) \geq \mathcal{D}(\mu^\epsilon, \mathfrak{s}) \geq \mathcal{J}_*(\mathfrak{s}) - \epsilon.$$

Lyapunov structure of $\mathcal{L}_{\tilde{\pi}}$, $\mathcal{J}_{\tilde{\pi}}$, $g_{\tilde{\pi}}$ under linear controller. As explained in the previous paragraph, we aim to show that the optimal policy of the inner minimization problem is linear in the state, and hence we will ultimately focus on such structured policies. Before entering the proof detail, we recall here an important technical result that allows us to express all the Lagrangian quantities $\mathcal{L}_{\tilde{\pi}}$, $\mathcal{J}_{\tilde{\pi}}$, $g_{\tilde{\pi}}$ as solutions of Lyapunov equation, as long as the policy $\tilde{\pi}$ is linear and stable, i.e., $\tilde{\pi}(x) = \tilde{K}x$ and $A^c(\tilde{K}) = A + \tilde{B}\tilde{K}$ is stable.

Proposition 19. For all $\mathfrak{s} \in \mathfrak{S}$, for all $\mu \geq 0$ and for all linear policy $\tilde{\pi}(x) = \tilde{K}x$ such that $A^c(\tilde{K}) = A + \tilde{B}\tilde{K}$ is stable,^{12,13}

$$\mathcal{J}_{\tilde{\pi}}(\mathfrak{s}) = \text{Tr}(P(\tilde{K})); \quad g_{\tilde{\pi}}(\mathfrak{s}) = \text{Tr}(G(\tilde{K})); \quad \mathcal{L}_{\tilde{\pi}}(\mu, \mathfrak{s}) = \text{Tr}(P_\mu(\tilde{K})); \quad (60)$$

¹²Notice the crucial difference between $\text{Tr}(P_\mu(\tilde{K}))$ that coincides with the average expected cost of the *Lagrangian* LQR, which is characterized by the cost matrix C_μ , and $\text{Tr}(P(\tilde{K}))$ that corresponds to the average expected cost of the *extended* LQR, which is characterized by the cost C_\dagger .

¹³For sake of readability, we omit the dependency in \mathfrak{s} of the Lyapunov solutions $P(\tilde{K})$, $G(\tilde{K})$ and $P_\mu(\tilde{K})$, although there are entirely characterized by \mathfrak{s} , μ and \tilde{K} .

where $P(\tilde{K})$, $G(\tilde{K})$ and $P_\mu(\tilde{K})$ are the unique solutions of the following Lyapunov equations:

$$\begin{aligned} P(\tilde{K}) &= (A^c(\tilde{K}))^\top P(\tilde{K}) A^c(\tilde{K}) + \begin{pmatrix} I \\ \tilde{K} \end{pmatrix}^\top C_\dagger \begin{pmatrix} I \\ \tilde{K} \end{pmatrix}, \\ G(\tilde{K}) &= (A^c(\tilde{K}))^\top G(\tilde{K}) A^c(\tilde{K}) + \begin{pmatrix} I \\ \tilde{K} \end{pmatrix}^\top C_g \begin{pmatrix} I \\ \tilde{K} \end{pmatrix}, \\ P_\mu(\tilde{K}) &= (A^c(\tilde{K}))^\top P_\mu(\tilde{K}) A^c(\tilde{K}) + \begin{pmatrix} I \\ \tilde{K} \end{pmatrix}^\top C_\mu \begin{pmatrix} I \\ \tilde{K} \end{pmatrix}. \end{aligned} \quad (61)$$

Proof. The proof relies on the fact that for any linear extended controller \tilde{K} , let $A^c(\tilde{K})$ be the induced closed-loop matrix and $\Sigma(\tilde{K}) = A^c(\tilde{K})\Sigma(\tilde{K})A^c(\tilde{K})^\top + I$ the steady-state associated variance, for any cost matrix C ,

$$\text{Tr} \left(\begin{pmatrix} I \\ \tilde{K} \end{pmatrix}^\top C \begin{pmatrix} I \\ \tilde{K} \end{pmatrix} \Sigma(\tilde{K}) \right) = \text{Tr}(X),$$

where X is the solution of the Lyapunov equation $X = A^c(\tilde{K})^\top X A^c(\tilde{K}) + \begin{pmatrix} I \\ \tilde{K} \end{pmatrix}^\top C \begin{pmatrix} I \\ \tilde{K} \end{pmatrix}$. □

Notice that from $C_\mu = C_\dagger + \mu C_g$, standard Lyapunov algebraic manipulation ensures that

$$P_\mu(\tilde{K}) = P(\tilde{K}) + \mu G(\tilde{K}). \quad (62)$$

As a result, for any $\tilde{\pi}$ (resp. for any \tilde{K}),

$$\mathcal{L}_{\tilde{\pi}}(\mu, \mathfrak{s}) = \mathcal{J}_{\tilde{\pi}}(\mathfrak{s}) + \mu g_{\tilde{\pi}}(\mathfrak{s}) = \text{Tr}(P_\mu(\tilde{K})) = \text{Tr}(P(\tilde{K})) + \mu \text{Tr}(G(\tilde{K})), \quad (63)$$

which summarizes the three different views of the Lagrangian cost in (58), (59), and (62).

G. Proof of Lem. 9

In this section, we consider an arbitrary $\mathfrak{s} \in \mathfrak{S}$ and focus on the minimization problem

$$\tilde{\pi}_\mu(\mathfrak{s}) = \arg \min_{\tilde{\pi}} \mathcal{L}_{\tilde{\pi}}(\mu; \mathfrak{s}); \quad \mathcal{D}(\mu; \mathfrak{s}) = \min_{\tilde{\pi}} \mathcal{L}_{\tilde{\pi}}(\mu; \mathfrak{s}). \quad (64)$$

We study the conditions such that for a given Lagrangian parameter μ , the optimal policy is linear and can be obtained by solving a Riccati equation. Further, we characterize the properties of $\mathcal{D}(\mu, \mathfrak{s})$ for those *admissible* Lagrangian parameters.

G.1. Riccati characterization of the inner minimization

We first notice that for any $\mathfrak{s} \in \mathfrak{S}$ the pair (A, \tilde{B}) is controllable even if the original one (A, B) is not.¹⁴ Further, C_μ is a symmetric matrix, which can be decomposed into a more “standard” LQ form isolating the cost matrices related to the state, the extended control, and the cross terms as

$$C_\mu = \begin{pmatrix} n & n+d \\ Q_\mu & N_\mu^\top \\ N_\mu & R_\mu \end{pmatrix} \begin{matrix} n \\ n+d \end{matrix}. \quad (65)$$

As a result, it is clear from Eqs. (59) and (65) that minimizing $\mathcal{L}_{\tilde{\pi}}(\mu, \mathfrak{s})$ resembles to solving an LQR problem. However, C_μ may not be p.s.d. which would violate the standard Riccati assumptions (see Eq. 4). Nonetheless, Riccati theory extends to a more general setting, when C_μ is not p.s.d.

¹⁴This directly follows from the fact that \tilde{B} is full column rank even if B may not.

Formally, for any $\mathfrak{s} \in \mathfrak{S}$, we say that the associated Lagrangian inner minimization corresponds to an admissible Riccati solution at μ if it exists a symmetric real matrix P_μ ¹⁵ that satisfies the set of conditions:

$$\begin{aligned} R_\mu + \tilde{B}^\top P_\mu \tilde{B} &\succ 0 \\ P_\mu &= Q_\mu + A^\top P_\mu A - [A^\top P_\mu \tilde{B} + N_\mu^\top][R_\mu + \tilde{B}^\top P_\mu \tilde{B}]^{-1}[\tilde{B}^\top P_\mu A + N_\mu] \\ \rho(A_\mu^c) &< 1 \quad \text{where} \quad A_\mu^c := A - \tilde{B}[R_\mu + \tilde{B}^\top P_\mu \tilde{B}]^{-1}[\tilde{B}^\top P_\mu A + N_\mu]. \end{aligned} \quad (66)$$

The following lemma maps the existence and uniqueness of a solution satisfying Eq. (66) and the solution of the Lagrangian LQR.

Lemma 10. *For any $\mathfrak{s} \in \mathfrak{S}$, if for some Lagrangian parameter $\mu \geq 0$ it exists a symmetric matrix P_μ satisfying the set of conditions (66) defined w.r.t. \mathfrak{s} and μ , then:*

1. P_μ is unique,
2. P_μ the solution of the Lagrangian LQR (64), i.e., $\mathcal{D}(\mu, \mathfrak{s}) = \text{Tr}(P_\mu)$,
3. The optimal policy is linear in the state, i.e.,

$$\tilde{\pi}_\mu(\mathfrak{s})(x) = \tilde{K}_\mu x, \quad \text{where} \quad \tilde{K}_\mu = -[R_\mu + \tilde{B}^\top P_\mu \tilde{B}]^{-1}[\tilde{B}^\top P_\mu A + N_\mu], \quad (67)$$

4. $\mathcal{J}_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}) = \text{Tr}(P(\tilde{K}_\mu))$, $g_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}) = \text{Tr}(G(\tilde{K}_\mu))$, $\mathcal{L}_{\tilde{\pi}_\mu(\mathfrak{s})}(\mu, \mathfrak{s}) = \text{Tr}(P_\mu)$.

Proof. The first assertion directly follows from Thm. 1 in (Molinari, 1975) which ensures that if a solution to the Riccati equation exists, then it is unique. Furthermore, as it satisfies the Bellman equation, this ensures that $\mathcal{D}(\mu, \mathfrak{s}) = \text{Tr}(P_\mu)$ and that $\tilde{\pi}_\mu(\mathfrak{s})(x) = \tilde{K}_\mu x$ (see [Bertsekas, Vol.2, Prop.5.6.1]). Finally, since the optimal policy is linear, we can use Prop. 19 to obtain the Lyapunov characterization of $\mathcal{J}_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s})$, $g_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s})$, $\mathcal{L}_{\tilde{\pi}_\mu(\mathfrak{s})}(\mu, \mathfrak{s})$. \square

In essence, Lem. 10 indicates that *when* (66) admits a solution, then the inner minimization is *nice* at μ , since it corresponds to the solution of a Riccati equation, that is unique and leads to a linear optimal controller. As a result, Lem. 10 provide us with a highly implicit characterization of *good* Lagrangian parameters μ .

G.2. Domain of Admissible Riccati Solutions

In the previous section we identified the condition for which at a specific Lagrangian parameter μ , the solution of (64) is obtained by solving a Riccati equation. We now proceed with characterizing the set of \mathcal{M} in which (66) holds.

Clearly, for $\{\mu \in \mathbb{R}^+ \text{ s.t. } C_\mu \succ 0\}$, (66) holds as it coincides with the "standard" LQR setting. However, this set can be extended to a larger one, that we denote as \mathcal{M} . We first relate the existence of a Riccati solution to a certain Popov criterion. Following Molinari (1975), we introduce the Popov functions:

Definition 3. *The Popov function $\Psi_\mu(\cdot, \mathfrak{s})$ associated with $\mathfrak{s} \in \mathfrak{S}$ and Lagrangian parameter μ is defined as:*

$$\Psi_\mu(z; \mathfrak{s}) = \begin{pmatrix} (Iz^{-1} - A)^{-1} \tilde{B} \\ I \end{pmatrix}^\top C_\mu \begin{pmatrix} (Iz - A)^{-1} \tilde{B} \\ I \end{pmatrix}, \quad \forall z \in \mathbb{C}, |z| = 1. \quad (68)$$

For any controller \tilde{K} , the controlled Popov function $\Psi_\mu^{\tilde{K}}(\cdot, \mathfrak{s})$ associated with $\mathfrak{s} \in \mathfrak{S}$ and Lagrangian parameter μ is defined as:

$$\Psi_\mu^{\tilde{K}}(z; \mathfrak{s}) = \begin{pmatrix} (Iz^{-1} - A^c(\tilde{K}))^{-1} \tilde{B} \\ I \end{pmatrix}^\top \begin{pmatrix} I & \tilde{K}^\top \\ 0 & I \end{pmatrix} C_\mu \begin{pmatrix} I & 0 \\ \tilde{K} & I \end{pmatrix} \begin{pmatrix} (Iz - A^c(\tilde{K}))^{-1} \tilde{B} \\ I \end{pmatrix} \quad (69)$$

In essence, the controlled Popov function by \tilde{K} simply consists in the Popov function when \mathfrak{s} is pre-controlled by \tilde{K} . The link between those two is given by

$$\Psi_\mu^{\tilde{K}}(z; \mathfrak{s}) = Y^{\tilde{K}}(z^{-1}; \mathfrak{s})^\top \Psi_\mu(z; \mathfrak{s}) Y^{\tilde{K}}(z; \mathfrak{s}) \quad \text{where} \quad Y^{\tilde{K}}(z; \mathfrak{s}) = I + \tilde{K}[Iz - A^c(\tilde{K})]^{-1} \tilde{B}. \quad (70)$$

¹⁵We drop the dependency in \mathfrak{s} for sake of readability.

The interest of considering Popov functions rather than solution of (66) is that the latter are well defined for all $\mu \geq 0$ while the former may not. Further, one can map the existence of a solution to (66) to the positive definiteness of the Popov functions. Those results are provided in (Molinari, 1975).

Proposition 20 (Thm.1&2 Molinari (1975)). *For any controllable system \mathfrak{s} , the following conditions are equivalent:*

1. *There exists a real symmetric solution satisfying Eq. (66), necessarily unique.*
2. *For some (and hence all) \tilde{K} such that $|\lambda(A_\mu^c(\tilde{K}))| \neq 1$, $\Psi_\mu^{\tilde{K}}(z; \mathfrak{s}) \succ 0$ on the unit circle $|z| = 1$.*

We also recall two important identities can be extracted from Molinari (1975) that links the Popov function to the optimal quantities $A_\mu^c, D_\mu, \tilde{K}_\mu$ whenever they are defined:

$$\begin{aligned} \forall |z| = 1, \quad \Psi_\mu^{\tilde{K}}(z; \mathfrak{s}) &= D_\mu; \\ \forall |z| = 1, \quad \forall \tilde{K}_1, \tilde{K}_2 \text{ s.t. } |\lambda(A^c(\tilde{K}_1))| \neq 1, |\lambda(A^c(\tilde{K}_2))| \neq 1, \Psi_\mu^{\tilde{K}_1}(z; \mathfrak{s}) &= Y^{\tilde{K}_1, \tilde{K}_2}(z^{-1}; \mathfrak{s})^\top \Psi_\mu^{\tilde{K}_2}(z; \mathfrak{s}) Y^{\tilde{K}_1, \tilde{K}_2}(z; \mathfrak{s}), \\ \text{where } Y^{\tilde{K}_1, \tilde{K}_2}(z; \mathfrak{s}) &= I + (\tilde{K}_1 - \tilde{K}_2)[Iz - A^c(\tilde{K}_1)]^{-1} \tilde{B}. \end{aligned} \quad (71)$$

Thanks to Prop. 20, we can now associate to any $\mathfrak{s} \in \mathfrak{S}$ a set \mathcal{M} such that for all $\mu \in \mathcal{M}$, a symmetric solution to (66) exists and is strictly stabilizing, which, by Lem. 10 implies that the optimal solution to (64) is a linear policy.

Lemma 11. *For any $\mathfrak{s} \in \mathfrak{S}$, let $\Psi_\mu^{\tilde{K}}(z, \mathfrak{s})$ be defined by Def. 3. Let*

$$\tilde{\mu} := \sup \{ \mu \geq 0 \text{ s.t. } \Psi_\mu^{\tilde{K}}(z, \mathfrak{s}) \succ 0, \text{ for all } |z| = 1 \text{ and } \tilde{K} \text{ such that } |\lambda(A_\mu^c(\tilde{K}))| \neq 1 \}. \quad (72)$$

Then, the set $\mathcal{M} = [0, \tilde{\mu}]$ is non-empty and for all $\mu \in \mathcal{M}$, Eq. (66) admits a solution such that $\rho(A_\mu^c) < 1$.

Proof. 1. For $\mu = 0$,

$$C_0 = C_\dagger + 0 \times C_g = \begin{pmatrix} C_\dagger^1 & 0 \\ 0 & 0 \end{pmatrix} + C_\dagger^2.$$

Let $\bar{K} = \begin{pmatrix} 0 \\ -A \end{pmatrix}$ which is such that $A^c(\bar{K}) = 0$. From $C_\dagger^1 \succ 0$ and $C_\dagger^2 \succcurlyeq 0$, one has for all $|z| = 1$,

$$\begin{aligned} \Psi_0^{\bar{K}}(z, \mathfrak{s}) &= \left((Iz^{-1} - A^c(\bar{K}))^{-1} \tilde{B} \right)^\top \begin{pmatrix} I & \bar{K}^\top \\ 0 & I \end{pmatrix} C_\dagger \begin{pmatrix} I & 0 \\ \bar{K} & I \end{pmatrix} \left((Iz - A^c(\bar{K}))^{-1} \tilde{B} \right) \\ &= \begin{pmatrix} z\tilde{B} \\ I \end{pmatrix}^\top \begin{pmatrix} I & \bar{K}^\top \\ 0 & I \end{pmatrix} C_\dagger \begin{pmatrix} I & 0 \\ \bar{K} & I \end{pmatrix} \begin{pmatrix} z^{-1}\tilde{B} \\ I \end{pmatrix} \\ &\succcurlyeq \begin{pmatrix} z\tilde{B} \\ I \end{pmatrix}^\top \begin{pmatrix} I & \bar{K}^\top \\ 0 & I \end{pmatrix} \begin{pmatrix} C_\dagger^1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} I & 0 \\ \bar{K} & I \end{pmatrix} \begin{pmatrix} z^{-1}\tilde{B} \\ I \end{pmatrix} \\ &\succcurlyeq \begin{pmatrix} zB & I_n \\ I_d & 0 \end{pmatrix}^\top C_\dagger^1 \begin{pmatrix} z^{-1}B & I_n \\ I_d & 0 \end{pmatrix} \\ &\succ 0 \end{aligned}$$

where we used that $\left| \det \begin{pmatrix} zB & I_n \\ I_d & 0 \end{pmatrix} \right| = \left| \det \begin{pmatrix} I_n & z\tilde{B} \\ 0 & I_d \end{pmatrix} \right| = 1$ and that $C_\dagger^1 \succ 0$ to obtain the last inequality. Since we exhibit a \bar{K} such that $\Psi_0^{\bar{K}}(z, \mathfrak{s}) \succ 0$ for all $|z| = 1$, Prop. 20 ensures that it holds for all \tilde{K} such that $|\lambda(A^c(\tilde{K}))| \neq 1$ and thus that $0 \in \mathcal{M}$ and \mathcal{M} is non-empty.

2. To show that all $\mu \in \mathcal{M}$ are associated with a strictly stabilizing solution of (66), we will invoke Prop. 20. As a result, we are left to prove that for all $\mu \in \mathcal{M}$, it exists \tilde{K} such that $|\lambda(A_\mu^c(\tilde{K}))| \neq 1$ and $\Psi_\mu^{\tilde{K}}(z, \mathfrak{s}) \succ 0$ for all $|z| = 1$.

Exploiting the linearity of Ψ_μ (as C_μ is linear in μ , for all $|z| = 1$ and \tilde{K}), so is $\Psi_\mu^{\tilde{K}}(z, \mathfrak{s})$. Hence, for all $|z| = 1$, for all $v \in \mathbb{C}^{n+d} \setminus \{0\}$, the function

$$f : \mu \rightarrow v^* \Psi_\mu^{\tilde{K}}(z, \mathfrak{s}) v,$$

is linear in μ . Moreover, $\lim_{\mu \rightarrow \tilde{\mu}} f(\mu) \geq 0$ and $f(0) > 0$, which implies that for all $\mu \in \mathcal{M}$, for all $|z| = 1$,

$$\forall v \in \mathbb{C}^{n+d} \setminus \{0\}, \quad v^* \Psi_\mu^{\tilde{K}}(z, \mathfrak{s}) v > 0 \quad \Rightarrow \quad \Psi_\mu^{\tilde{K}}(z, \mathfrak{s}) \succ 0.$$

□

G.3. Characterization of $\mathcal{D}(\mu, \mathfrak{s})$

We are now ready to characterize the dual function \mathcal{D} . For any $\mathfrak{s} \in \mathfrak{S}$, the use of the Popov functions allows us to define a set \mathcal{M} (Lem. 11) that depends explicitly (yet in a non-trivial fashion) in \mathfrak{s} , and on which the inner minimization problem (64) corresponds to a well-defined (yet non-standard) Riccati equation. We summarize lemmas 10 and 11 in Prop. 21 which proves the first statements of Lem. 9.

Proposition 21. *For any $\mathfrak{s} \in \mathfrak{S}$, let \mathcal{M} be the admissible Riccati set associated with \mathfrak{s} as defined in (72) and $\mathcal{D}(\mu, \mathfrak{s})$ the dual function defined in (64). Then,*

1. for all $\mu \in \mathcal{M}$, it exists a unique symmetric matrix P_μ satisfying (66),
2. the optimal policy $\tilde{\pi}_\mu(\mathfrak{s}) = \arg \min_{\tilde{\pi}} \mathcal{L}_{\tilde{\pi}}(\mu; \mathfrak{s})$ is linear in the state, i.e., $\tilde{\pi}_\mu(\mathfrak{s})(x) = \tilde{K}_\mu x$, where

$$\tilde{K}_\mu = -[R_\mu + \tilde{B}^\top P_\mu \tilde{B}]^{-1} [\tilde{B}^\top P_\mu A + N_\mu],$$

3. $\mathcal{J}_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}) = \text{Tr}(P(\tilde{K}_\mu))$, $g_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}) = \text{Tr}(G(\tilde{K}_\mu))$, $\mathcal{D}(\mu, \mathfrak{s}) = \mathcal{L}_{\tilde{\pi}_\mu(\mathfrak{s})}(\mu, \mathfrak{s}) = \text{Tr}(P_\mu)$.

Proof. The proof of Prop. 21 directly follows from merging Lem. 10 and Lem. 11. □

Now that Prop. 21 guarantees that for any $\mathfrak{s} \in \mathfrak{S}$, $\mathcal{D}(\cdot, \mathfrak{s})$ is well defined on some explicit convex set \mathcal{M} , we are able to prove the last statements of Lem. 9.

Proposition 22. *For any $\mathfrak{s} \in \mathfrak{S}$, let \mathcal{M} be the admissible Riccati set associated with \mathfrak{s} as defined in (72). Let \mathcal{D}' denote the derivative of \mathcal{D} w.r.t. μ . Then,*

1. \mathcal{M} is a non-empty convex set,
2. $\mu \rightarrow \mathcal{D}(\mu, \mathfrak{s})$ is concave on \mathcal{M} and $\mathcal{D}(\mu, \mathfrak{s}) > -\infty$ on \mathcal{M} .
3. $\mu \rightarrow \mathcal{D}(\mu, \mathfrak{s}) \in \mathcal{C}^1(\mathcal{M})$.
4. $\mathcal{D}'(\mu, \mathfrak{s}) = g_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}) = \text{Tr}(G(\tilde{K}_\mu))$, where \tilde{K}_μ and $\tilde{\pi}_\mu(\mathfrak{s})(x) = \tilde{K}_\mu x$ be respectively the optimal controller and policy as defined in Prop. 21 and $G(\tilde{K}_\mu)$ is given in (61). As a result,

$$\mathcal{D}(\mu, \mathfrak{s}) = \mathcal{J}_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}) + \mu g_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}) = \mathcal{J}_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}) + \mu \mathcal{D}'(\mu, \mathfrak{s}). \quad (73)$$

Proof. 1. \mathcal{M} is non-empty from Lem. 11. Further, $\mathcal{M} = [0, \tilde{\mu})$ where $\tilde{\mu}$ is defined in (72). Thus, \mathcal{M} is convex.

2. The concavity of $\mathcal{D}(\cdot, \mathfrak{s})$ is provided by construction, as \mathcal{D} is the minimum of affine function. The lower boundedness property of $\mathcal{D}(\cdot, \mathfrak{s})$ follows from $D_\mu = R_\mu + \tilde{B}^\top P_\mu \tilde{B} \succ 0$ for all $\mu \in \mathcal{M}$. Since $R_\mu \prec +\infty$ and \tilde{B} is full column rank,

$$D_\mu \succ 0 \quad \Rightarrow \quad \tilde{B}^\top P_\mu \tilde{B} \succ -\infty \quad \Rightarrow \quad P_\mu \succ -\infty \quad \Rightarrow \quad \mathcal{D}(\mu, \mathfrak{s}) = \text{Tr}(P_\mu) > -\infty.$$

3. The smoothness of $\mathcal{D}(\cdot, \mathfrak{s})$ follows from Alexandrov's Theorem that guarantees that $\mathcal{D}(\cdot, \mathfrak{s})$ is twice differentiable on \mathcal{M} and hence in $\mathcal{C}^1(\mathcal{M})$.

4. $g_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}) = \text{Tr}(G(\tilde{K}_\mu))$ by Prop. 21. We are thus left to prove that $\mathcal{D}'(\mu, \mathfrak{s}) = g_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s})$, which is a structural property inherited from the Lagrangian formulation and that is directly implied by the fact that $\mathcal{D}(\mu, \mathfrak{s}) > -\infty$ for all $\mu \in \mathcal{M}$ and that $\mathcal{D}(\cdot, \mathfrak{s})$ is smooth and concave. Observe that by construction, for any $(\mu, \bar{\mu}) \in \mathcal{M}^2$,

$$\mathcal{D}(\bar{\mu}, \mathfrak{s}) = \mathcal{J}_{\tilde{\pi}_{\bar{\mu}}(\mathfrak{s})}(\mathfrak{s}) + \bar{\mu}g_{\tilde{\pi}_{\bar{\mu}}(\mathfrak{s})}(\mathfrak{s}) = \min_{\tilde{\pi}} (\mathcal{J}_{\tilde{\pi}}(\mathfrak{s}) + \bar{\mu}g_{\tilde{\pi}}(\mathfrak{s})) \geq \mathcal{J}_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}) + \bar{\mu}g_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}) = \mathcal{D}(\mu, \mathfrak{s}) + (\bar{\mu} - \mu)g_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}).$$

Thus, $\mathcal{D}(\bar{\mu}, \mathfrak{s}) \geq \mathcal{D}(\mu, \mathfrak{s}) + (\bar{\mu} - \mu)g_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s})$ which means by definition that $g_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s})$ belongs to the super-gradient set at μ (which is defined for concave function that are not $-\infty$, see (Rockafellar, 1970) for a survey). However, since $\mathcal{D}(\cdot, \mathfrak{s}) \in \mathcal{C}^1(\mathcal{M})$, the super-gradient set is reduced to a singleton, which coincides with the derivative. As a result, $\mathcal{D}'(\mu, \mathfrak{s}) = g_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s})$.

□

H. Proof of Thm. 4

We are now ready to proceed with the proof of Thm. 4 and aim at proving strong duality, i.e., that for any $\mathfrak{s} \in \mathfrak{S}$ associated with admissible Riccati set \mathcal{M} ,

$$\mathcal{J}_*(\mathfrak{s}) = \sup_{\mu \in \mathcal{M}} \mathcal{D}(\mu; \mathfrak{s}) = \sup_{\mu \in \mathcal{M}} \min_{\tilde{\pi}} \mathcal{L}_{\tilde{\pi}}(\mu; \mathfrak{s}).$$

Notice that the definition of \mathfrak{S} guarantees that it exists a feasible stable linear policy for \mathfrak{s} , and hence that

$$\mathcal{J}_*(\mathfrak{s}) < \infty. \quad (74)$$

Further, from weak-duality,

$$\mathcal{J}_*(\mathfrak{s}) \geq \sup_{\mu \in \mathcal{M}} \mathcal{D}(\mu; \mathfrak{s}), \quad (75)$$

thus the proof consists in showing the converse inequality.

The reason why we write the strong-duality in a sup – min sense is that strong duality might hold on the frontier of the open set \mathcal{M} . More in details, we aim to show that for any $\mathfrak{s} \in \mathfrak{S}$, for any arbitrary $\epsilon > 0$,

$$\mathcal{J}_*(\mathfrak{s}) \geq \sup_{\mu \in \mathcal{M}} \mathcal{D}(\mu; \mathfrak{s}) = \sup_{\mu \in \mathcal{M}} \min_{\tilde{\pi}} \mathcal{L}_{\tilde{\pi}}(\mu; \mathfrak{s}) \geq \mathcal{J}_*(\mathfrak{s}) - \epsilon.$$

Since the l.h.s. inequality is always true by weak-duality, we focus on proving the r.h.s. inequality.

The proof distinguishes between two distinct cases, based on the monotonicity of the dual function \mathcal{D} . Notice that Lem. 9 ensures that \mathcal{D} is a concave function over the real convex set \mathcal{M} . As a result, it can be either **1)** increasing then non-increasing or non-increasing everywhere **2)** (strictly) increasing everywhere on \mathcal{M} .

In case **1)**, we show that strong duality holds in a “strict” sense¹⁶, that is it exists $\mu_* \in \mathcal{M}$ such that

$$\begin{aligned} \text{strong duality} \quad & \mathcal{J}_*(\mathfrak{s}) = \max_{\mu \in \mathcal{M}} \mathcal{D}(\mu; \mathfrak{s}) = \mathcal{D}(\mu_*; \mathfrak{s}) \\ \text{primal feasibility} \quad & g_{\tilde{\pi}_{\mu_*}(\mathfrak{s})} \leq 0 \\ \text{complementary slackness} \quad & \mu_* g_{\tilde{\pi}_{\mu_*}(\mathfrak{s})} = 0. \end{aligned} \quad (76)$$

In case **2)**, we show that strong duality holds in a “weak” sense¹⁷, that is that

$$\mathcal{J}_*(\mathfrak{s}) = \sup_{\mu \in \mathcal{M}} \mathcal{D}(\mu; \mathfrak{s}). \quad (77)$$

More in detail, we show that for any arbitrary $\epsilon > 0$, it exists $\mu^\epsilon \in \mathcal{M}$ such that $\mathcal{D}(\mu^\epsilon, \mathfrak{s}) \geq \mathcal{J}_*(\mathfrak{s}) - \epsilon$. As a by-product, we provide an explicit linear policy $\tilde{\pi}^\epsilon(x) = \tilde{K}^\epsilon x$ which satisfies the primal feasibility i.e., $g_{\tilde{\pi}^\epsilon}(\mathfrak{s}) \leq 0$ and is ϵ -optimal i.e., $\mathcal{J}_{\tilde{\pi}^\epsilon}(\mathfrak{s}) - \mathcal{J}_*(\mathfrak{s}) \leq \epsilon$, either in closed-form, either as the solution of a well-defined Riccati equation.

¹⁶We say that strong duality holds in a strict sense when the sup – min is replaced by a max – min

¹⁷We say that strong duality holds in a weak sense when max is not attain on the domain, and hence has to be replaced by a sup.

H.1. Case 1). The dual function $\mathcal{D}(\cdot, \mathfrak{s})$ is non-increasing on a subset of \mathcal{M}

By Lem. 9, $\mathcal{D}(\cdot, \mathfrak{s})$ is concave and $\mathcal{C}^1(\mathcal{M})$. As a result, $\mathcal{D}(\cdot, \mathfrak{s})$ is non-increasing on a subset of \mathcal{M} if and only if it exists $\mu_- \in \mathcal{M}$ such that $\mathcal{D}'(\mu_-, \mathfrak{s}) \leq 0$. In this case, we have the following result.

Lemma 12. *For any $\mathfrak{s} \in \mathfrak{S}$, let \mathcal{M} be the admissible Riccati set associated with \mathfrak{s} as defined in (72). If it exists $\mu_- \in \mathcal{M}$ such that $\mathcal{D}'(\mu_-, \mathfrak{s}) \leq 0$, then, $\mathcal{D}(\cdot, \mathfrak{s})$ attains its maximum on \mathcal{M} and it exists $\mu_* \in \arg \max_{\mu \in \mathcal{M}} \mathcal{D}(\mu, \mathfrak{s})$ such that*

1. $\mathcal{J}_*(\mathfrak{s}) = \max_{\mu \in \mathcal{M}} \mathcal{D}(\mu; \mathfrak{s}) = \mathcal{D}(\mu_*; \mathfrak{s})$,
2. $g_{\tilde{\pi}_{\mu_*}(\mathfrak{s})} \leq 0$
3. $\mu_* g_{\tilde{\pi}_{\mu_*}(\mathfrak{s})} = 0$.

Proof. The proof of Lem. 12 follows from Lem. 9. Since it exists $\mu_- \in \mathcal{M}$ such that $\mathcal{D}'(\mu_-, \mathfrak{s}) \leq 0$, by concavity, $\mathcal{D}(\cdot, \mathfrak{s})$ attains its maximum on \mathcal{M} . Let μ_* be the smallest Lagrangian parameter in $\{\arg \max_{\mu \in \mathcal{M}} \mathcal{D}(\mu, \mathfrak{s})\}$. Then,

$$\text{Either } \mu_* = 0 \text{ and } \mathcal{D}'(\mu_*, \mathfrak{s}) \leq 0; \quad \text{Either } \mu_* > 0 \text{ and } \mathcal{D}'(\mu_*, \mathfrak{s}) = 0. \quad (78)$$

Since $g_{\tilde{\pi}_{\mu}(\mathfrak{s})}(\mathfrak{s}) = \mathcal{D}'(\mu, \mathfrak{s})$ for all $\mu \in \mathcal{M}$, (78) translates into

$$\text{Either } \mu_* = 0 \text{ and } g_{\tilde{\pi}_{\mu_*}(\mathfrak{s})}(\mathfrak{s}) \leq 0; \quad \text{Either } \mu_* > 0 \text{ and } g_{\tilde{\pi}_{\mu_*}(\mathfrak{s})}(\mathfrak{s}) = 0. \quad (79)$$

In both case, we obtain the primal feasibility and complementary slackness.

We conclude the proof noting that $\mathcal{J}_*(\mathfrak{s}) \geq \mathcal{D}(\mu_*, \mathfrak{s})$ from weak duality, that $\mathcal{D}(\mu_*, \mathfrak{s}) = \mathcal{L}_{\tilde{\pi}_{\mu_*}(\mathfrak{s})}(\mu_*, \mathfrak{s}) = \mathcal{J}_{\tilde{\pi}_{\mu_*}(\mathfrak{s})}(\mathfrak{s}) + \mu_* g_{\tilde{\pi}_{\mu_*}(\mathfrak{s})}(\mathfrak{s}) = \mathcal{J}_{\tilde{\pi}_{\mu_*}(\mathfrak{s})}(\mathfrak{s})$ from complementary slackness, and that $\mathcal{J}_{\tilde{\pi}_{\mu_*}(\mathfrak{s})}(\mathfrak{s}) \geq \mathcal{J}_*(\mathfrak{s})$ from primal feasibility. Hence, $\mathcal{J}_*(\mathfrak{s}) = \mathcal{D}(\mu_*, \mathfrak{s})$. \square

H.2. Case 2). The dual function $\mathcal{D}(\cdot, \mathfrak{s})$ is strictly increasing on \mathcal{M}

When $\mathcal{D}(\cdot, \mathfrak{s})$ is strictly increasing on \mathcal{M} , we have

$$\forall \mu \in \mathcal{M}, \quad \mathcal{D}'(\mu, \mathfrak{s}) > 0, \quad (80)$$

which implies that the maximum is not attained on \mathcal{M} , and prevents us from using a similar path as in App. H.1 but indicates that we should study the behavior of $\mathcal{D}(\mu, \mathfrak{s})$ when $\mu \rightarrow \tilde{\mu}$ that is the closure of \mathcal{M} .

We first stress in App. H.2.1 the implications of (80) in term of boundedness of \mathcal{M} and in term of stability of the optimal policy $\tilde{\pi}_{\mu}(\mathfrak{s})$. In light of those properties, we discuss in App. H.2.2 the behavior of $\mathcal{D}(\mu, \mathfrak{s})$ when μ approaches $\tilde{\mu}$ and show that there exists suboptimal linear policies $\tilde{\pi} \neq \tilde{\pi}_{\mu}(\mathfrak{s})$ that yield almost the same performance, i.e., such that $\mathcal{L}_{\tilde{\pi}}(\mu, \mathfrak{s}) \simeq \mathcal{D}(\mu, \mathfrak{s})$. Finally, in App. H.2.3 we provide an explicit feasible policy $\tilde{\pi}^\epsilon$ whose performance can be made arbitrarily close to the optimal one, i.e., such that $\mathcal{L}_{\tilde{\pi}^\epsilon}(\mu^\epsilon, \mathfrak{s}) - \mathcal{D}(\mu^\epsilon, \mathfrak{s}) \leq \epsilon$ for some $\mu^\epsilon \in \mathcal{M}$ which implies that $\mathcal{J}_{\tilde{\pi}^\epsilon}(\mathfrak{s}) - \mathcal{J}_*(\mathfrak{s}) \leq \epsilon$.

H.2.1. STABILITY AND BOUNDEDNESS PROPERTIES

The condition (80) provides us with information both on the value of $\tilde{\mu}$ and on the stability of $\tilde{\pi}_{\mu}(\mathfrak{s})$ over \mathcal{M} . Those are formalized by in the following statements.

Proposition 23. *For any $\mathfrak{s} \in \mathfrak{S}$, let $\mathcal{M} = [0, \tilde{\mu})$ be the admissible Riccati set associated with \mathfrak{s} as defined in (72). If for all $\mu \in \mathcal{M}$, $\mathcal{D}'(\mu, \mathfrak{s}) > 0$, then,*

$$\tilde{\mu} \leq \lambda_{\max} \left(\begin{pmatrix} I_{n+d} & 0 \\ 0 & 0 \end{pmatrix} C_{\dagger} \begin{pmatrix} I_{n+d} & 0 \\ 0 & 0 \end{pmatrix} \right) / \lambda_{\min}(V^{-1}). \quad (81)$$

Proposition 24. *For any $\mathfrak{s} \in \mathfrak{S}$, let \mathcal{M} be the admissible Riccati set associated with \mathfrak{s} as defined in (72). Let \tilde{K}_{μ} , $A_{\mu}^c = A + \tilde{B}\tilde{K}_{\mu}$ and $\Sigma_{\mu} = A_{\mu}^c \Sigma_{\mu} (A_{\mu}^c)^{\top} + I$ be the optimal controller, closed-loop matrix and steady-state covariance matrix at $\mu \in \mathcal{M}$ respectively. Then, for all $\mu \in \mathcal{M}$,*

$$\mathcal{D}'(\mu, \mathfrak{s}) \geq 0 \quad \Rightarrow \quad \text{Tr}(\Sigma_{\mu}) \leq \kappa(\mu, \mathfrak{s}) \leq \kappa(\mathfrak{s}) \quad \Rightarrow \quad \rho(A_{\mu}^c)^2 \leq 1 - 1/\kappa(\mu, \mathfrak{s}) \leq 1 - 1/\kappa(\mathfrak{s}), \quad (82)$$

where $1 \leq \kappa(\mu, \mathfrak{s}) := \mathcal{D}(\mu, \mathfrak{s}) / \lambda_{\min}(C_{\dagger}^1)$ and $1 \leq \kappa(\mathfrak{s}) := \mathcal{J}_*(\mathfrak{s}) / \lambda_{\min}(C_{\dagger}^1)$.

The next proposition follows directly as a corollary from (82).

Proposition 25. For any $\mathfrak{s} \in \mathfrak{S}$, let \mathcal{M} be the admissible Riccati set associated with \mathfrak{s} as defined in (72). Let \tilde{K}_μ , $A_\mu^c = A + \tilde{B}\tilde{K}_\mu$ and $\Sigma_\mu = A_\mu^c \Sigma_\mu (A_\mu^c)^\top + I$ be the optimal controller, closed-loop matrix and steady-state covariance matrix at $\mu \in \mathcal{M}$ respectively. Then, for all $\mu \in \mathcal{M}$ such that $\mathcal{D}'(\mu, \mathfrak{s}) \geq 0$,

$$\mathcal{D}'(\mu, \mathfrak{s}) \geq 0 \Rightarrow \begin{cases} \forall k \in \mathbb{N}, \left\| (A_\mu^c)^k \right\|_2 \leq \sqrt{\kappa(\mu, \mathfrak{s})} (1 - 1/\kappa(\mu, \mathfrak{s}))^{-k} \leq \sqrt{\kappa(\mathfrak{s})} (1 - 1/\kappa(\mathfrak{s}))^{-k} \\ \left\| \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix} \right\|_2 \leq \sqrt{\kappa(\mu, \mathfrak{s})} (2 + \|A\|_2 \|B\|_2) \leq \sqrt{\kappa(\mathfrak{s})} (2 + \|A\|_2 \|B\|_2) \end{cases} \quad (83)$$

where $1 \leq \kappa(\mu, \mathfrak{s}) := \mathcal{D}(\mu, \mathfrak{s})/\lambda_{\min}(C_\dagger^1)$ and $1 \leq \kappa(\mathfrak{s}) := \mathcal{J}_*(\mathfrak{s})/\lambda_{\min}(C_\dagger^1)$.

Proof of Prop. 23. The proof is done by contradiction. Suppose that (81) do not hold. Then, it exists $\mu \in \mathcal{M}$ such that

$$\mu > \lambda_{\max} \left(\begin{pmatrix} I_{n+d} & 0 \\ 0 & 0 \end{pmatrix} C_\dagger \begin{pmatrix} I_{n+d} & 0 \\ 0 & 0 \end{pmatrix} \right) / \lambda_{\min}(V^{-1}).$$

As a result, from $C_\mu = C_\dagger + \mu C_g$ we obtain

$$\begin{pmatrix} I_{n+d} & 0 \\ 0 & 0 \end{pmatrix} C_\mu \begin{pmatrix} I_{n+d} & 0 \\ 0 & 0 \end{pmatrix} \prec \lambda_{\max} \left(\begin{pmatrix} I_{n+d} & 0 \\ 0 & 0 \end{pmatrix} C_\dagger \begin{pmatrix} I_{n+d} & 0 \\ 0 & 0 \end{pmatrix} \right) I - \mu \lambda_{\min}(V^{-1}) I \prec 0.$$

Hence, for all linear policy $\tilde{\pi}(x) = \tilde{K}x$ of the form $\tilde{K} = \begin{pmatrix} K_u \\ 0 \end{pmatrix}$

$$\begin{pmatrix} I \\ \tilde{K} \end{pmatrix}^\top C_\mu \begin{pmatrix} I \\ \tilde{K} \end{pmatrix} = \begin{pmatrix} I \\ K_u \\ 0 \end{pmatrix}^\top C_\mu \begin{pmatrix} I \\ K_u \\ 0 \end{pmatrix} = \begin{pmatrix} I \\ K_u \end{pmatrix}^\top \begin{pmatrix} I_{n+d} & 0 \\ 0 & 0 \end{pmatrix} C_\mu \begin{pmatrix} I_{n+d} & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} I \\ K_u \end{pmatrix} \prec 0.$$

Thus, for such $\tilde{\pi}$,

$$\mathcal{L}_{\tilde{\pi}}(\mu; \mathfrak{s}) = \lim_{S \rightarrow \infty} \frac{1}{S} \mathbb{E} \left[\sum_{s=0}^{S-1} \begin{pmatrix} x_s^\top & \tilde{u}_s^\top \end{pmatrix} C_\mu \begin{pmatrix} x_s \\ \tilde{u}_s \end{pmatrix} \right] \leq \lim_{S \rightarrow \infty} \frac{1}{S} \mathbb{E} \left[\sum_{s=0}^{S-1} x_s^\top \begin{pmatrix} I \\ K_u \end{pmatrix}^\top \begin{pmatrix} I_{n+d} & 0 \\ 0 & 0 \end{pmatrix} C_\mu \begin{pmatrix} I_{n+d} & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} I \\ K_u \end{pmatrix} x_s \right] < 0$$

subject to $x_{s+1} = A^c(\tilde{K})x_s + \epsilon_{s+1}$.

As a result,

- either $A^c(\tilde{K})$ is not stable and $\mathcal{D}(\mu, \mathfrak{s}) \leq \mathcal{L}_{\tilde{\pi}}(\mu; \mathfrak{s}) = -\infty$, that contradicts the second statement in Prop. 22,
- either $A^c(\tilde{K})$ is stable and $-\infty < \mathcal{D}(\mu, \mathfrak{s}) < 0$. However, from $C_0 \succ 0$, we know that $\mathcal{D}(0, \mathfrak{s}) \geq 0$, which contradicts the fact that $\mathcal{D}(\cdot, \mathfrak{s})$ is (strictly) increasing on \mathcal{M} .

□

Proof of Prop. 24. Let $\mu \in \mathcal{M}$ such that $\mathcal{D}'(\mu, \mathfrak{s}) \geq 0$. Then, we have

$$\mathcal{J}_*(\mathfrak{s}) \geq \mathcal{D}(\mu, \mathfrak{s}) = \mathcal{J}_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}) + \mu g_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}) = \mathcal{J}_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}) + \mu \mathcal{D}'(\mu, \mathfrak{s}) \geq \mathcal{J}_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}),$$

where the l.h.s. inequality follows from weak duality while the equalities follows from (73). Moreover, let Σ_μ be the steady-state covariance matrix defined by

$$\Sigma_\mu = A_\mu^c \Sigma_\mu (A_\mu^c)^\top + I,$$

we have by Eqs. (60) and (61),

$$\mathcal{J}_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s}) = \text{Tr}(P(\tilde{K}_\mu)) \geq \text{Tr}(\Sigma_\mu) \lambda_{\min}(C_\dagger^1).$$

As a result,

$$\frac{1}{1 - \rho(A_\mu^c)^2} = \|\Sigma_\mu\|_2 \leq \text{Tr}(\Sigma_\mu) \leq \mathcal{D}(\mu, \mathfrak{s}) / \lambda_{\min}(C_\dagger^1) := \kappa(\mu, \mathfrak{s}).$$

Finally, weak duality ensures that $\kappa(\mu, \mathfrak{s}) \leq \kappa(\mathfrak{s}) := \mathcal{J}_*(\mathfrak{s}) / \lambda_{\min}(C_\dagger^1)$.

□

Proof of Prop. 25. 1. The first assertion follows from propositions 6 and 24. Since by Prop. 24, $\kappa(\mu, \mathfrak{s}) \succcurlyeq \Sigma_\mu \succcurlyeq I$, and $A_\mu^c \Sigma_\mu (A_\mu^c)^\top - \Sigma_\mu = -I$, Prop. 6 ensures that for any $k \in \mathbb{N}$,

$$\left\| (A_\mu^c)^k \right\|_2 \leq \sqrt{\kappa(\mu, \mathfrak{s})} (1 - 1/\kappa(\mu, \mathfrak{s}))^{-k}.$$

2. Decomposing $\tilde{K}_\mu = \begin{pmatrix} K_\mu^u \\ K_\mu^w \end{pmatrix}$, where $K_\mu^u \in \mathbb{R}^{d,n}$, $K_\mu^w \in \mathbb{R}^{n,n}$ we have that

$$\begin{aligned} \begin{pmatrix} I \\ \tilde{K}_\mu^u \end{pmatrix}^\top \begin{pmatrix} I \\ \tilde{K}_\mu^u \end{pmatrix} &= \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix}^\top \begin{pmatrix} I_{n+d} & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix} \preccurlyeq \frac{1}{\lambda_{\min}(C_\dagger^1)} \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix}^\top C_\dagger^1 \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix} \preccurlyeq \frac{1}{\lambda_{\min}(C_\dagger^1)} \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix}^\top C_\dagger \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix} \\ &\preccurlyeq \frac{1}{\lambda_{\min}(C_\dagger^1)} P(\tilde{K}_\mu). \end{aligned}$$

Hence, $\left\| \begin{pmatrix} I \\ \tilde{K}_\mu^u \end{pmatrix} \right\|_2^2 \leq \frac{\text{Tr}(P(\tilde{K}_\mu))}{\lambda_{\min}(C_\dagger^1)} \leq \kappa(\mathfrak{s}, \mu)$. Further, we have $K_\mu^w = A_\mu^c - (A \ B) \begin{pmatrix} I \\ \tilde{K}_\mu^u \end{pmatrix}$, we obtain

$$\|K_\mu^w\|_2 \leq \|A_\mu^c\|_2 + \|A\|_2 \|B\|_2 \left\| \begin{pmatrix} I \\ \tilde{K}_\mu^u \end{pmatrix} \right\|_2 \leq \sqrt{\kappa(\mu, \mathfrak{s})} (1 + \|A\|_2 \|B\|_2).$$

As a result, since $\|A_\mu^c\|_2 \leq \sqrt{\kappa(\mu, \mathfrak{s})}$, we obtain $\left\| \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix} \right\|_2 \leq \left\| \begin{pmatrix} I \\ \tilde{K}_\mu^u \end{pmatrix} \right\|_2 + \|K_\mu^w\|_2 \leq \sqrt{\kappa(\mu, \mathfrak{s})} (2 + \|A\|_2 \|B\|_2)$. \square

H.2.2. CLOSURE OF \mathcal{M}

From condition (80), it is clear that all optimal policies $\{\tilde{\pi}_\mu(\mathfrak{s})\}_{\mu \in \mathcal{M}}$ will violate the constraint since it is equivalent to $g_{\tilde{\pi}_\mu}(\mathfrak{s}) > 0$ for all $\mu \in \mathcal{M}$. However, we show here that when $\tilde{\mu} - \mu$ is small, there exist other linear policies $\tilde{\pi}$ whose constraint value $g_{\tilde{\pi}}(\mathfrak{s})$ may be very different than $g_{\tilde{\pi}_\mu(\mathfrak{s})}(\mathfrak{s})$ but such that $\mathcal{L}_{\tilde{\pi}}(\mu, \mathfrak{s}) - \mathcal{D}(\mu, \mathfrak{s})$ is small.

Proposition 26. *For any $\mathfrak{s} \in \mathfrak{S}$, let $\mathcal{M} = [0, \tilde{\mu}]$ be the admissible Riccati set associated with \mathfrak{s} as defined in (72). For all $\mu \in \mathcal{M}$, let A_μ^c, D_μ be the associated Riccati quantities defined in (66). Then, whenever $\tilde{\mu} < +\infty$, it exists $|z| = 1$ such that*

$$\lim_{\mu \rightarrow \tilde{\mu}} |\det(Iz - A_\mu^c)|^2 \det(D_\mu) = 0. \quad (84)$$

Prop. 26 indicates that when μ approaches $\tilde{\mu}$, the stability and positivity conditions in (66) saturate, that is $\rho(A_\mu^c) \rightarrow 1$ and/or $\lambda_{\min}(D_\mu) \rightarrow 0$. However, according to Prop. 24, the optimal policies $\{\tilde{\pi}_\mu(\mathfrak{s})\}_{\mu \in \mathcal{M}}$ are uniformly stable when (80) holds. As a result, we have:

Corollary 2. *For any $\mathfrak{s} \in \mathfrak{S}$, let $\mathcal{M} = [0, \tilde{\mu}]$ be the admissible Riccati set associated with \mathfrak{s} as defined in (72). For all $\mu \in \mathcal{M}$, let A_μ^c, D_μ be the associated Riccati quantities defined in (66). Then, if for all $\mu \in \mathcal{M}$, $\mathcal{D}'(\mu, \mathfrak{s}) > 0$,*

$$\lim_{\mu \rightarrow \tilde{\mu}} \det(D_\mu) = 0. \quad (85)$$

Further,

$$\lambda_{\min}(D_\mu) \leq \alpha_D(\mathfrak{s}) |\tilde{\mu} - \mu|, \text{ where } \alpha_D(\mathfrak{s}) = 8 \|C_g\|_2 \kappa(\mathfrak{s})^4 ((2 + \|A\|_2 \|B\|_2)(1 + \|B\|_2))^2. \quad (86)$$

Comment on the role of $D_\mu = R_\mu + \tilde{B}^\top P_\mu \tilde{B}$. As stressed in Cor. 2, we are facing here the case where $\lim_{\mu \rightarrow \tilde{\mu}} \lambda_{\min}(D_\mu) = 0$. This paragraph aims at providing some insight on why this is equivalent to the fact that the optimal controller \tilde{K}_μ become less and less "unique" as μ approaches $\tilde{\mu}$.

First, notice that (66) focuses on particular solution of the Riccati equation: **1)** the condition $\rho(A_\mu^c) < 1$ imposes the solution

to be stabilizing; **2)** the condition $D_\mu \succ 0$ imposes some sort of global and unique optimality of the solution. Indeed, for any linear controller $\tilde{\pi}(x) = \tilde{K}x$, let $P_\mu(\tilde{K})$ be the solution of the Lyapunov equation

$$P_\mu(\tilde{K}) = A^c(\tilde{K})^\top P_\mu(\tilde{K}) A^c(\tilde{K}) + \begin{pmatrix} I \\ \tilde{K} \end{pmatrix}^\top C_\mu \begin{pmatrix} I \\ \tilde{K} \end{pmatrix}$$

Then, $\mathcal{L}_{\tilde{\pi}}(\mu) = \text{Tr}(P_\mu(\tilde{K}))$ and the sub-optimality gap is given by $\mathcal{L}_{\tilde{\pi}}(\mu) - \mathcal{D}(\mu) = \text{Tr}(\Delta)$ where $\Delta = P_\mu(\tilde{K}) - P_\mu$. Moreover, algebraic manipulations ensure that Δ is the solution of the Lyapunov equation:

$$\Delta = A^c(\tilde{K})^\top \Delta A^c(\tilde{K}) + (\tilde{K}_\mu - \tilde{K})^\top D_\mu (\tilde{K}_\mu - \tilde{K}). \quad (87)$$

(87) links the sub-optimality gap to the quantity $(\tilde{K}_\mu - \tilde{K})^\top D_\mu (\tilde{K}_\mu - \tilde{K})$. As a result, global optimality is asserted as soon as $D_\mu \succ 0$ since this implies that $\Delta \succ 0$. Further, $D_\mu \succ 0$ guarantees the uniqueness, while there exists a subspace of optimal solution if D_μ is rank deficient. Finally, when D_μ is arbitrarily close to 0, it may exist controllers $\tilde{K} \neq \tilde{K}_\mu$ with arbitrarily close value $\mathcal{L}_{\tilde{\pi}}(\mu) \approx \mathcal{D}(\mu)$.

Proof of Prop. 26. The proof follows from (69), (71) and (72). Let $\bar{K} = \begin{pmatrix} 0 \\ -A \end{pmatrix}$, which is such that $A^c(\bar{K}) = A + \bar{B}\bar{K} = A - A = 0$. The definition of $\tilde{\mu}$ in (72) and the fact that $\tilde{\mu} < +\infty$ ensure that it exists at least a $|z| = 1$ such that

$$\det(\Psi_{\tilde{\mu}}^{\bar{K}}(z, \mathfrak{s})) = 0,$$

hence the linearity of Ψ w.r.t. μ and (69) leads to $\lim_{\mu \rightarrow \tilde{\mu}} \det(\Psi_{\tilde{\mu}}^{\bar{K}}(z, \mathfrak{s})) = 0$. Further, (71) provides that for all $\mu \in \mathcal{M}$,

$$\begin{aligned} \det(\Psi_{\tilde{\mu}}^{\bar{K}}(z, \mathfrak{s})) &= \det\left(\left(I + z(\bar{K} - \tilde{K}_\mu)\tilde{B}\right)^\top D_\mu \left(I + z^{-1}(\bar{K} - \tilde{K}_\mu)\tilde{B}\right)\right) \\ &= \det(D_\mu) \left| \det\left(Iz + (\bar{K} - \tilde{K}_\mu)\tilde{B}\right) \right|^2. \end{aligned}$$

Using Sylvester's determinant identity ($\det(I+XY) = \det(I+YX)$ for any matrices X, Y), and the fact that $\tilde{B}(\bar{K} - \tilde{K}_\mu) = -A_\mu^c$, we obtain

$$\det(\Psi_{\tilde{\mu}}^{\bar{K}}(z, \mathfrak{s})) = \det(D_\mu) \left| \det(Iz - A_\mu^c) \right|^2 \quad (88)$$

Thus,

$$\lim_{\mu \rightarrow \tilde{\mu}} \det(D_\mu) \left| \det(Iz - A_\mu^c) \right|^2 = 0.$$

□

Proof of Cor. 2. The proof of Cor. 2 directly follows from propositions 23, 24 and 26.

1. Since $\rho(A_\mu^c) < 1$ for all $\mu \in \mathcal{M}$, $|\det(Iz - A_\mu^c)|^2 > 0$ for all $|z| = 1$. Hence, we have necessarily that $\lim_{\mu \rightarrow \tilde{\mu}} \det(D_\mu) = 0$.
2. $\lim_{\mu \rightarrow \tilde{\mu}} \det(D_\mu) = 0$ implies that for any \tilde{K} such that $|\lambda(A^c(\tilde{K}))| \neq 1$,

$$\det(\Psi_{\tilde{\mu}}^{\tilde{K}}(z, \mathfrak{s})) = 0 \quad \forall |z| = 1 \quad \Rightarrow \quad \det(\Psi_{\tilde{\mu}}^{\tilde{K}}(1, \mathfrak{s})) = 0 \quad \Rightarrow \quad \lambda_{\min}(\Psi_{\tilde{\mu}}^{\tilde{K}}(1, \mathfrak{s})) = 0.$$

Fix $\mu \in \mathcal{M}$. By definition $|\lambda(A^c(\tilde{K}_\mu))| \neq 1$, and thus $\lambda_{\min}(\Psi_{\tilde{\mu}}^{\tilde{K}_\mu}(1, \mathfrak{s})) = 0$. Let v_{\min} be an eigenvector associated with the zero eigenvalue of $\Psi_{\tilde{\mu}}^{\tilde{K}_\mu}(1, \mathfrak{s})$. The linearity of Ψ leads to

$$0 = v_{\min}^\top \Psi_{\tilde{\mu}}^{\tilde{K}_\mu}(1, \mathfrak{s}) v_{\min} = v_{\min}^\top \Psi_{\tilde{\mu}}^{\tilde{K}_\mu}(1, \mathfrak{s}) v_{\min} + (\tilde{\mu} - \mu) v_{\min}^\top \begin{pmatrix} (I - A_\mu^c)^{-1} \tilde{B} \\ I \end{pmatrix}^\top \begin{pmatrix} I & \tilde{K}_\mu^\top \\ 0 & I \end{pmatrix} C_g \begin{pmatrix} I & 0 \\ \tilde{K}_\mu & I \end{pmatrix} \begin{pmatrix} (I - A_\mu^c)^{-1} \tilde{B} \\ I \end{pmatrix} v_{\min}.$$

Thus,

$$\begin{aligned} v_{\min}^{\top} \Psi_{\mu}^{\tilde{K}_{\mu}}(1, \mathfrak{s}) v_{\min} &\leq |\tilde{\mu} - \mu| \|C_g\|_2 \left\| \begin{pmatrix} I & 0 \\ \tilde{K}_{\mu} & I \end{pmatrix} \right\|_2^2 v_{\min}^{\top} (I + \tilde{B}^{\top} (I - A_{\mu}^c)^{-\top} (I - A_{\mu}^c)^{-1} \tilde{B}) v_{\min} \\ &\leq 2|\tilde{\mu} - \mu| \|C_g\|_2 \left(1 + \left\| \begin{pmatrix} I \\ \tilde{K}_{\mu} \end{pmatrix} \right\|_2^2 \right) (1 + \|\tilde{B}\|_2^2 \|(I - A_{\mu}^c)^{-1}\|_2^2). \end{aligned}$$

From Prop. 25, we have $\left\| \begin{pmatrix} I \\ \tilde{K}_{\mu} \end{pmatrix} \right\|_2^2 \leq \kappa(\mathfrak{s})(2 + \|A\|_2 \|B\|_2)^2$. Further, since A_{μ}^c is stable,

$$\|(I - A_{\mu}^c)^{-1}\|_2 = \left\| \sum_{k=0}^{\infty} (A_{\mu}^c)^k \right\|_2 \leq \kappa(\mathfrak{s})^{3/2}.$$

As a result,

$$\begin{aligned} v_{\min}^{\top} \Psi_{\mu}^{\tilde{K}_{\mu}}(1, \mathfrak{s}) v_{\min} &\leq 2|\tilde{\mu} - \mu| \|C_g\|_2 (1 + \|\tilde{B}\|_2^2 \kappa(\mathfrak{s})^3) (1 + \kappa(\mathfrak{s})(2 + \|A\|_2 \|B\|_2)^2) \\ &\leq 8|\tilde{\mu} - \mu| \|C_g\|_2 \kappa(\mathfrak{s})^4 ((2 + \|A\|_2 \|B\|_2)(1 + \|B\|_2))^2. \end{aligned}$$

We conclude noting that by (71), we have

$$\lambda_{\min}(D_{\mu}) \leq v_{\min}^{\top} D_{\mu} v_{\min} = v_{\min}^{\top} \Psi_{\mu}^{\tilde{K}_{\mu}}(1, \mathfrak{s}) v_{\min}.$$

□

H.2.3. FEASIBLE AND ϵ -OPTIMAL LINEAR POLICIES

When for all $\mu \in \mathcal{M}$, $\mathcal{D}'(\mu, \mathfrak{s}) > 0$, Cor. 2 shows that $\lim_{\mu \rightarrow \tilde{\mu}} \lambda_{\min}(D_{\mu}) = 0$. Further, the previous comment suggests that it exists linear policies $\tilde{\pi}$ that yield close-to-optimal performances but can differ significantly from $\tilde{\pi}_{\mu}$. Lem. 13 shows that this is enough to guarantee the existence of ϵ -optimal feasible linear policy, while such policies are explicitly given in propositions 27 and 28, either in closed-form, either as the optimal solution of a modified well-posed Riccati equation.

Lemma 13. *For any $\mathfrak{s} \in \mathfrak{S}$, let \mathcal{M} be the admissible Riccati set associated with \mathfrak{s} as defined in (72). If for all $\mu \in \mathcal{M}$, $\mathcal{D}'(\mu, \mathfrak{s}) > 0$, then, for any $\epsilon > 0$, it exists $\tilde{\pi}^{\epsilon}(x) = \tilde{K}^{\epsilon} x$ and $\mu^{\epsilon} \in \mathcal{M}$ such that*

1. $g_{\tilde{\pi}^{\epsilon}}(\mathfrak{s}) \leq 0$,
2. $\mu^{\epsilon} g_{\tilde{\pi}^{\epsilon}}(\mathfrak{s}) = 0$,
3. $\mathcal{D}(\mu^{\epsilon}, \mathfrak{s}) \geq \mathcal{L}_{\tilde{\pi}^{\epsilon}}(\mu^{\epsilon}, \mathfrak{s}) - \epsilon$.

Lem. 13 proves strong-duality (in a weak "sup – min" sense) when $\mathcal{D}'(\mu, \mathfrak{s}) > 0$ for all $\mu \in \mathcal{M}$. Indeed, it implies that

$$\mathcal{J}_{*}(\mathfrak{s}) \stackrel{(1)}{\geq} \sup_{\mu \in \mathcal{M}} \mathcal{D}(\mu, \mathfrak{s}) \stackrel{(2)}{\geq} \mathcal{D}(\mu^{\epsilon}, \mathfrak{s}) \stackrel{(3)}{\geq} \mathcal{L}_{\tilde{\pi}^{\epsilon}}(\mu^{\epsilon}, \mathfrak{s}) - \epsilon \stackrel{(4)}{=} \mathcal{J}_{\tilde{\pi}^{\epsilon}} \stackrel{(5)}{\geq} \mathcal{J}_{*}(\mathfrak{s}) - \epsilon,$$

where (1) follows from weak duality, (2) from $\mu^{\epsilon} \in \mathcal{M}$, (3) from the ϵ -optimality of $\tilde{\pi}^{\epsilon}$, (4) from complementary slackness and (5) from primal feasibility.

As a result, Lem. 13 in conjunction with Lem. 12 proves Thm. 4.

We are left to prove Lem. 13. The analysis is conducted differently depending on the limiting null-space of D_{μ} . Formally, we distinguish between the following cases:

Case A. $\lambda_{\ker(\tilde{B})}(D_{\mu^{\epsilon}}) < \sqrt{\lambda_{\min}(D_{\mu^{\epsilon}})}$,

Case B). $\lambda_{\ker(\tilde{B})}(D_{\mu^\epsilon}) \geq \sqrt{\lambda_{\min}(D_{\mu^\epsilon})}$,

where $\lambda_{\ker(\tilde{B})}(D_{\mu^\epsilon}) = \min_{\|v\|=1, v \in \ker(\tilde{B})} v^\top D_{\mu^\epsilon} v$. Since $\mathcal{D}'(\mu, \mathfrak{s}) > 0$ for all $\mu \in \mathcal{M}$, propositions 23 and 26 ensure that for any arbitrary $1 > \epsilon > 0$, it exists μ^ϵ such that

$$\lambda_{\min}(D_{\mu^\epsilon}) = \nu(\epsilon)^2, \text{ where } \nu(\epsilon) = \min \left(\frac{\lambda_{\min}(C_\dagger^1)}{2\|\tilde{B}\|_2^2 \max(\mathcal{J}_*(\mathfrak{s}), 1)}, \frac{1}{8^{2n+1} \kappa(\mathfrak{s})^{2n}} \min \left(1, \frac{\lambda_{\min}(C_\dagger^1) \sigma_{\tilde{B}}^2}{2\kappa(\mathfrak{s})^2 c(\bar{\mu})} \right) \right)^2 \epsilon^2.$$

In **Case A**), i.e., when $\lambda_{\ker(\tilde{B})}(D_{\mu^\epsilon}) < \nu(\epsilon)$, we explicitly provide an ϵ -optimal feasible controller. Formally, we use Prop. 27 (we postpone the proof in to the next subsection), which explicit a modified policy linear $\tilde{\pi}^\epsilon$ given by the linear controller \tilde{K}^ϵ .

1. $g_{\tilde{\pi}^\epsilon}(\mathfrak{s}) = 0$,
2. $\mathcal{D}(\mu^\epsilon, \mathfrak{s}) \geq \mathcal{L}_{\tilde{\pi}^\epsilon}(\mu^\epsilon, \mathfrak{s}) - \epsilon$.

Proposition 27. For any $\mathfrak{s} \in \mathfrak{S}$, let \mathcal{M} be the admissible Riccati set associated with \mathfrak{s} as defined in (72). For any $\mu \in \mathcal{M}$, let \tilde{K}_μ and A_μ^c be respectively the optimal controller and closed-loop matrix defined by (66), and let $\Sigma_\mu = A_\mu^c \Sigma (A_\mu^c)^\top + I$ be its associated steady-state covariance matrix. If it exists $\bar{\mu} \in \mathcal{M}$ such that

1. $\mathcal{D}(\mu, \mathfrak{s}) > 0$
2. $\lambda_{\ker(\tilde{B})}(D_{\bar{\mu}}) \leq \frac{\lambda_{\min}(C_\dagger^1)}{2\|\tilde{B}\|_2^2 \max(\mathcal{J}_*(\mathfrak{s}), 1)}$ where $\lambda_{\ker(\tilde{B})}(D_{\bar{\mu}}) = \min_{\|v\|=1, v \in \ker(\tilde{B})} v^\top D_{\bar{\mu}} v$,

then, the modified linear policy $\bar{\pi}$ given by the modified controller \bar{K} defined as:

$$\begin{aligned} \bar{K} &= \tilde{K}_{\bar{\mu}} + \delta K; \quad \delta K = \alpha v x^\top, \\ v &= \arg \min_{\|v\|=1, v \in \ker(\tilde{B})} v^\top D_{\bar{\mu}} v, \\ x &\in \mathbb{R}^n \text{ is such that } \|x\| = 1 \text{ and } v^\top Y \Sigma_{\bar{\mu}} x = 0, \text{ where } Y = \begin{pmatrix} 0 \\ I_{n+d} \end{pmatrix}^\top C_g \begin{pmatrix} I \\ \tilde{K}_{\bar{\mu}} \end{pmatrix}, \\ \alpha^2 &= \frac{g_{\tilde{\pi}_{\bar{\mu}}}(\mathfrak{s})}{-v^\top Z v x^\top \Sigma_{\bar{\mu}} x} \quad \text{where } Z = \begin{pmatrix} 0 \\ I_{n+d} \end{pmatrix}^\top C_g \begin{pmatrix} 0 \\ I_{n+d} \end{pmatrix}. \end{aligned} \tag{89}$$

is well defined. Further, $\mathcal{D}(\bar{\mu}, \mathfrak{s}) \geq \mathcal{L}_{\bar{\pi}}(\bar{\mu}, \mathfrak{s}) - \lambda_{\ker(\tilde{B})}(D_{\bar{\mu}}) \frac{2\|\tilde{B}\|_2^2 \mathcal{J}_*(\mathfrak{s})}{\lambda_{\min}(C_\dagger^1)}$ and $g_{\bar{\pi}}(\mathfrak{s}) = 0$.

In **Case B**), i.e., when $\lambda_{\ker(\tilde{B})}(D_{\mu^\epsilon}) \geq \nu(\epsilon)$, we have that $\lambda_{\min}(D_{\mu^\epsilon})/\lambda_{\ker(\tilde{B})}(D_{\mu^\epsilon}) \leq \nu(\epsilon)$. We implicitly provide an ϵ -optimal feasible controller as the solution of the dual problem of a modified system. Formally, we use Prop. 28 (we postpone the proof in to the next subsection), which guarantees the existence of an optimal linear policy $\tilde{\pi}^\epsilon$ w.r.t. to a modified system for a Lagrangian parameter $\mu_*^\epsilon \in \mathcal{M}$ such that

1. $g_{\tilde{\pi}^\epsilon}(\mathfrak{s}) \leq 0$,
2. $\mu_*^\epsilon g_{\tilde{\pi}^\epsilon}(\mathfrak{s}) = 0$,
3. $\mathcal{D}(\mu^\epsilon, \mathfrak{s}) \geq \mathcal{L}_{\tilde{\pi}^\epsilon}(\mu^\epsilon, \mathfrak{s}) - \epsilon$.

Proposition 28. For any $\mathfrak{s} \in \mathfrak{S}$, let \mathcal{M} be the admissible Riccati set associated with \mathfrak{s} as defined in (72). For any $\mu \in \mathcal{M}$, let A_μ^c , D_μ be defined by (66) and \tilde{K}_μ be the optimal controller. We assume that it exists $\bar{\mu} \in \mathcal{M}$ such that

1. $\mathcal{D}(\bar{\mu}, \mathfrak{s}) > 0$

$$2. \frac{\lambda_{\min}(D_{\bar{\mu}})}{\lambda_{\ker(\tilde{B})}(D_{\bar{\mu}})} \leq \frac{1}{8^{2n+1}\kappa(\mathfrak{s})^{2n}} \min\left(1, \frac{\lambda_{\min}(C_{\dagger}^1)\sigma_{\tilde{B}}^2}{2c(\bar{\mu})\kappa(\mathfrak{s})}\right),$$

where

$$\begin{aligned} \lambda_{\ker(\tilde{B})}(D_{\bar{\mu}}) &= \min_{\|v\|=1, v \in \ker(\tilde{B})} v^{\top} D_{\bar{\mu}} v, & \sigma_{\tilde{B}}^2 &= \min_{\|v\|=1, v \in \text{Im}(\tilde{B})} v^{\top} \tilde{B}^{\top} \tilde{B} v, \\ \kappa(\mathfrak{s}) &= \mathcal{J}_*(\mathfrak{s})/\lambda_{\min}(C_{\dagger}^1), & c(\bar{\mu}) &= (\lambda_{\max}(C_{\dagger}) + \bar{\mu})(1 + \|\tilde{B}\|_2^2(1 + \|A\|_2^2)). \end{aligned}$$

Then, it exists a linear policy $\tilde{\pi}(\mathfrak{s}^n)$ and a Lagrangian parameter $\mu^n \in \mathcal{M}$ such that

1. $g_{\tilde{\pi}(\mathfrak{s}^n)}(\mathfrak{s}) \leq 0$,
2. $\mu^n g_{\tilde{\pi}(\mathfrak{s}^n)}(\mathfrak{s}) = 0$,
3. $\mathcal{D}(\mu^n, \mathfrak{s}) \geq \mathcal{L}_{\tilde{\pi}(\mathfrak{s}^n)}(\mu^n, \mathfrak{s}) - \frac{\lambda_{\min}(D_{\bar{\mu}})}{\lambda_{\ker(\tilde{B})}(D_{\bar{\mu}})} 2\kappa(\mathfrak{s})^2 \frac{c(\bar{\mu})}{\sigma_{\tilde{B}}^2} 8^{2n+1} \kappa(\mathfrak{s})^{2n}$.

For **Case A**) and **Case B**), we constructed a feasible linear policy $\tilde{\pi}^{\epsilon}$ that is ϵ -optimal, which concludes the proof of Lem. 13

H.2.4. PROOFS

Proof of Prop. 27. We first show that \bar{K} is well-defined, then quantify its sub-optimality gap and finally show that it satisfies primal feasibility.

1. Clearly, the existence of x is asserted (it is simply a unitary vector orthogonal to $y = Y \Sigma_{\bar{\mu}} Y^{\top} v$ that is well defined). Further, $g_{\tilde{\pi}_{\bar{\mu}}(\mathfrak{s})}(\mathfrak{s}) = \mathcal{D}(\bar{\mu}, \mathfrak{s}) > 0$ and $x^{\top} \Sigma_{\bar{\mu}} x \geq 1 > 0$. Thus, we only have to show that $-v^{\top} Z v > 0$ to assert the existence of \bar{K} .

From $v \in \ker(\tilde{B})$, we have

$$v^{\top} R_{\bar{\mu}} v = v^{\top} (R_{\bar{\mu}} + \tilde{B}^{\top} P_{\bar{\mu}} \tilde{B}) v = v^{\top} D_{\bar{\mu}} v = \lambda_{\ker(\tilde{B})}(D_{\bar{\mu}}) \leq \frac{\lambda_{\min}(C_{\dagger}^1)}{2\|\tilde{B}\|_2^2 \max(\mathcal{J}_*(\mathfrak{s}), 1)}.$$

Further,

$$R_{\bar{\mu}} = \begin{pmatrix} 0 \\ I_{n+d} \end{pmatrix}^{\top} C_{\bar{\mu}} \begin{pmatrix} 0 \\ I_{n+d} \end{pmatrix} = \begin{pmatrix} 0 \\ I_{n+d} \end{pmatrix}^{\top} C_{\dagger} \begin{pmatrix} 0 \\ I_{n+d} \end{pmatrix} + \bar{\mu} \begin{pmatrix} 0 \\ I_{n+d} \end{pmatrix}^{\top} C_g \begin{pmatrix} 0 \\ I_{n+d} \end{pmatrix} = \begin{pmatrix} 0 \\ I_{n+d} \end{pmatrix}^{\top} C_{\dagger} \begin{pmatrix} 0 \\ I_{n+d} \end{pmatrix} + \bar{\mu} Z.$$

Since $\mathfrak{s} \in \mathfrak{S}$, we have that $C_{\dagger} \succcurlyeq \begin{pmatrix} C_{\dagger}^1 & 0_{n,n+d} \\ 0_{n+d,n} & 0_n \end{pmatrix}$ and $C_{\dagger}^1 \in \mathbb{S}_{++}^{n+d}$. Thus,

$$R_{\bar{\mu}} \succcurlyeq \begin{pmatrix} 0_{n,d} & 0_{n,n} \\ I_d & 0_{d,n} \\ 0_{n,d} & I_n \end{pmatrix}^{\top} \begin{pmatrix} C_{\dagger}^1 & 0_{n,n+d} \\ 0_{n+d,n} & 0_n \end{pmatrix} \begin{pmatrix} 0_{n,d} & 0_{n,n} \\ I_d & 0_{d,n} \\ 0_{n,d} & I_n \end{pmatrix} + \bar{\mu} Z = \begin{pmatrix} 0_{n,d} & 0_{n,n} \\ I_d & 0_{d,n} \end{pmatrix}^{\top} C_{\dagger}^1 \begin{pmatrix} 0_{n,d} & 0_{n,n} \\ I_d & 0_{d,n} \end{pmatrix} + \bar{\mu} Z.$$

Now, decomposing $v \in \mathbb{R}^{n+d}$ into control $u \in \mathbb{R}^d$ and perturbation $w \in \mathbb{R}^n$ as $v^{\top} = (u^{\top} \quad w^{\top})$, we have

$$v^{\top} R_{\bar{\mu}} v \geq \begin{pmatrix} 0 \\ u \end{pmatrix}^{\top} C_{\dagger}^1 \begin{pmatrix} 0 \\ u \end{pmatrix} + \bar{\mu} v^{\top} Z v \geq \|u\|^2 \lambda_{\min}(C_{\dagger}^1) + \bar{\mu} v^{\top} Z v.$$

From $\tilde{B}v = 0$ we have that $Bu + w = 0$ and hence that $v = \begin{pmatrix} u \\ -Bu \end{pmatrix} = \begin{pmatrix} I_d \\ -B \end{pmatrix} u$. Further, $\|v\|^2 = 1$ implies that

$$1 = u^{\top} \begin{pmatrix} I_d & -B^{\top} \\ -B \end{pmatrix} \begin{pmatrix} I_d \\ -B \end{pmatrix} u = u^{\top} (I + B^{\top} B) u \leq \lambda_{\max}(I + B^{\top} B) \|u\|^2 = \lambda_{\max}(I + BB^{\top}) \|u\|^2 = \|\tilde{B}\|_2^2 \|u\|^2.$$

Summarizing, we have that

$$\begin{aligned} \frac{\lambda_{\min}(C_{\dagger}^1)}{2\|\tilde{B}\|_2^2 \max(\mathcal{J}_*(\mathfrak{s}), 1)} &\geq v^\top R_{\bar{\mu}} v \geq \lambda_{\min}(C_{\dagger}^1) / \|\tilde{B}\|_2^2 + \bar{\mu} v^\top Z v \\ \Rightarrow \bar{\mu} v^\top Z v &\leq \frac{\lambda_{\min}(C_{\dagger}^1)}{2\|\tilde{B}\|_2^2} - \lambda_{\min}(C_{\dagger}^1) / \|\tilde{B}\|_2^2 \leq -\frac{\lambda_{\min}(C_{\dagger}^1)}{2\|\tilde{B}\|_2^2} < 0. \end{aligned}$$

Hence, necessarily, $\bar{\mu} \neq 0$ and $v^\top Z v < 0$.

2. We now quantify the sub-optimality of \bar{K} . From (87)

$$P_{\bar{\mu}}(\bar{K}) - P_{\bar{\mu}} = (A_{\bar{\mu}}^c)^\top (P_{\bar{\mu}}(\bar{K}) - P_{\bar{\mu}}) A_{\bar{\mu}}^c + \delta K^\top D_{\bar{\mu}} \delta K.$$

Further, $\delta K^\top D_{\bar{\mu}} \delta K = \alpha^2 v^\top D_{\bar{\mu}} v x x^\top$, thus we obtain

$$\mathcal{L}_{\bar{\pi}}(\bar{\mu}, \mathfrak{s}) - \mathcal{D}(\bar{\mu}, \mathfrak{s}) \leq \alpha^2 v^\top D_{\bar{\mu}} v \text{Tr}(\Sigma_{\bar{\mu}} x x^\top) \leq \alpha^2 v^\top D_{\bar{\mu}} v x^\top \Sigma_{\bar{\mu}} x \leq v^\top D_{\bar{\mu}} v \frac{g_{\bar{\pi}_{\bar{\mu}}(\mathfrak{s})}(\mathfrak{s})}{-v^\top Z v}.$$

From $v^\top Z v \leq -\frac{\lambda_{\min}(C_{\dagger}^1)}{2\bar{\mu}\|\tilde{B}\|_2^2}$ and $\bar{\mu} g_{\bar{\pi}_{\bar{\mu}}(\mathfrak{s})}(\mathfrak{s}) \leq \mathcal{J}_{\bar{\pi}_{\bar{\mu}}(\mathfrak{s})}(\mathfrak{s}) + \bar{\mu} g_{\bar{\pi}_{\bar{\mu}}(\mathfrak{s})}(\mathfrak{s}) = \mathcal{D}(\bar{\mu}, \mathfrak{s}) \leq \mathcal{J}_*(\mathfrak{s})$ we obtain

$$\mathcal{L}_{\bar{\pi}}(\bar{\mu}, \mathfrak{s}) - \mathcal{D}(\bar{\mu}, \mathfrak{s}) \leq v^\top D_{\bar{\mu}} v \frac{2\|\tilde{B}\|_2^2 \mathcal{J}_*(\mathfrak{s})}{\lambda_{\min}(C_{\dagger}^1)} \leq \lambda_{\ker(\tilde{B})}(D_{\bar{\mu}}) \frac{2\|\tilde{B}\|_2^2 \mathcal{J}_*(\mathfrak{s})}{\lambda_{\min}(C_{\dagger}^1)}.$$

3. We now show that \bar{K} is feasible as $g_{\bar{\pi}}(\mathfrak{s}) = 0$. From Eqs. (60) and (61),

$$G(\bar{K}) = (A_{\bar{\mu}}^c)^\top G(\bar{K}) (A_{\bar{\mu}}^c) + \begin{pmatrix} I \\ \bar{K} \end{pmatrix}^\top C_g \begin{pmatrix} I \\ \bar{K} \end{pmatrix}.$$

Further,

$$\begin{pmatrix} I \\ \bar{K} \end{pmatrix}^\top C_g \begin{pmatrix} I \\ \bar{K} \end{pmatrix} = \begin{pmatrix} I \\ \tilde{K}_{\bar{\mu}} \end{pmatrix}^\top C_g \begin{pmatrix} I \\ \tilde{K}_{\bar{\mu}} \end{pmatrix} + \delta K^\top Y + Y^\top \delta K + \delta K^\top Z \delta K.$$

As a result,

$$\begin{aligned} g_{\bar{\pi}}(\mathfrak{s}) &= \text{Tr}(G(\bar{K})) = \text{Tr}(G(\tilde{K}_{\bar{\mu}})) + \text{Tr}(\Sigma_{\bar{\mu}}(\delta K^\top Y + Y^\top \delta K)) + \text{Tr}(\Sigma_{\bar{\mu}} \delta K^\top Z \delta K) \\ &= g_{\bar{\pi}_{\bar{\mu}}(\mathfrak{s})}(\mathfrak{s}) + 2\text{Tr}(Y \Sigma_{\bar{\mu}} \delta K) + \text{Tr}(\Sigma_{\bar{\mu}} \delta K^\top Z \delta K) \\ &= g_{\bar{\pi}_{\bar{\mu}}(\mathfrak{s})}(\mathfrak{s}) + 2\alpha v^\top Y \Sigma_{\bar{\mu}} x + \text{Tr}(\Sigma_{\bar{\mu}} \delta K^\top Z \delta K) \\ &= g_{\bar{\pi}_{\bar{\mu}}(\mathfrak{s})}(\mathfrak{s}) + \text{Tr}(\Sigma_{\bar{\mu}} \delta K^\top Z \delta K) \\ &= g_{\bar{\pi}_{\bar{\mu}}(\mathfrak{s})}(\mathfrak{s}) + \alpha^2 v^\top Z v \text{Tr}(\Sigma_{\bar{\mu}} x x^\top) \\ &= g_{\bar{\pi}_{\bar{\mu}}(\mathfrak{s})}(\mathfrak{s}) + \alpha^2 v^\top Z v x^\top \Sigma_{\bar{\mu}} x = g_{\bar{\pi}_{\bar{\mu}}(\mathfrak{s})}(\mathfrak{s}) - g_{\bar{\pi}_{\bar{\mu}}(\mathfrak{s})}(\mathfrak{s}) = 0. \end{aligned}$$

Proof of Prop. 28. We prove Prop. 28 as a corollary of the following proposition.

Proposition 29. For any $\mathfrak{s} \in \mathfrak{S}$, let \mathcal{M} be the admissible Riccati set associated with \mathfrak{s} as defined in (72). For any $\mu \in \mathcal{M}$, let A_{μ}^c , D_{μ} be defined by (66) and \tilde{K}_{μ} be the optimal controller. We assume that it exists $\bar{\mu} \in \mathcal{M}$ such that

1. $\mathcal{D}(\bar{\mu}, \mathfrak{s}) > 0$

2. $\frac{\lambda_{\min}(D_{\bar{\mu}})}{\lambda_{\ker(\tilde{B})}(D_{\bar{\mu}})} \leq \frac{1}{8^{2n+1} \kappa(\mathfrak{s})^{2n}} \min \left(1, \frac{\lambda_{\min}(C_{\dagger}^1) \sigma_{\tilde{B}}^2}{2c(\bar{\mu}) \kappa(\mathfrak{s})} \right),$

where

$$\begin{aligned}\lambda_{\ker(\tilde{B})}(D_{\bar{\mu}}) &= \min_{\|v\|=1, v \in \ker(\tilde{B})} v^\top D_{\bar{\mu}} v, & \sigma_{\tilde{B}}^2 &= \min_{\|v\|=1, v \in \text{Im}(\tilde{B})} v^\top \tilde{B}^\top \tilde{B} v, \\ \kappa(\mathfrak{s}) &= \mathcal{J}_*(\mathfrak{s})/\lambda_{\min}(C_\dagger^1), & c(\bar{\mu}) &= (\lambda_{\max}(C_\dagger) + \bar{\mu})(1 + \|\tilde{B}\|_2^2(1 + \|A\|_2^2)).\end{aligned}$$

Let \mathfrak{s}^η be a modified system w.r.t. \mathfrak{s} as $\mathfrak{s}^\eta = (A, \tilde{B}, C_\dagger + \eta\Delta, C_g)$ where

$$\begin{aligned}\Delta &= \begin{pmatrix} I & -\bar{K}^\top \\ 0 & I \end{pmatrix} \begin{pmatrix} I & \\ & -\tilde{B}^\top \end{pmatrix} \begin{pmatrix} I & -\tilde{B} \\ & -\bar{K} \end{pmatrix} \begin{pmatrix} I & 0 \\ & I \end{pmatrix}, \\ \bar{K} &= \begin{pmatrix} 0 \\ -A \end{pmatrix} \\ \frac{\lambda_{\min}(C_\dagger^1)}{2\kappa(\mathfrak{s})} &\geq \eta \geq \frac{c(\bar{\mu})}{\sigma_{\tilde{B}}^2} \frac{8^{2n+1} \kappa(\mathfrak{s})^{2n} \lambda_{\min}(D_{\bar{\mu}})}{\lambda_{\ker(\tilde{B})}(D_{\bar{\mu}})}.\end{aligned}\tag{90}$$

Let \mathcal{M}^η be the admissible Riccati set associated with \mathfrak{s}^η as defined in (72). Then,

1. $\mathfrak{s}^\eta \in \mathfrak{S}$,
2. $[0, \bar{\mu}] \subset \mathcal{M} \subset \mathcal{M}^\eta$,
3. for all $\mu \in [0, \bar{\mu}]$, $\mathcal{D}(\mu, \mathfrak{s}^\eta) \leq \mathcal{D}(\mu, \mathfrak{s}) + 2\eta\kappa(\mathfrak{s})^2$ and $\mathcal{D}(\mu, \mathfrak{s}^\eta) \geq \mathcal{L}_{\tilde{\pi}_\mu(\mathfrak{s}^\eta)}(\mu, \mathfrak{s})$ where $\tilde{\pi}_\mu(\mathfrak{s}^\eta) = \arg \min_{\tilde{\pi}} \mathcal{L}_{\tilde{\pi}}(\mu, \mathfrak{s}^\eta)$,
4. $\mathcal{D}'(\bar{\mu}, \mathfrak{s}^\eta) < 0$,
5. for all $\mu \in [0, \bar{\mu}]$, $D_\mu^\eta \succ \min\left(\lambda_{\min}(D_0), \min\left(1, \eta\sigma_{\tilde{B}}^2/c(\bar{\mu})\right) \frac{\lambda_{\ker(\tilde{B})}(D_{\bar{\mu}})}{8}\right) I$, where D_μ^η is given in (66) for the modified system \mathfrak{s}^η .

Prop. 29 explicit a set modified system $\{\mathfrak{s}^\eta\}_\eta \in \mathfrak{S}$ such that the dual function $\mathcal{D}(\cdot, \mathfrak{s}^\eta)$ is non-increasing on a subset of its domain. Hence, we can invoke Lem. 12 to obtain that it exists $\mu^\eta \in [0, \bar{\mu}] \subset \mathcal{M}$ and $\tilde{\pi}(\mathfrak{s})^\eta = \tilde{\pi}_{\mu^\eta}(\mathfrak{s}^\eta)$ such that:

1. $g_{\tilde{\pi}(\mathfrak{s}^\eta)}(\mathfrak{s}) \leq 0$,
2. $\mu^\eta g_{\tilde{\pi}(\mathfrak{s}^\eta)}(\mathfrak{s}) = 0$.

This proves the first two assertions in Prop. 28. The third assertion in Prop. 29 leads to

$$\mathcal{D}(\mu^\eta, \mathfrak{s}) \geq \mathcal{D}(\mu^\eta, \mathfrak{s}^\eta) - 2\kappa(\mathfrak{s})^2\eta = \mathcal{L}_{\tilde{\pi}(\mathfrak{s}^\eta)}(\mu^\eta, \mathfrak{s}^\eta) - 2\kappa(\mathfrak{s})^2\eta \geq \mathcal{L}_{\tilde{\pi}(\mathfrak{s}^\eta)}(\mu^\eta, \mathfrak{s}) - 2\kappa(\mathfrak{s})^2\eta.$$

Setting $\eta = \frac{c(\bar{\mu})}{\sigma_{\tilde{B}}^2} \frac{8^{2n+1} \kappa(\mathfrak{s})^{2n} \lambda_{\min}(D_{\bar{\mu}})}{\lambda_{\ker(\tilde{B})}(D_{\bar{\mu}})}$ proves the last assertion in Prop. 28 and concludes the proof. We are thus left to prove Prop. 29.

Proof of Prop. 29. 1. First, notice that from $\mathcal{D}'(\bar{\mu}, \mathfrak{s}) > 0$, Prop. 24 ensures that $\rho(A_{\bar{\mu}}^c) < 1$ and hence $(I - A_{\bar{\mu}}^c)$ is full rank. Thus, Δ is well defined.

The modified \mathfrak{s}^η consists in adding a p.s.d perturbation $\eta\Delta$ to the cost matrix of the original system C_\dagger , while leaving the dynamics parametrization (A, \tilde{B}) and the constraint cost C_g unchanged. Thus, the structure is still compatible with \mathfrak{S} while the existence of a stable linear feasible policy for \mathfrak{s}^η is inherited from $\mathfrak{s} \in \mathfrak{S}$. As a result, $\mathfrak{s}^\eta \in \mathfrak{S}$.

2. Since $\eta\Delta \succcurlyeq 0$, $C_\dagger + \eta\Delta \succcurlyeq C_\dagger$, the modification enforces the positive definiteness of \mathfrak{s} , and as a result, extends its Riccati admissible domain. For all $\mu \geq 0$, for any \tilde{K} such that $|\lambda(A_\mu^c(\tilde{K}))| \neq 1$, for all $|z| = 1$,

$$\Psi_\mu^{\tilde{K}}(z, \mathfrak{s}^\eta) = \Psi_\mu^{\tilde{K}}(z, \mathfrak{s}) + \eta Y^{\tilde{K}}(z^{-1})^\top \Delta Y^{\tilde{K}}(z) \succcurlyeq \Psi_\mu^{\tilde{K}}(z, \mathfrak{s}).$$

Thus, $\mu \in \mathcal{M}$ implies that $\Psi_\mu^{\tilde{K}}(z, \mathfrak{s}) \succ 0$ which implies that $\Psi_\mu^{\tilde{K}}(z, \mathfrak{s}^\eta) \succ 0$ which leads to $\mu \in \mathcal{M}^\eta$. As a result, $\mathcal{M} \subset \mathcal{M}^\eta$.

3. For all $\mu \in [0, \bar{\mu}]$, let $\tilde{\pi}_\mu(\mathfrak{s}^\eta)$ and $\tilde{\pi}_\mu(\mathfrak{s})$ the optimal policy at μ for the system \mathfrak{s}^η and \mathfrak{s} . Since Δ is p.s.d., we immediately have that $\mathcal{L}_{\tilde{\pi}_\mu(\mathfrak{s}^\eta)}(\mu, \mathfrak{s}) \leq \mathcal{D}(\mu, \mathfrak{s}^\eta)$. To prove the r.h.s. inequality, we rely on Prop. 24.

Let $\tilde{K}_\mu^\eta, \tilde{K}_\mu, A_\mu^{\eta,c}$ and A_μ^c be the linear controller and closed-loop matrix associated with $\tilde{\pi}_\mu(\mathfrak{s}^\eta)$ and $\tilde{\pi}_\mu(\mathfrak{s})$. Then,

$$\mathcal{D}(\mu, \mathfrak{s}^\eta) = \mathcal{L}_{\tilde{\pi}_\mu(\mathfrak{s}^\eta)}(\mu, \mathfrak{s}^\eta) \leq \mathcal{L}_{\tilde{\pi}_\mu(\mathfrak{s})}(\mu, \mathfrak{s}^\eta) = \mathcal{L}_{\tilde{\pi}_\mu(\mathfrak{s})}(\mu, \mathfrak{s}) + \eta \text{Tr}(X) = \mathcal{D}(\mu, \mathfrak{s}) + \eta \text{Tr}(X), \quad (91)$$

where X is the solution of the Lyapunov equation

$$X = (A_\mu^c)^\top X A_\mu^c + \eta \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix}^\top \Delta \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix}.$$

From $\tilde{B}(\tilde{K}_\mu - \bar{K}) = \tilde{B}\tilde{K}_\mu + (B \quad I) \begin{pmatrix} 0 \\ -A \end{pmatrix} = A_\mu^c$, we have

$$\begin{aligned} \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix}^\top \Delta \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix} &= (I - \tilde{B}(\tilde{K}_\mu - \bar{K}))^\top (I - \tilde{B}(\tilde{K}_\mu - \bar{K})) \\ &= (I - A_\mu^c)^\top (I - A_\mu^c) \\ &\leq \|I - A_\mu^c\|_2^2 I \\ &\leq 2(1 + \|A_\mu^c\|_2^2) I \end{aligned}$$

As a result, $\text{Tr}(X) \leq 2(1 + \|A_\mu^c\|_2^2) \text{Tr}(\Sigma_\mu)$ where $\Sigma_\mu = A_\mu^c \Sigma_\mu (A_\mu^c)^\top + I$ is the steady-state covariance of the state process driven by A_μ^c . Since $\mathcal{D}(\bar{\mu}, \mathfrak{s}) > 0$, the concavity of $\mathcal{D}(\cdot, \mathfrak{s})$ ensures that $\mathcal{D}(\mu, \mathfrak{s}) > 0$ for all $\mu \in [0, \bar{\mu}]$. Hence, Prop. 24 guarantees that

$$\text{Tr}(\Sigma_\mu) \leq \kappa(\mathfrak{s}) := \mathcal{J}_*(\mathfrak{s}) / \lambda_{\min}(C_\dagger^\dagger).$$

Further, noticing that $\Sigma_\mu \succcurlyeq I$, we have

$$\kappa(\mathfrak{s}) \geq \text{Tr}(\Sigma_\mu) \geq \|\Sigma_\mu\|_2 = \max_{\|x\|=1} x^\top \Sigma_\mu x = \max_{\|x\|=1} x^\top A_\mu^c \Sigma_\mu (A_\mu^c)^\top x + 1 \geq 1 + \max_{\|x\|=1} x^\top A_\mu^c (A_\mu^c)^\top x = 1 + \|A_\mu^c\|_2^2,$$

and we obtain that $\text{Tr}(X) \leq 2\kappa(\mathfrak{s})^2$, which we use in (91) to prove the third statement.

4. Let $\Psi_\mu^{\tilde{K}}(z; \mathfrak{s}^\eta)$ and $\Psi_\mu^{\tilde{K}}(z; \mathfrak{s})$ be the Popov functions associated with \mathfrak{s}^η and \mathfrak{s} as defined in (69). The structure of the modification Δ ensures that

$$\begin{aligned} \Psi_\mu^{\tilde{K}}(1; \mathfrak{s}^\eta) &= \Psi_\mu^{\tilde{K}}(1; \mathfrak{s}) + \eta \begin{pmatrix} \tilde{B} \\ I \end{pmatrix}^\top \begin{pmatrix} I & \bar{K}^\top \\ 0 & I \end{pmatrix} \Delta \begin{pmatrix} I & 0 \\ \bar{K} & I \end{pmatrix} \begin{pmatrix} \tilde{B} \\ I \end{pmatrix} \\ &= \Psi_\mu^{\tilde{K}}(1; \mathfrak{s}) + \eta \begin{pmatrix} \tilde{B} \\ I \end{pmatrix}^\top \begin{pmatrix} I \\ -\tilde{B}^\top \end{pmatrix} (I \quad -\tilde{B}) \begin{pmatrix} \tilde{B} \\ I \end{pmatrix} \\ &= \Psi_\mu^{\tilde{K}}(1; \mathfrak{s}). \end{aligned} \quad (92)$$

Further, (71) provides that

$$\begin{aligned} \Psi_\mu^{\tilde{K}}(1; \mathfrak{s}^\eta) &= (I + (\bar{K} - \tilde{K}_\mu^\eta) \tilde{B})^\top D_\mu^\eta (I + (\bar{K} - \tilde{K}_\mu^\eta) \tilde{B}), \\ \Psi_\mu^{\tilde{K}}(1; \mathfrak{s}) &= (I + (\bar{K} - \tilde{K}_\mu) \tilde{B})^\top D_\mu (I + (\bar{K} - \tilde{K}_\mu) \tilde{B}), \end{aligned} \quad (93)$$

where D_μ and D_μ^η are given in (66) for \mathfrak{s} and \mathfrak{s}^η respectively. Thus combining Eqs. (92) and (93) at $\bar{\mu}$ leads to

$$(I + (\bar{K} - \tilde{K}_\mu^\eta) \tilde{B})^\top D_\mu^\eta (I + (\bar{K} - \tilde{K}_\mu^\eta) \tilde{B}) = (I + (\bar{K} - \tilde{K}_\mu) \tilde{B})^\top D_\mu (I + (\bar{K} - \tilde{K}_\mu) \tilde{B}).$$

Taking the determinant and using Sylvester's determinant identity ($\det(I + XY) = \det(I + YX)$) for any rectangular matrices X, Y , we obtain

$$\begin{aligned} \det(I + (\bar{K} - \tilde{K}_\mu^\eta) \tilde{B})^2 \det(D_\mu^\eta) &= \det(I + (\bar{K} - \tilde{K}_\mu) \tilde{B})^2 \det(D_\mu), \\ \Leftrightarrow \det(I - A_\mu^{\eta,c})^2 \det(D_\mu^\eta) &= \det(I - A_\mu^c)^2 \det(D_\mu). \end{aligned}$$

Let $\{\lambda_i^\eta\}_{i=1,\dots,n}$ and $\{\lambda_i\}_{i=1,\dots,n}$ be the eigenvalues of $A_{\bar{\mu}}^{\eta,c}$ and $A_{\bar{\mu}}^c$ respectively. Since $A_{\bar{\mu}}^c$ is stable,

$$\det(I - A_{\bar{\mu}}^c)^2 \leq \prod_{i=1}^n |1 - \lambda_i|^2 \leq 2^{2n},$$

$$\det(I - A_{\bar{\mu}}^{\eta,c})^2 \geq \prod_{i=1}^n |1 - \lambda_i|^2 \geq \prod_{i=1}^n (1 - |\lambda_i|)^2 \geq (1 - \max_{i=1,\dots,n} |\lambda_i|)^{2n} = (1 - \rho(A_{\bar{\mu}}^{\eta,c}))^{2n}.$$

Hence,

$$(1 - \rho(A_{\bar{\mu}}^{\eta,c}))^{2n} \leq 2^{2n} \det(D_{\bar{\mu}}) / \det(D_{\bar{\mu}}^\eta).$$

Let $P_{\bar{\mu}}^\eta$ and $P_{\bar{\mu}}$ be the solution of the Riccati equation at $\bar{\mu}$ associated with \mathfrak{s}^η and \mathfrak{s} respectively. Since Δ is p.s.d. we have that $P_{\bar{\mu}}^\eta \succcurlyeq P_{\bar{\mu}}$. Algebraic manipulations show that $D_{\bar{\mu}}^\eta = R_{\bar{\mu}} + \eta \tilde{B}^\top \tilde{B} + \tilde{B}^\top P_{\bar{\mu}}^\eta \tilde{B} \succcurlyeq D_{\bar{\mu}} + \eta \tilde{B}^\top \tilde{B}$. Thus,

$$\frac{\det(D_{\bar{\mu}})}{\det(D_{\bar{\mu}}^\eta)} \leq \frac{\det(D_{\bar{\mu}})}{\det(D_{\bar{\mu}} + \eta \tilde{B}^\top \tilde{B})}.$$

Further, for any p.s.d matrix X and Y such that $X \succcurlyeq Y \succ 0$, we have that $\frac{\det(Y)}{\det(X)} \leq \inf_{x \neq 0} \frac{x^\top Y x}{x^\top X x}$. This identity can be found in (Abbasi-Yadkori & Szepesvári, 2011, Lem.11) for instance. Clearly, $D_{\bar{\mu}} + \eta \tilde{B}^\top \tilde{B} \succcurlyeq D_{\bar{\mu}}$, and applying the given identity at v_{\min} where v_{\min} is a eigenvector associated with the smallest eigenvalue of $D_{\bar{\mu}}$, we obtain

$$(1 - \rho(A_{\bar{\mu}}^{\eta,c}))^{2n} \leq 2^{2n} \frac{\lambda_{\min}(D_{\bar{\mu}})}{v_{\min}^\top (D_{\bar{\mu}} + \eta \tilde{B}^\top \tilde{B}) v_{\min}} \leq 2^{2n} \frac{\lambda_{\min}(D_{\bar{\mu}})}{\lambda_{\min}(D_{\bar{\mu}} + \eta \tilde{B}^\top \tilde{B})}.$$

However, we have by property 3 that $\lambda_{\min}(D_{\bar{\mu}} + \eta \tilde{B}^\top \tilde{B}) > \lambda_{\ker(\tilde{B})}(D_{\bar{\mu}}) \min(1/8, \eta \sigma_{\tilde{B}}^2 / (8\lambda_{\max}(D_{\bar{\mu}})))$, which together with Prop. 30 ensures that

$$\lambda_{\min}(D_{\bar{\mu}} + \eta \tilde{B}^\top \tilde{B}) > \lambda_{\ker(\tilde{B})}(D_{\bar{\mu}}) \min(1/8, \eta \sigma_{\tilde{B}}^2 / (8c(\bar{\mu}))).$$

Thus,

$$(1 - \rho(A_{\bar{\mu}}^{\eta,c}))^{2n} < 2^{2n} \frac{\lambda_{\min}(D_{\bar{\mu}})}{\lambda_{\ker(\tilde{B})}(D_{\bar{\mu}}) \min(1/8, \eta \sigma_{\tilde{B}}^2 / (8c(\bar{\mu})))}.$$

Since, $\eta \geq \frac{c(\bar{\mu})}{\sigma_{\tilde{B}}^2} \frac{8^{2n+1} \kappa(\mathfrak{s})^{2n} \lambda_{\min}(D_{\bar{\mu}})}{\lambda_{\ker(\tilde{B})}(D_{\bar{\mu}})}$ and $\frac{\lambda_{\min}(D_{\bar{\mu}})}{\lambda_{\ker(\tilde{B})}(D_{\bar{\mu}})} \leq \frac{1}{8^{2n+1} \kappa(\mathfrak{s})^{2n}}$,

- when $\frac{\sigma_{\tilde{B}}^2 \eta}{c(\bar{\mu})} \leq 1$, $(1 - \rho(A_{\bar{\mu}}^{\eta,c}))^{2n} < 2^{2n} \frac{8c(\bar{\mu}) \lambda_{\min}(D_{\bar{\mu}})}{\lambda_{\ker(\tilde{B})}(D_{\bar{\mu}}) \eta \sigma_{\tilde{B}}^2} \leq \frac{1}{4^{2n} \kappa(\mathfrak{s})^{2n}}$,
- when $\frac{\sigma_{\tilde{B}}^2 \eta}{c(\bar{\mu})} \geq 1$, $(1 - \rho(A_{\bar{\mu}}^{\eta,c}))^{2n} < 2^{2n} 8 \frac{\lambda_{\min}(D_{\bar{\mu}})}{\lambda_{\ker(\tilde{B})}(D_{\bar{\mu}})} \leq \frac{1}{4^{2n} \kappa(\mathfrak{s})^{2n}}$.

Thus,

$$(1 - \rho(A_{\bar{\mu}}^{\eta,c})) < \frac{1}{4\kappa(\mathfrak{s})}. \quad (94)$$

Now, suppose that $\mathcal{D}'(\bar{\mu}, \mathfrak{s}^\eta) \geq 0$, by Prop. 24 it implies that $1 - \rho(A_{\bar{\mu}}^{\eta,c})^2 \geq 1/\kappa(\bar{\mu}, \mathfrak{s}^\eta)$ where $\kappa(\bar{\mu}, \mathfrak{s}^\eta) = \mathcal{D}(\bar{\mu}, \mathfrak{s}^\eta) / \lambda_{\min}(C_{\dagger}^1)$. Further, from $\mathcal{D}(\bar{\mu}, \mathfrak{s}^\eta) \leq \mathcal{D}(\bar{\mu}, \mathfrak{s}) + 2\eta \kappa(\mathfrak{s})^2$ and $\eta \leq \frac{\lambda_{\min}(C_{\dagger}^1)}{2\kappa(\mathfrak{s})}$, we obtain

$$\mathcal{D}(\bar{\mu}, \mathfrak{s}^\eta) \leq \mathcal{J}_*(\mathfrak{s}) + \kappa(\mathfrak{s}) \lambda_{\min}(C_{\dagger}^1) \Rightarrow \kappa(\bar{\mu}, \mathfrak{s}^\eta) \leq 2\kappa(\mathfrak{s}).$$

Thus,

$$2(1 - \rho(A_{\bar{\mu}}^{\eta,c})) \geq 1 - \rho(A_{\bar{\mu}}^{\eta,c})^2 \geq \frac{1}{2\kappa(\mathfrak{s})} \Rightarrow 1 - \rho(A_{\bar{\mu}}^{\eta,c}) \geq \frac{1}{4\kappa(\mathfrak{s})},$$

which contradicts (94). As a result, $\mathcal{D}'(\bar{\mu}, \mathfrak{s}^\eta) < 0$ which proves the statement.

5. To prove the final statement, we use Prop. 30, which states that for any $v \in \mathbb{R}^{n+d}$, the function $\mu \in \mathcal{M} \rightarrow v^\top D_\mu^\eta v$ is quasi-concave. As a result, for all $\mu \in [0, \bar{\mu}]$, let $\lambda_{\min}(D_\mu^\eta)$ be the smallest eigenvalue of D_μ^η and let v_{\min} be a corresponding eigenvector, we have that

$$\lambda_{\min}(D_\mu^\eta) = v_{\min}^\top D_\mu^\eta v_{\min} \geq \min \left(v_{\min}^\top D_0^\eta v_{\min}, v_{\min}^\top D_{\bar{\mu}}^\eta v_{\min} \right) \geq \min \left(\lambda_{\min}(D_0^\eta), \lambda_{\min}(D_{\bar{\mu}}^\eta) \right).$$

Thus, we conclude the proof using that

$$\begin{aligned} \lambda_{\min}(D_{\bar{\mu}}^\eta) &\geq \lambda_{\min}(D_{\bar{\mu}} + \eta \tilde{B} \tilde{B}^\top) > \min \left(1, \eta \sigma_{\tilde{B}}^2 / c(\bar{\mu}) \right) \frac{\lambda_{\ker(\tilde{B})}(D_{\bar{\mu}})}{8}, \\ \lambda_{\min}(D_0^\eta) &\geq \lambda_{\min}(D_0 + \eta \tilde{B} \tilde{B}^\top) \geq \lambda_{\min}(D_0). \end{aligned}$$

Notice that $D_0 \succ 0$ since $0 \in \mathcal{M}$ and is fully characterized by \mathfrak{s} i.e., it has no dependency on μ . As a result, we obtain

$$\lambda_{\min}(D_\mu^\eta) \geq \min \left(\lambda_{\min}(D_0), \min \left(1, \eta \sigma_{\tilde{B}}^2 / c(\bar{\mu}) \right) \frac{\lambda_{\ker(\tilde{B})}(D_{\bar{\mu}})}{8} \right).$$

□

Property 3. Let $D \in \mathbb{S}_{++}^{n+d}$ and $\tilde{B} \in \mathbb{R}^{n,n+d}$. Let $\lambda_{\max}(D)$ denote the maximum eigenvalue of D and

$$\lambda_{\ker(\tilde{B})}(D) = \min_{\|v\|=1, v \in \ker(\tilde{B})} v^\top D v, \quad \sigma_{\tilde{B}}^2 = \min_{\|v\|=1, v \in \text{Im}(\tilde{B})} v^\top \tilde{B}^\top \tilde{B} v,$$

be the singular value of D and \tilde{B} on the null space and row space of \tilde{B} respectively¹⁸. Then, for any $\eta \geq 0$,

$$\lambda_{\min}(D + \eta \tilde{B}^\top \tilde{B}) > \lambda_{\ker(\tilde{B})}(D) \min \left(1/8, \eta \sigma_{\tilde{B}}^2 / (8 \lambda_{\max}(D)) \right).$$

Proof of property 3. First, we have that

$$\lambda_{\min}(D + \eta \tilde{B}^\top \tilde{B}) \geq \min_{\|v\|=1} v^\top (D + \eta \tilde{B}^\top \tilde{B}) v \geq \min_{\|v\|=1} \max \left(v^\top D v, \eta \|\tilde{B} v\|^2 \right). \quad (95)$$

We thus provide a lower bound for $v^\top D v$ and $\eta \|\tilde{B} v\|^2$ separately and then show a lower bound for the minimum of the maximum between those lower bound. Let $P_{\tilde{B}}$ be defined the orthogonal projection matrix onto \tilde{B} , which satisfies $\tilde{B} P_{\tilde{B}} = 0$, $P_{\tilde{B}} = P_{\tilde{B}}^\top = P_{\tilde{B}}^2$. Then,

1. for any $v \in \mathbb{R}^{n+d}$, $\|\tilde{B} v\| = \|\tilde{B}(v - P_{\tilde{B}} v)\|$. Since $v - P_{\tilde{B}} v \in \text{Im}(\tilde{B})$, $\|\tilde{B}(v - P_{\tilde{B}} v)\| \geq \sigma_{\tilde{B}} \|v - P_{\tilde{B}} v\|$. Thus,

$$\eta \|\tilde{B} v\|^2 \geq \sigma_{\tilde{B}}^2 \|v - P_{\tilde{B}} v\|^2. \quad (96)$$

2. the function $v \rightarrow v^\top D v$ is convex since $D \in \mathbb{S}_{++}^{n+d}$. As a result, for any $v \in \mathbb{R}^{n+d}$,

$$v^\top D v \geq v^\top P_{\tilde{B}}^\top D P_{\tilde{B}} v + 2v^\top P_{\tilde{B}}^\top D (v - P_{\tilde{B}} v) \geq \|P_{\tilde{B}} v\|_D^2 - 2\sqrt{\lambda_{\max}(D)} \|P_{\tilde{B}} v\|_D \|v - P_{\tilde{B}} v\|. \quad (97)$$

Plugging Eqs. (96) and (97) in (95) we obtain

$$\lambda_{\min}(D + \eta \tilde{B}^\top \tilde{B}) \geq \min_{\|v\|=1} \max \left(\|P_{\tilde{B}} v\|_D^2 - 2\sqrt{\lambda_{\max}(D)} \|P_{\tilde{B}} v\|_D \|v - P_{\tilde{B}} v\|, \eta \sigma_{\tilde{B}}^2 \|v - P_{\tilde{B}} v\|^2 \right). \quad (98)$$

- When $\|v - P_{\tilde{B}} v\|^2 > \frac{\lambda_{\ker(\tilde{B})}(D)}{8 \lambda_{\max}(D)}$, we use the r.h.s lower bound and get

$$\max \left(\|P_{\tilde{B}} v\|_D^2 - 2\sqrt{\lambda_{\max}(D)} \|P_{\tilde{B}} v\|_D \|v - P_{\tilde{B}} v\|, \eta \sigma_{\tilde{B}}^2 \|v - P_{\tilde{B}} v\|^2 \right) > \eta \sigma_{\tilde{B}}^2 \frac{\lambda_{\ker(\tilde{B})}(D)}{8 \lambda_{\max}(D)}.$$

¹⁸Notice that $\sigma_{\tilde{B}}^2 > 0$ as it is the smallest non-zero eigenvalue of $\tilde{B}^\top \tilde{B}$.

- When $\|v - P_{\tilde{B}}v\|^2 \leq \frac{\lambda_{\ker(\tilde{B})}(D)}{8\lambda_{\max}(D)}$, we use the l.h.s. lower bound and get

$$\|P_{\tilde{B}}v\|_D^2 - 2\sqrt{\lambda_{\max}(D)}\|P_{\tilde{B}}v\|_D\|v - P_{\tilde{B}}v\| \geq \|P_{\tilde{B}}v\|_D^2 - \sqrt{\lambda_{\ker(\tilde{B})}(D)/2}\|P_{\tilde{B}}v\|_D.$$

The r.h.s. is now lower bounded by a second order polynomial in $\|P_{\tilde{B}}v\|_D$, whose roots are given by 0 and $\sqrt{\lambda_{\ker(\tilde{B})}(D)/2}$. However, from $\|v - P_{\tilde{B}}v\|^2 \leq \frac{\lambda_{\ker(\tilde{B})}(D)}{8\lambda_{\max}(D)}$ and $\lambda_{\ker(\tilde{B})}(D) \leq \lambda_{\max}(D)$, we have that

$$\|v - P_{\tilde{B}}v\|^2 \leq 1/8 \quad \Rightarrow \quad \|P_{\tilde{B}}\|^2 \geq 1 - 1/8 > 1/2 \quad \Rightarrow \quad \|P_{\tilde{B}}v\|_D \geq \sqrt{\lambda_{\ker(\tilde{B})}(D)7/8} > \sqrt{\lambda_{\ker(\tilde{B})}(D)/2}.$$

As a result, the polynomial is lower bounded as

$$\|P_{\tilde{B}}v\|_D^2 - \sqrt{\lambda_{\ker(\tilde{B})}(D)/2}\|P_{\tilde{B}}v\|_D \geq \lambda_{\ker(\tilde{B})}(D)\sqrt{\frac{7}{16}}(\sqrt{7/4} - 1) > \lambda_{\ker(\tilde{B})}(D)/8.$$

Summarizing, we have that

$$\lambda_{\min}(D + \eta\tilde{B}^T\tilde{B}) > \lambda_{\ker(\tilde{B})}(D) \min\left(1/8, \eta\sigma_{\tilde{B}}^2/8\lambda_{\max}(D)\right). \quad (99)$$

□

Proposition 30. For any $\mathfrak{s} \in \mathfrak{S}$, let \mathcal{M} be the admissible Riccati set associated with \mathfrak{s} as defined in (72). For any $\mu \in \mathcal{M}$, let \tilde{K}_μ be the optimal controller given in (67), P_μ , A_μ^c and D_μ be given in (66) and for any \tilde{K} , let $P(\tilde{K})$, $G(\tilde{K})$ and $P_\mu(\tilde{K})$ be respectively the matrix representation of the objective, the constraint and the Lagrangian, defined in (61). Then,

1. for any $(\mu_1, \mu_2) \in \mathcal{M}^2$, $(\mu_1 - \mu_2)(G(\tilde{K}_{\mu_2}) - G(\tilde{K}_{\mu_1})) \succcurlyeq 0$,
2. for any $v \in \mathbb{R}^{n+d}$, the function $\mu \in \mathcal{M} \rightarrow v^T D_\mu v$ is quasi-concave,
3. for any $\mu \in \mathcal{M}$, $\lambda_{\max}(D_\mu) \leq c(\mu)$, where $c(\mu) = (\lambda_{\max}(C_\dagger) + \mu)(1 + \|\tilde{B}\|_2^2(1 + \|A\|_2^2))$.

Proof of Prop. 30. 1. From (87), we have that

$$\begin{aligned} P_{\mu_1}(\tilde{K}_{\mu_2}) - P_{\mu_1} &= (A_{\mu_2}^c)^T (P_{\mu_1}(\tilde{K}_{\mu_2}) - P_{\mu_1}) A_{\mu_2}^c + (\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})^T D_{\mu_1} (\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2}) \\ P_{\mu_2}(\tilde{K}_{\mu_1}) - P_{\mu_2} &= (A_{\mu_1}^c)^T (P_{\mu_2}(\tilde{K}_{\mu_1}) - P_{\mu_2}) A_{\mu_1}^c + (\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})^T D_{\mu_2} (\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2}) \end{aligned}$$

which since $P_{\mu_1}(\tilde{K}_{\mu_2}) - P_{\mu_1} \succcurlyeq 0$, $P_{\mu_2}(\tilde{K}_{\mu_1}) - P_{\mu_2} \succcurlyeq 0$, leads to

$$\begin{aligned} P_{\mu_1}(\tilde{K}_{\mu_2}) - P_{\mu_1} &\succcurlyeq (\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})^T D_{\mu_1} (\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2}), \\ P_{\mu_2}(\tilde{K}_{\mu_1}) - P_{\mu_2} &\succcurlyeq (\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})^T D_{\mu_2} (\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2}). \end{aligned} \quad (100)$$

From (62), we have that

$$\begin{aligned} P_{\mu_1}(\tilde{K}_{\mu_2}) &= P(\tilde{K}_{\mu_2}) + \mu_1 G(\tilde{K}_{\mu_2}) = P_{\mu_2}(\tilde{K}_{\mu_2}) + (\mu_1 - \mu_2) G(\tilde{K}_{\mu_2}) = P_{\mu_2} + (\mu_1 - \mu_2) G(\tilde{K}_{\mu_2}), \\ P_{\mu_2}(\tilde{K}_{\mu_1}) &= P(\tilde{K}_{\mu_1}) + \mu_2 G(\tilde{K}_{\mu_1}) = P_{\mu_1}(\tilde{K}_{\mu_1}) + (\mu_2 - \mu_1) G(\tilde{K}_{\mu_1}) = P_{\mu_1} + (\mu_2 - \mu_1) G(\tilde{K}_{\mu_1}). \end{aligned}$$

Combining these with (100) and summing the two inequalities, one obtains:

$$(\mu_1 - \mu_2)(G(\tilde{K}_{\mu_2}) - G(\tilde{K}_{\mu_1})) \succcurlyeq (\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})^T (D_{\mu_1} + D_{\mu_2})(\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2}) \succcurlyeq 0.$$

2. For any $v \in \mathbb{R}^{n+d}$, let $f : \mu \rightarrow v^T D_\mu v$. Recalling that $D_\mu = R_\mu + \tilde{B}^T P_\mu \tilde{B}$ where $R_\mu = R_0 + \mu dR$ (with R_0 constructed from C_\dagger and dR constructed from C_g), we have that $f'(\mu) = v^T (dR + \tilde{B}^T dP_\mu \tilde{B}) v$ where dP_μ is the (matrix) derivative of P_μ w.r.t. μ . Direct matrix differentiation of (66) shows, after tedious algebraic manipulations, that $dP_\mu = G(\tilde{K}_\mu)$. As a result,

$$f'(\mu) = v^T (dR + \tilde{B}^T G(\tilde{K}_\mu) \tilde{B}) v,$$

which, according to the first statement is a decreasing function. Thus, f can be either increasing then decreasing, increasing or decreasing, which is enough to ensure quasi-concavity over $\mathcal{M} \subset \mathbb{R}_+$.

3. For any $\mu \in \mathcal{M}$, $P_\mu \preceq P_\mu(\tilde{K})$ for any \tilde{K} . Thus, for $\bar{K} = \begin{pmatrix} 0 \\ -A \end{pmatrix}$ that is such that $A^c(\bar{K}) = 0$, one has:

$$P_\mu \preceq P_\mu(\bar{K}) = \begin{pmatrix} I \\ \bar{K} \end{pmatrix}^\top C_\mu \begin{pmatrix} I \\ \bar{K} \end{pmatrix} \preceq \lambda_{\max}(C_\mu) \|\tilde{B}\|_2^2 (I + A^\top A).$$

Notice that $\lambda_{\max}(C_\mu) \geq \mu \geq 0$. From $D_\mu = R_\mu + \tilde{B}^\top P_\mu \tilde{B}$, and $\lambda_{\max}(R_\mu) \leq \lambda_{\max}(C_\mu)$ we obtain:

$$D_\mu \preceq \lambda_{\max}(C_\mu) \left(I + \|\tilde{B}\|_2^2 (I + A^\top A) \right) \Rightarrow \lambda_{\max}(D_\mu) \leq \lambda_{\max}(C_\mu) (1 + \|\tilde{B}\|_2^2 (1 + \|A\|_2^2)).$$

The proof concludes by noting that $\lambda_{\max}(C_\mu) \leq \lambda_{\max}(C_\dagger) + \mu$.

□

H.3. Smoothness of \mathcal{D}

We conclude the characterization of the dual function \mathcal{D} discussing its smoothness, that is, its Lipschitz and gradient Lipschitz continuity. Indeed, since the objective is ultimately to maximize \mathcal{D} over \mathcal{M} , one has to first guarantee that the function behaves 'nicely'. Unfortunately, Prop. 26 shows that the closed-loop matrix may be arbitrarily close-to-unstable when $\mu \rightarrow \tilde{\mu}$, that is $\rho(A_\mu^c) \rightarrow 1$. This result in $\mathcal{D}'(\mu, \mathfrak{s}) \rightarrow -\infty$ and indicates that \mathcal{D} cannot be smooth *everywhere* on \mathcal{M} . Nevertheless, we show that \mathcal{D} is still smooth on a subset of \mathcal{M} , which will be sufficient to provide a practical algorithm to maximize the dual function.

Lemma 14. *For any $\mathfrak{s} \in \mathfrak{S}$, let $\mathcal{M} = [0, \tilde{\mu})$ be the admissible Riccati set associated with \mathfrak{s} as defined in (72). Let \mathcal{M}_+ be a subset of \mathcal{M} such that*

$$\mathcal{M}_+ = \{\mu \in \mathcal{M} \text{ s.t. } \mathcal{D}'(\mu, \mathfrak{s}) \geq 0\}.$$

Then, $\mathcal{D}(\cdot, \mathfrak{s})$ has Lipschitz gradient, i.e., for all $(\mu_1, \mu_2) \in \mathcal{M}_+^2$,

$$|\mathcal{D}'(\mu_1, \mathfrak{s}) - \mathcal{D}'(\mu_2, \mathfrak{s})| \leq |\mu_1 - \mu_2| \frac{\|C_g\|_2 \alpha_D(\mathfrak{s})}{2 \max(\lambda_{\min}(D_{\mu_1}), \lambda_{\min}(D_{\mu_2}))},$$

where $\alpha_D(\mathfrak{s}) := 8 \|C_g\|_2 \kappa(\mathfrak{s})^4 ((2 + \|A\|_2 \|B\|_2)(1 + \|B\|_2))^2$.

Proof of Lem. 14. For any $\mu \in \mathcal{M}$, let \tilde{K}_μ , A_μ^c , D_μ and P_μ be the optimal quantities at μ associated with the Riccati equation for \mathfrak{s} . Without loss of generality, we assume that $\mu_2 \geq \mu_1$. To prove Lem. 14, we first show $\mu \rightarrow P_\mu$ is Lipschitz and that $\left\| \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix} \right\|_2$ is bounded over \mathcal{M}_+ . Then, we show that \mathcal{D} has Lipschitz gradient on \mathcal{M}_+ .

1. From (62), we have that

$$\begin{aligned} P_{\mu_2} - P_{\mu_1} &\preceq P_{\mu_2}(\tilde{K}_{\mu_1}) - P_{\mu_1} = P(\tilde{K}_{\mu_1}) + \mu_2 G(\tilde{K}_{\mu_1}) - P_{\mu_1} = (\mu_2 - \mu_1) G(\tilde{K}_{\mu_1}), \\ P_{\mu_1} - P_{\mu_2} &\preceq P_{\mu_1}(\tilde{K}_{\mu_2}) - P_{\mu_2} = P(\tilde{K}_{\mu_2}) + \mu_1 G(\tilde{K}_{\mu_2}) - P_{\mu_2} = (\mu_1 - \mu_2) G(\tilde{K}_{\mu_2}). \end{aligned}$$

Thus, $(\mu_2 - \mu_1) G(\tilde{K}_{\mu_2}) \preceq P_{\mu_2} - P_{\mu_1} \preceq (\mu_2 - \mu_1) G(\tilde{K}_{\mu_1})$. Decomposing the positive and negative definite part of C_g , we obtain that for any $\mu \in \mathcal{M}_+$,

$$G(\tilde{K}_\mu) \succcurlyeq -X, \text{ where } X = (A_{\mu_1}^c)^\top X A_{\mu_1}^c + \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix}^\top \begin{pmatrix} V^{-1} & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix}.$$

Further, standard Lyapunov algebra ensures that $\lambda_{\max}(X) \leq \text{Tr}(X) \leq \frac{\lambda_{\max}(V^{-1})}{\lambda_{\min}(C_\dagger)} \text{Tr}(P(\tilde{K}_\mu))$. Since $\mu \in \mathcal{M}_+$, $\text{Tr}(P(\tilde{K}_\mu)) = \mathcal{J}_{\tilde{\pi}_\mu}(\mathfrak{s}) \leq \mathcal{D}(\mu, \mathfrak{s}) \leq \mathcal{J}_*(\mathfrak{s})$ by weak duality and $g_{\tilde{\pi}}(\mathfrak{s}) \geq 0$. Finally, $\lambda_{\max}(V^{-1}) \leq \|C_g\|_2$ leads to

$$\lambda_{\min}(G(\tilde{K}_\mu)) \geq -\|C_g\|_2 \kappa(\mathfrak{s}). \quad (101)$$

To prove an upper bound on $\lambda_{\max}(G(\tilde{K}_\mu))$, we use Prop. 30 which implies that $G(\tilde{K}_\mu) \preceq G(\tilde{K}_0)$. \tilde{K}_0 is the optimal controller at $\mu = 0$, and one can check that it is such that $A_0^c = 0$ (by algebraic manipulations). Hence, decomposing the positive and negative definite part of C_g , we obtain that for any $\mu \in \mathcal{M}_+$,

$$G(\tilde{K}_\mu) \preceq G(\tilde{K}_0) \preceq \begin{pmatrix} I \\ \tilde{K}_0 \end{pmatrix}^\top \begin{pmatrix} 0_{n+d} & 0 \\ 0 & I_n \end{pmatrix} \begin{pmatrix} I \\ \tilde{K}_0 \end{pmatrix}.$$

However, for all $\mu \in \mathcal{M}_+$, we have by Prop. 25 that $\|A_\mu^c\|_2 \leq \sqrt{\kappa(\mathfrak{s})}$ and $\left\| \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix} \right\|_2 \leq \sqrt{\kappa(\mathfrak{s})}(2 + \|A\|_2\|B\|_2)$.

Thus,

$$\lambda_{\max}(G(\tilde{K}_\mu)) \preceq \sqrt{\kappa(\mathfrak{s})}(2 + \|A\|_2\|B\|_2). \quad (102)$$

Summarizing, Eqs. (101) and (102) leads to

$$\|P_{\mu_1} - P_{\mu_2}\|_2 \leq |\mu_1 - \mu_2| \sqrt{\kappa(\mathfrak{s})} \max\left(\|C_g\|_2 \sqrt{\kappa(\mathfrak{s})}, (2 + \|A\|_2\|B\|_2)\right). \quad (103)$$

2. We now show that \mathcal{D} has Lipschitz gradient on \mathcal{M}_+ . From (87), we have that

$$\begin{aligned} P_{\mu_1}(\tilde{K}_{\mu_2}) - P_{\mu_1} &= (A_{\mu_2}^c)^\top (P_{\mu_1}(\tilde{K}_{\mu_2}) - P_{\mu_1}) A_{\mu_2}^c + (\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})^\top D_{\mu_1} (\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2}), \\ P_{\mu_2}(\tilde{K}_{\mu_1}) - P_{\mu_2} &= (A_{\mu_1}^c)^\top (P_{\mu_2}(\tilde{K}_{\mu_1}) - P_{\mu_2}) A_{\mu_1}^c + (\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})^\top D_{\mu_2} (\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2}). \end{aligned}$$

Thus, standard Lyapunov manipulation shows that

$$\begin{aligned} 0 &\leq \text{Tr}\left(P_{\mu_1}(\tilde{K}_{\mu_2}) - P_{\mu_1}\right) \leq \|D_{\mu_1}^{1/2}(\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})\|_2^2 \text{Tr}(\Sigma_{\mu_2}), \\ 0 &\leq \text{Tr}\left(P_{\mu_2}(\tilde{K}_{\mu_1}) - P_{\mu_2}\right) \leq \|D_{\mu_2}^{1/2}(\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})\|_2^2 \text{Tr}(\Sigma_{\mu_1}), \end{aligned}$$

where Σ_{μ_1} and Σ_{μ_2} are the steady-state covariance matrix of the state processes driven by $A_{\mu_1}^c$ and $A_{\mu_2}^c$ respectively. Further, since $(\mu_1, \mu_2) \in \mathcal{M}_+^2$, Prop. 24 ensures that $\text{Tr}(\Sigma_{\mu_1}) \leq \kappa(\mathfrak{s})$ and $\text{Tr}(\Sigma_{\mu_2}) \leq \kappa(\mathfrak{s})$.

As a result,

$$\begin{aligned} 0 &\leq \text{Tr}\left(P_{\mu_1}(\tilde{K}_{\mu_2}) - P_{\mu_1}\right) \leq \kappa(\mathfrak{s}) \|D_{\mu_1}^{1/2}(\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})\|_2^2, \\ 0 &\leq \text{Tr}\left(P_{\mu_2}(\tilde{K}_{\mu_1}) - P_{\mu_2}\right) \leq \kappa(\mathfrak{s}) \|D_{\mu_2}^{1/2}(\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})\|_2^2. \end{aligned} \quad (104)$$

From (62), we have that

$$\begin{aligned} P_{\mu_1}(\tilde{K}_{\mu_2}) &= P(\tilde{K}_{\mu_2}) + \mu_1 G(\tilde{K}_{\mu_2}) = P_{\mu_2}(\tilde{K}_{\mu_2}) + (\mu_1 - \mu_2) G(\tilde{K}_{\mu_2}) = P_{\mu_2} + (\mu_1 - \mu_2) G(\tilde{K}_{\mu_2}), \\ P_{\mu_2}(\tilde{K}_{\mu_1}) &= P(\tilde{K}_{\mu_1}) + \mu_2 G(\tilde{K}_{\mu_1}) = P_{\mu_1}(\tilde{K}_{\mu_1}) + (\mu_2 - \mu_1) G(\tilde{K}_{\mu_1}) = P_{\mu_1} + (\mu_2 - \mu_1) G(\tilde{K}_{\mu_1}). \end{aligned}$$

Combining the later identities with (104) and summing the two inequalities we obtain:

$$|\mu_1 - \mu_2| \left| \text{Tr}\left(G(\tilde{K}_{\mu_2}) - G(\tilde{K}_{\mu_1})\right) \right| \leq \kappa(\mathfrak{s}) \|(D_{\mu_1} + D_{\mu_2})^{1/2}(\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})\|_2^2$$

Finally, from

$$\begin{aligned} \|(D_{\mu_1} + D_{\mu_2})^{1/2}(\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})\|_2^2 &\leq \lambda_{\max}((D_{\mu_1} + D_{\mu_2})^{-1}) \|(D_{\mu_1} + D_{\mu_2})(\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})\|_2^2 \\ &\leq \frac{1}{\lambda_{\min}(D_{\mu_1} + D_{\mu_2})} \|(D_{\mu_1} + D_{\mu_2})(\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})\|_2^2 \\ &\leq \frac{1}{\max(\lambda_{\min}(D_{\mu_1}), \lambda_{\min}(D_{\mu_2}))} \|(D_{\mu_1} + D_{\mu_2})(\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})\|_2^2, \end{aligned}$$

one has, noticing that $(\mathcal{D}'(\mu_2, \mathfrak{s}) - \mathcal{D}'(\mu_1, \mathfrak{s})) = \text{Tr}\left(G(\tilde{K}_{\mu_2}) - G(\tilde{K}_{\mu_1})\right)$ (see Proof of Prop. 30),

$$|\mu_1 - \mu_2| |\mathcal{D}'(\mu_2, \mathfrak{s}) - \mathcal{D}'(\mu_1, \mathfrak{s})| \leq \frac{\kappa(\mathfrak{s})}{\max(\lambda_{\min}(D_{\mu_1}), \lambda_{\min}(D_{\mu_2}))} \|(D_{\mu_1} + D_{\mu_2})(\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})\|_2^2. \quad (105)$$

(105) indicates that the smoothness of \mathcal{D}' is tied to the one of $\mu \rightarrow \tilde{K}_\mu$. Using the expression of \tilde{K}_μ in Prop. 21, we have $D_{\mu_1} \tilde{K}_{\mu_1} = -\tilde{B}^\top P_{\mu_1} A - N_{\mu_1}$ and $D_{\mu_2} \tilde{K}_{\mu_2} = -\tilde{B}^\top P_{\mu_2} A - N_{\mu_2}$. Combining those two expressions leads to

$$\begin{aligned} D_{\mu_1}(\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2}) &= \tilde{B}^\top (P_{\mu_2} - P_{\mu_1})A + (N_{\mu_2} - N_{\mu_1}) + (D_{\mu_2} - D_{\mu_1})\tilde{K}_{\mu_2}, \\ D_{\mu_2}(\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2}) &= \tilde{B}^\top (P_{\mu_2} - P_{\mu_1})A + (N_{\mu_2} - N_{\mu_1}) + (D_{\mu_2} - D_{\mu_1})\tilde{K}_{\mu_1}, \\ (D_{\mu_1} + D_{\mu_2})(\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2}) &= 2\tilde{B}^\top (P_{\mu_2} - P_{\mu_1})A + 2(N_{\mu_2} - N_{\mu_1}) + (D_{\mu_2} - D_{\mu_1})(\tilde{K}_{\mu_1} + \tilde{K}_{\mu_2}). \end{aligned}$$

Since $D_\mu = R_\mu + \tilde{B}^\top P_\mu \tilde{B}$ and $2A + \tilde{B}(\tilde{K}_{\mu_1} + \tilde{K}_{\mu_2}) = A_{\mu_1}^c + A_{\mu_2}^c$ we obtain

$$(D_{\mu_1} + D_{\mu_2})(\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2}) = \tilde{B}^\top (P_{\mu_2} - P_{\mu_1})(A_{\mu_1}^c + A_{\mu_2}^c) + 2(N_{\mu_2} - N_{\mu_1}) + (R_{\mu_2} - R_{\mu_1})(\tilde{K}_{\mu_1} + \tilde{K}_{\mu_2})$$

From Prop. 25 and (103), we have

$$\left\| \tilde{B}^\top (P_{\mu_2} - P_{\mu_1})(A_{\mu_1}^c + A_{\mu_2}^c) \right\|_2 \leq 2|\mu_1 - \mu_2|(1 + \|B\|_2)\kappa(\mathfrak{s}) \max\left(\|C_g\|_2 \sqrt{\kappa(\mathfrak{s})}, (2 + \|A\|_2 \|B\|_2)\right).$$

Further, the linear structure of C_μ leads to

$$(N_{\mu_2} - N_{\mu_1}) + (R_{\mu_2} - R_{\mu_1})(\tilde{K}_{\mu_1} + \tilde{K}_{\mu_2}) = (\mu_2 - \mu_1) \begin{pmatrix} 0 \\ I_n \end{pmatrix} C_g \left(\begin{pmatrix} I \\ \tilde{K}_{\mu_1} \end{pmatrix} + \begin{pmatrix} I \\ \tilde{K}_{\mu_2} \end{pmatrix} \right).$$

Combined with Prop. 25, this leads to

$$\left\| (N_{\mu_2} - N_{\mu_1}) + (R_{\mu_2} - R_{\mu_1})(\tilde{K}_{\mu_1} + \tilde{K}_{\mu_2}) \right\|_2 \leq 2(\mu_2 - \mu_1) \|C_g\|_2 \sqrt{\kappa(\mathfrak{s})} (2 + \|A\|_2 \|B\|_2).$$

Putting everything together, we have

$$\|(D_{\mu_1} + D_{\mu_2})(\tilde{K}_{\mu_1} - \tilde{K}_{\mu_2})\|_2 \leq |\mu_1 - \mu_2| \kappa(\mathfrak{s})^{3/2} c(\mathfrak{s}), \text{ where } c(\mathfrak{s}) := 2\|C_g\|_2 (1 + \|B\|_2) (2 + \|A\|_2 \|B\|_2).$$

Finally, using this upper bound in (105) leads to

$$|\mathcal{D}'(\mu_2, \mathfrak{s}) - \mathcal{D}'(\mu_1, \mathfrak{s})| \leq |\mu_1 - \mu_2| \frac{\|C_g\|_2 \alpha_D(\mathfrak{s})}{2 \max(\lambda_{\min}(D_{\mu_1}), \lambda_{\min}(D_{\mu_2}))},$$

where $\alpha_D(\mathfrak{s}) := 8\|C_g\|_2 \kappa(\mathfrak{s})^4 ((2 + \|A\|_2 \|B\|_2)(1 + \|B\|_2))^2$ is the same problem dependent constant as in Cor. 2.

□

I. Algorithms and complexity bound

In this section, we provide the pseudo-code of OFU-LQ++ and LAGLQ. We further detail the procedure of DS-OFU which allows us to solve the Extended LQR with Relaxed Constraint efficiently by Dichotomy Search and then prove Thm. 3.

I.1. OFU-LQ++ and LAGLQ

Input: $\Theta_0, T, \delta, Q, R, D, \sigma^2$.

- 1: Set λ by (106), $V_0 = \lambda I$, $\hat{\theta}_0$ and \mathcal{C}_0 by (107).
- 2: $\theta_0 = \arg \min_{\theta \in \mathcal{C}_0} J(\theta)$.
- 3: Compute optimistic controller $K_0 = K(\hat{\theta}_0)$ by (108).
- 4: **for** $t=1, \dots, T$ **do**
- 5: **if** $\det(V_t) \geq 2 \det(V_0)$ **then**
- 6: Compute $\hat{\theta}_t$ and \mathcal{C}_t by (107).
- 7: Find $\theta_t = \arg \min_{\theta \in \mathcal{C}_t} J(\theta) + \frac{1}{\sqrt{t}}$.
- 8: Compute optimistic controller $K_t = K(\theta_t)$ by (108).
- 9: Let $V_0 = V_t$.
- 10: **else**
- 11: $K_t = K_{t-1}$.
- 12: **end if**
- 13: Compute control u_t based on the current controller K_t as $u_t = K_t x_t$.
- 14: Execute control u_t and observe next state x_{t+1} .
- 15: Save (z_t, x_{t+1}) in the dataset, where $z_t^\top = (x_t^\top, u_t^\top)$.
- 16: Update $V_{t+1} = V_t + z_t z_t^\top$.
- 17: **end for**

Figure 4. OFU-LQ++ algorithm

Input: $\Theta_0, T, \delta, Q, R, D, \sigma^2$.

- 1: Set λ by (106), $V_0 = \lambda I$, $\hat{\theta}_0, \beta_0$ and \mathcal{C}_0 by (107).
- 2: Compute the extended controller $\tilde{K}_0 = \text{DS-OFU}()$.
- 3: Extract $K_0 = [\tilde{K}_0]_{1:d,1:n}$
- 4: **for** $t=1, \dots, T$ **do**
- 5: **if** $\det(V_t) \geq 2 \det(V_0)$ **then**
- 6: Compute $\hat{\theta}_t, \beta_t$ and \mathcal{C}_t by (107).
- 7: Compute the extended controller $\tilde{K}_t = \text{DS-OFU}()$.
- 8: Extract $K_t = [\tilde{K}_t]_{1:d,1:n}$
- 9: Let $V_0 = V_t$.
- 10: **else**
- 11: $K_t = K_{t-1}$.
- 12: **end if**
- 13: Compute control u_t based on the current controller K_t as $u_t = K_t x_t$.
- 14: Execute control u_t and observe next state x_{t+1} .
- 15: Save (z_t, x_{t+1}) in the dataset, where $z_t^\top = (x_t^\top, u_t^\top)$.
- 16: Update $V_{t+1} = V_t + z_t z_t^\top$.
- 17: **end for**

Figure 5. LAGLQ algorithm

OFU-LQ++ and LAGLQ builds on the original OFU-LQ algorithm of Abbasi-Yadkori & Szepesvári (2011) but differ in the initialization (to ensure that the state remains bounded over the trajectory) and in the way the optimistic controller is computed (for LAGLQ).

Initialization. OFU-LQ and LAGLQ are provided with the LQR costs matrices Q, R , the upper bound D on $\text{Tr}(P_*)$ (see Asm. 3), the proxy-variance upper bound σ^2 (see Asm. 1), the time horizon T and the confidence level δ as well as the stabilizing set $\Theta_0 = \{\theta : \|\theta - \theta_0\| \leq \epsilon_0\}$, where ϵ_0 satisfies the requirement of Lem. 1 for OFU-LQ++ and (18) for LAGLQ. Formally, we require that

OFU-LQ++: $\frac{2\kappa\eta(2\epsilon_0)}{1-\eta(2\epsilon_0)} \leq 1$ and $\eta(2\epsilon_0) := 8\kappa\epsilon_0(1+2\epsilon_0) < 1$ which is guaranteed as soon as $\epsilon_0 \leq 1/(64\kappa^2)$,

LAGLQ: $\frac{2\kappa\eta(2\sqrt{\kappa}\epsilon_0)}{1-\eta(2\sqrt{\kappa}\epsilon_0)} \leq 1$ and $\eta(2\sqrt{\kappa}\epsilon_0) := 8\kappa\sqrt{\kappa}\epsilon_0(1+2\sqrt{\kappa}\epsilon_0) < 1$ which is guaranteed as soon as $\epsilon_0 \leq 1/(64\kappa^{5/2})$,

where $\kappa = D/\lambda_{\min}(C)$, $C = \begin{pmatrix} Q & 0 \\ 0 & R \end{pmatrix}$. The stabilizing set can be obtained in finite time following the procedure of Faradonbeh et al. (2018a) or following the procedure of Simchowit & Foster (2020) if a stabilizing controller K_0 is known a priori. Finally, the stabilizing set Θ_0 is used to regularize the RLS estimates as

$$\lambda = \frac{2n^2\sigma^2}{\epsilon_0^2} \left(\log\left(\frac{4T}{\delta}\right) + (n+d) \log(1 + \kappa X^2 T) \right). \quad (106)$$

Recursive Least-Square. After the initial phase, OFU-LQ++ and LAGLQ proceed through episodes following the update rule of OFU-LQ in (Abbasi-Yadkori & Szepesvári, 2011) i.e., update whenever $\det(V_t) \geq 2 \det(V_{t_k})$. At the beginning of each episode, they maintain an estimation of θ_* by RLS together with a confidence set given by:

$$\hat{\theta}_t = V_t^{-1} \left(\lambda \theta_0 + \sum_{s=0}^{t-1} z_s x_{s+1}^\top \right); \quad \beta_t = n\sigma \sqrt{2 \log\left(\frac{\det(V_t)^{1/2} 4T}{\det(\lambda I)^{1/2} \delta}\right)} + \lambda^{1/2} \epsilon_0; \quad \mathcal{C}_t = \{\theta : \|\theta - \hat{\theta}_t\|_{V_t} \leq \beta_t\}. \quad (107)$$

Optimistic controllers. OFU-LQ++ and LAGLQ differs in the way the compute the optimistic controller at each policy update.

OFU-LQ++ follows (Abbasi-Yadkori & Szepesvári, 2011) and has to find an optimistic parameter by a $1/\sqrt{t}$ margin as $\theta_t = \arg \min_{\theta \in \mathcal{C}_t} J(\theta) + \frac{1}{\sqrt{t}}$. Then, it computes the optimal controller associated with $\theta_t^\top = (A_t, B_t)$ by solving a standard Riccati equation:

$$\begin{aligned} P_t &= Q + A_t^\top P_t A_t - A_t^\top P_t B_t (R + B_t^\top P_t B_t)^{-1} B_t^\top P_t A_t, \\ K_t &= -(R + B_t^\top P_t B_t)^{-1} B_t^\top P_t B_t. \end{aligned} \quad (108)$$

Unfortunately, finding the optimistic parameter $\tilde{\theta}_t$ is intractable which makes OFU-LQ++ inefficient.

On the other hand, LAGLQ uses the Extended LQR with Relaxed Constraint approach to compute an extended optimistic controller \tilde{K}_t using the DS-OFU routine (described in detail in the next subsection). Then, it simply extracts the control part of $\tilde{K}_t^\top = (K_t^{u\top} \quad K_t^{w\top})$ to obtain the controller used for the whole episode. The complexity of DS-OFU is at most $O(n^3)$ and only requires solving Riccati equations associated with the extended LQR, which makes LAGLQ very efficient.

I.2. DS-OFU algorithm

In this subsection, we detail the algorithm DS-OFU and the back-up procedure BACKUPPROC used in LAGLQ to solve the Extended LQR with Relaxed Constraint problem (16). We leverage the strong duality results of Thm. 2 to find a solution to (16) by solving the associated Lagrangian extended LQR problem (21). The routines are summarized in figs. 6 and 7.

Input: $Q, R, D, \hat{\theta}, \beta, V, \epsilon,$

- 1: Set μ_{\max} by (110), α, λ_0 by (111).
- 2: **if** $\mathcal{D}(0) \leq 0$ **then**
- 3: Set $\bar{\mu} = 0$ and $\tilde{\pi}_\epsilon = \tilde{\pi}_{\bar{\mu}}$
- 4: **else**
- 5: Set $\mu_l = 0, \mu_r = \mu_{\max}$ (Lem. 7)
- 6: **while** $\alpha \frac{\mu_r - \mu_l}{\lambda_{\min}(D_{\mu_l})} \geq \epsilon$ **or** $\lambda_{\min}(D_{\mu_l}) \geq \lambda_0 \epsilon^2$ **do**
- 7: Set $\bar{\mu} = (\mu_l + \mu_r)/2$
- 8: **if** $\mathcal{D}'(\bar{\mu}) > 0$ **then**
- 9: $\mu_l = \bar{\mu}$
- 10: **else**
- 11: $\mu_r = \bar{\mu}$
- 12: **end if**
- 13: **end while**
- 14: **end if**
- 15: **if** $\alpha \frac{\mu_r - \mu_l}{\lambda_{\min}(D_{\mu_l})} \leq \epsilon$ **then**
- 16: Set $\bar{\mu} = \mu_l$ and $\tilde{\pi}_\epsilon = \tilde{\pi}_{\bar{\mu}}$
- 17: **else**
- 18: Compute $\tilde{\pi}_\epsilon = \text{BACKUPPROC}(Q, R, D, \hat{\theta}, \beta, V, \epsilon, \bar{\mu})$.
- 19: **end if**
- 20: **return** Control policy $\tilde{\pi}_\epsilon$

Figure 6. The DS-OFU algorithm to solve (21).

Input: $Q, R, D, \hat{\theta}, \beta, V, \epsilon, \bar{\mu}.$

- 1: Compute $D_{\bar{\mu}}$ by (109)
- 2: **if** $\min_{\|v\|=1, v \in \ker(\tilde{B})} (D_{\bar{\mu}}) \leq \sqrt{\lambda_0} \epsilon$ **then**
- 3: Compute \tilde{K}_ϵ by (112).
- 4: Set $\tilde{\pi}_\epsilon(x) = \tilde{K}_\epsilon x$.
- 5: **else**
- 6: Modify the LQR system according to (113).
- 7: Set $\mu_l = 0, \mu_r = \bar{\mu}, \alpha_{\text{mod}}$ by (114).
- 8: **while** $\alpha_{\text{mod}}(\mu_r - \mu_l) \geq \epsilon^3$ **do**
- 9: Set $\mu = (\mu_l + \mu_r)/2$
- 10: **if** $\mathcal{D}'_{\text{mod}}(\mu) > 0$ **then**
- 11: $\mu_l = \mu$
- 12: **else**
- 13: $\mu_r = \mu$
- 14: **end if**
- 15: **end while**
- 16: Set $\mu = \mu_l$ and $\tilde{\pi}_\epsilon = \tilde{\pi}_{\mu}^{\text{mod}}$.
- 17: **end if**
- 18: **return** Control policy $\tilde{\pi}_\epsilon$

Figure 7. Back up procedure BACKUPPROC

Extended Lagrangian LQR. DS-OFU aims at finding a solution to (21) up to an ϵ accuracy. It takes as input an estimated LQR instance, parametrized by the cost matrices Q, R and the RLS estimate $\hat{\theta}^\top = (\hat{A}, \hat{B})$ as well as the associated confidence set given by V and β . Then, for each Lagrangian parameters μ , it constructs the Lagrangian extended LQR system (A, \tilde{B}, C_μ) where $C_\mu = C_\dagger + \mu C_g$,

$$A = \hat{A}; \quad \tilde{B} = (\hat{B}, I); \quad C_\mu = \begin{pmatrix} Q_\mu & N_\mu \\ N_\mu^\top & R_\mu \end{pmatrix}; \quad C_\dagger = \begin{pmatrix} Q & 0 & 0 \\ 0 & R & 0 \\ 0 & 0 & 0 \end{pmatrix}; \quad C_g = \begin{pmatrix} \beta^2 V^{-1} & 0 \\ 0 & I \end{pmatrix}.$$

Lem. 5 ensures that whenever μ is in some admissible set \mathcal{M} , the dual function $\mathcal{D}(\mu)$ and its derivative $\mathcal{D}'(\mu)$ associated with the estimated extended LQR can be computed by Riccati and Lyapunov equations as $\mathcal{D}(\mu) = \text{Tr}(P_\mu)$ and $\mathcal{D}'(\mu) =$

$\text{Tr}(G(\tilde{K}_\mu))$ where

$$\begin{aligned} P_\mu &= Q_\mu + A^\top P_\mu A - (A^\top P_\mu \tilde{B} + N_\mu^\top) D_\mu^{-1} (\tilde{B}^\top P_\mu A + N_\mu) \\ D_\mu &= R_\mu + \tilde{B}^\top P_\mu \tilde{B}; \quad \tilde{K}_\mu = -D_\mu^{-1} (\tilde{B}^\top P_\mu A + N_\mu); \quad A_\mu^c = A + \tilde{B} \tilde{K}_\mu, \\ G(\tilde{K}_\mu) &= (A_\mu^c)^\top G(\tilde{K}_\mu) A_\mu^c + \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix}^\top C_g \begin{pmatrix} I \\ \tilde{K}_\mu \end{pmatrix}. \end{aligned} \quad (109)$$

Whenever $\mu \notin \mathcal{M}$ (or equivalently when (109) does not admit a well-defined solution in the sense of Lem. 10), we simply set $\mathcal{D}'(\mu) = -\infty$. Finally, since all the involved policies are linear in the state, we use either $\tilde{\pi}$ or the linear controller \tilde{K} to characterize them.

Initialization. DS-OFU performs a dichotomy search over some range $[0, \mu_{\max}]$ to find an optimistic feasible extended policy. Before starting the dichotomy, DS-OFU computes the upper bound μ_{\max} by (110) and Lem. 7 ensures that $\mathcal{M} \subset [0, \mu_{\max}]$

$$\mu_{\max} = \beta^{-2} \lambda_{\max}(C) \lambda_{\max}(V). \quad (110)$$

Finally, DS-OFU uses the upper bound D on $\text{Tr}(P_*)$ to compute *conservative* constants that will be used to set the termination rules.

$$\begin{aligned} \alpha &= \max(1, \|C_g\|_2/2) 8 \|C_g\| \kappa^4 \left((2 + \|\hat{A}\|_2 \|\hat{B}\|_2) (1 + \|\hat{B}\|_2) \right)^2, \quad \kappa = D/\lambda_{\min}(C) \\ \lambda_0 &= \min \left(\frac{\lambda_{\min}(C)}{2 \|\tilde{B}\|_2^2 \max(D, 1)}, \frac{1}{8^{2n+1} \kappa^{2n}} \min \left(1, \frac{\min(1, \lambda_{\min}(C)/2\kappa) \sigma_{\tilde{B}}^2}{2\kappa^2 c(\mu_{\max})} \right) \right)^2, \\ \sigma_{\tilde{B}}^2 &= \min_{\|v\|=1, v \in \text{Im}(\tilde{B})} v^\top \tilde{B}^\top \tilde{B} v; \quad c(\mu_{\max}) = (\lambda_{\max}(C) + \mu_{\max}) (1 + \|\tilde{B}\|_2^2 (1 + \|A\|_2^2)). \end{aligned} \quad (111)$$

Dichotomy Search. DS-OFU proceeds with the dichotomy search until $\alpha \frac{\mu_r - \mu_l}{\lambda_{\min}(D_{\mu_l})} \leq \epsilon$ or $\lambda_{\min}(D_{\mu_l}) \leq \lambda_0 \epsilon^2$. When the first condition is met, a valid solution $\tilde{\pi}_*$ is found by the dichotomy search. When the second condition is met, the dichotomy search failed to find a valid solution, and we use the *backup procedure* BACKUPPROC.

BACKUPPROC. When the back-up procedure starts, it means that $\lambda_{\min}(D_{\bar{\mu}}) \leq \lambda_0 \epsilon^2$. BACKUPPROC then distinguishes between two cases, depending on the spectrum of $D_{\bar{\mu}}$ w.r.t. the null space of \tilde{B} .

- When $\min_{\|v\|=1, v \in \ker(\tilde{B})} (D_{\bar{\mu}}) \leq \sqrt{\lambda_0} \epsilon$, a feasible optimistic linear policy can be computed explicitly as:

$$\tilde{K}_\epsilon = \tilde{K}_{\bar{\mu}} + \eta v x^\top, \quad (112)$$

where

$$\eta^2 = \frac{\mathcal{D}'(\bar{\mu})}{-v^\top Z v x^\top \Sigma_{\bar{\mu}} x}; \quad Z = \begin{pmatrix} 0 \\ I_{n+d} \end{pmatrix}^\top C_g \begin{pmatrix} 0 \\ I_{n+d} \end{pmatrix}; \quad Y = \begin{pmatrix} 0 \\ I_{n+d} \end{pmatrix}^\top C_g \begin{pmatrix} I \\ \tilde{K}_{\bar{\mu}} \end{pmatrix}; \quad \Sigma_{\bar{\mu}} = (A_{\bar{\mu}}^c)^\top \Sigma_{\bar{\mu}} A_{\bar{\mu}}^c + I.$$

and v is an eigenvector associated with $\lambda_{\min}(D_{\bar{\mu}})$, x is a unitary vector such that $v^\top Y \Sigma_{\bar{\mu}} x = 0$.

- When $\min_{\|v\|=1, v \in \ker(\tilde{B})} (D_{\bar{\mu}}) > \sqrt{\lambda_0} \epsilon$, we construct a modified extended Lagrangian LQR system as:

$$A^{mod} = A; \quad \tilde{B}^{mod} = \tilde{B}; \quad C_g^{mod} = C_g; \quad C_\dagger^{mod} = C_\dagger + \eta \Delta; \quad C_\mu^{mod} = C_\dagger^{mod} + \mu C_g^{mod}, \quad (113)$$

where

$$\begin{aligned} \Delta &= \begin{pmatrix} I & -\bar{K}^\top \\ 0 & I \end{pmatrix} \begin{pmatrix} I \\ -\tilde{B}^\top \end{pmatrix} \begin{pmatrix} I & -\tilde{B} \end{pmatrix} \begin{pmatrix} I & 0 \\ -\bar{K} & I \end{pmatrix}; \quad \bar{K} = \begin{pmatrix} 0 \\ -A \end{pmatrix}, \\ \eta &= \min \left(c(\mu_{\max})/\sigma_{\tilde{B}}^2, \min(1, \lambda_{\min}(C)/2\kappa)/(2\kappa^2) \right) \epsilon. \end{aligned}$$

We then consider \mathcal{D}_{mod} and \mathcal{D}'_{mod} , the dual function and its derivative associated with the modified system, and run a standard dichotomy search over $[0, \bar{\mu}]$ on the modified system. The system modification ensures \mathcal{D}_{mod} to have a bounded curvature. As a result, the termination criterion of the dichotomy search for the modified system is simplified to $\alpha_{mod}(\mu_r - \mu_l) \leq \epsilon^3$, where

$$\alpha_{mod} = \frac{64\|C_g\|_2^2 \kappa^4 ((2 + \|A\|_2 \|B\|_2)(1 + \|B\|_2))^2}{\min((1 + \|B\|_2)^{-2} \lambda_{\min}(C), \sqrt{\lambda_0}/8)}. \quad (114)$$

I.3. Guarantees

Theorem 3. For any LQR parametrized by $\hat{\theta}$, V , β , and psd cost matrices Q , R , and any accuracy $\epsilon \in (0, 1/2)$, whenever $\theta_* \in \mathcal{C}$,

1. DS-OFU outputs an ϵ -optimistic and ϵ -feasible policy $\tilde{\pi}_\epsilon$ given by the linear controller \tilde{K}_ϵ such that

$$\mathcal{J}_{\tilde{\pi}_\epsilon} \leq \mathcal{J}_* + \epsilon \quad \text{and} \quad g_{\tilde{\pi}_\epsilon} \leq \epsilon.$$

2. DS-OFU terminates within at most $N = O(\log(\mu_{\max}/\epsilon))$ iterations, each solving one Riccati and one Lyapunov equation for the extended Lagrangian LQR, both with complexity $O(n^3)$.

Proof of Thm. 3. We first prove that DS-OFU actually finds an ϵ -optimal feasible solution and then characterize its complexity.

Performance guarantees for DS-OFU. The algorithm proceeds in two phases **1)** it performs a dichotomy search on the derivative of the dual function \mathcal{D} associated with original extended LQR problem. **2)** if the solution of the first dichotomy search is not satisfactory, it enters the *backup procedure* BACKUPPROC, and obtain modified valid solution, either explicitly, either performing a dichotomy search for a modified extended LQR problem. The proof is conducted in two steps, one for each phase.

1. The first dichotomy search is done w.r.t. \mathcal{D}' . By Lem. 5, \mathcal{D} is concave over \mathcal{M} , and the initialization procedure ensures by Lem. 7 that $\mathcal{M} \subset [0, \mu_{\max}]$, $\mathcal{D}'(0) > 0$ and $\mathcal{D}'(\mu_{\max}) \leq 0$. As a result, the dichotomy maintains a non-empty interval $[\mu_l, \mu_r]$, which shrinks with the number of iterations. Further, at the end of the dichotomy, we have that

$$\mathcal{D}'(\mu_l) > 0; \quad \mathcal{D}'(\mu_r) \leq 0.$$

- If $\alpha \frac{\mu_r - \mu_l}{\lambda_{\min}(D_{\mu_l})} < \epsilon$, then, strong-duality holds in a 'strict' sense (see Lem. 12) and it exists $\mu_* \in \mathcal{M}$ such that $\mathcal{D}'(\mu_*) = 0$. Further, by construction, $\mu_* \in [\mu_l, \mu_r]$. We prove this assertion by contradiction. Suppose that strong-duality holds in a 'weak-sense', which corresponds to the case where $\mathcal{D}'(\mu) > 0$ for all $\mu \in \mathcal{M}$. Then, $\tilde{\mu} \in [\mu_l, \mu_r]$ and by Cor. 2,

$$\begin{aligned} \lambda_{\min}(D_{\mu_l}) &\leq 8\|C_g\|_2^2 \kappa^4 \left((2 + \|\hat{A}\|_2 \|\hat{B}\|_2)(1 + \|\hat{B}\|_2) \right)^2 |\mu_l - \tilde{\mu}| \\ &\leq 8\|C_g\|_2^2 \kappa^4 \left((2 + \|\hat{A}\|_2 \|\hat{B}\|_2)(1 + \|\hat{B}\|_2) \right)^2 |\mu_l - \mu_r| \\ &\leq \epsilon \frac{\lambda_{\min}(D_{\mu_l})}{\max(1, \|C_g\|_2/2)} \leq \epsilon \lambda_{\min}(D_{\mu_l}). \end{aligned}$$

Since $\epsilon < 1$, we obtain the contradiction.

Now that the existence of $\mu_* \in [\mu_l, \mu_r]$ such that $\mathcal{D}'(\mu_*)$ is asserted, we have from Lem. 14,

$$\begin{aligned} \mathcal{D}'(\mu_l) &= |\mathcal{D}'(\mu_l) - \mathcal{D}'(\mu_*)| \leq \frac{4\|C_g\|^2\kappa^4 \left((2 + \|\widehat{A}\|_2\|\widehat{B}\|_2)(1 + \|\widehat{B}\|_2) \right)^2}{\lambda_{\min}(D_{\mu_l})} |\mu_l - \mu_*| \\ &\leq \frac{4\|C_g\|^2\kappa^4 \left((2 + \|\widehat{A}\|_2\|\widehat{B}\|_2)(1 + \|\widehat{B}\|_2) \right)^2}{\lambda_{\min}(D_{\mu_l})} |\mu_l - \mu_r| \\ &\leq \frac{4\|C_g\|^2\kappa^4 \left((2 + \|\widehat{A}\|_2\|\widehat{B}\|_2)(1 + \|\widehat{B}\|_2) \right)^2}{\alpha} \epsilon \leq \epsilon. \end{aligned}$$

As a result, from $g_{\tilde{\pi}_{\mu_l}} = \mathcal{D}'(\mu_l) \leq \epsilon$, we obtain the ϵ -feasibility. Further, from $g_{\tilde{\pi}_{\mu_l}} = \mathcal{D}'(\mu_l) > 0$, we have by weak-duality,

$$\mathcal{J}_* \geq \mathcal{J}_{\tilde{\pi}_{\mu_l}} + \mu_l g_{\tilde{\pi}_{\mu_l}} > \mathcal{J}_{\tilde{\pi}_{\mu_l}}.$$

- If $\alpha \frac{\mu_r - \mu_l}{\lambda_{\min}(D_{\mu_l})} \geq \epsilon$ we enter the *back-up procedure*. Note that in this case, necessarily, $\lambda_{\min}(D_{\mu_l}) < \lambda_0 \epsilon^2$.

2. If DS-OFU triggers BACKUPPROC, then $\lambda_{\min}(D_{\mu_l}) < \lambda_0 \epsilon^2$.

- If $\min_{\|v\|=1, v \in \ker(\tilde{B})} (D_{\mu}) \leq \sqrt{\lambda_0} \epsilon$, the conditions of Prop. 27 are satisfied. This guarantees that $\tilde{\pi}_{\epsilon}(x) = \tilde{K}_{\epsilon} x$, where \tilde{K}_{ϵ} is obtained by (112), is such that $g_{\tilde{\pi}_{\epsilon}} = 0$ and

$$\mathcal{J}_{\tilde{\pi}_{\epsilon}} = \mathcal{L}_{\tilde{\pi}_{\epsilon}}(\bar{\mu}) \leq \mathcal{D}(\bar{\mu}) + \epsilon \leq \mathcal{J}_* + \epsilon.$$

- If $\min_{\|v\|=1, v \in \ker(\tilde{B})} (D_{\mu}) > \sqrt{\lambda_0} \epsilon$, the conditions of Prop. 28 are satisfied. This guarantees that strong duality holds in a 'strict' sense for the modified LQR system in (113). Formally, we have that

- it exists $\mu_*^{mod} \in [0, \bar{\mu}]$ such that $\mathcal{D}'_{mod}(\mu_*^{mod}) = 0$,
- $\mathcal{D}'_{mod}(0) > 0$, $\mathcal{D}'_{mod}(\bar{\mu}) \leq 0$,
- $\mathcal{J}_{\tilde{\pi}_{\mu_*^{mod}}} \leq \mathcal{J}_* + \epsilon$.

As a result, when the dichotomy on the modified system terminates, we have

$$\begin{aligned} \mathcal{D}'_{mod}(\mu_l) &= |\mathcal{D}'_{mod}(\mu_l) - \mathcal{D}'_{mod}(\mu_*^{mod})| \leq \frac{4\|C_g\|^2\kappa_{mod}^4 \left((2 + \|\widehat{A}\|_2\|\widehat{B}\|_2)(1 + \|\widehat{B}\|_2) \right)^2}{\lambda_{\min}(D_{\mu_l}^{mod})} |\mu_l - \mu_*| \\ &\leq \frac{4\|C_g\|^2\kappa_{mod}^4 \left((2 + \|\widehat{A}\|_2\|\widehat{B}\|_2)(1 + \|\widehat{B}\|_2) \right)^2}{\lambda_{\min}(D_{\mu_l}^{mod})} |\mu_l - \mu_r| \\ &\leq \frac{4\|C_g\|^2\kappa_{mod}^4 \left((2 + \|\widehat{A}\|_2\|\widehat{B}\|_2)(1 + \|\widehat{B}\|_2) \right)^2}{\lambda_{\min}(D_{\mu_l}^{mod})\alpha_{mod}} \epsilon^3, \end{aligned}$$

where κ_{mod} is the constant associated with the modified system and D_{μ}^{mod} is the matrix obtained by (109) on the modified system. Further, one can check that $\kappa_{mod} \leq 2\kappa$ while Prop. 29 ensures that $\lambda_{\min}(D_{\mu_l}^{mod}) \geq \min(\lambda_{\min}(D_0), \sqrt{\lambda_0} \epsilon^2 / 8)$. Using that $\lambda_{\min}(D_0) \geq (1 + \|B\|_2)^{-2} \lambda_{\min}(C)$ and the value of α_{mod} in (114), we obtain by Lem. 14,

$$\begin{aligned} \mathcal{D}'_{mod}(\mu_l) &\leq \frac{4\|C_g\|^2\kappa_{mod}^4 \left((2 + \|\widehat{A}\|_2\|\widehat{B}\|_2)(1 + \|\widehat{B}\|_2) \right)^2}{\lambda_{\min}(D_{\mu_l}^{mod})\alpha_{mod}} \epsilon^3 \\ &\leq \frac{\min((1 + \|B\|_2)^{-2} \lambda_{\min}(C), \sqrt{\lambda_0} / 8)}{\min((1 + \|B\|_2)^{-2} \lambda_{\min}(C), \sqrt{\lambda_0} / 8)} \epsilon \leq \epsilon. \end{aligned}$$

As a result, we have $g_{\tilde{\pi}_{\mu_l}^{mod}} \leq \epsilon$ and $\mathcal{J}_{\tilde{\pi}_{\mu_l}^{mod}} \leq \mathcal{J}_{\tilde{\pi}_{\mu_*^{mod}}} \leq \mathcal{J}_* + \epsilon$.

Complexity of DS-OFU. Since the dichotomy reduces the interval by a factor two at each iterations, we have:

1. The first dichotomy performed by DS-OFU terminates within at most $N_1 = \log_2(\alpha/\lambda_0) + 3 \log_2(1/\epsilon)$ iterations,
2. The back up procedure terminates within at most $N_2 = \log_2(\alpha_{mod}) + 2 \log_2(1/\epsilon)$ iterations.

Overall, DS-OFU performs at most $N = N_1 + N_2 = O(\log_2(1/\epsilon))$ iterations. Finally, each iteration require solving a Lyapunov and a Riccati equation for the Extended Lagrangian LQR problem, that can be solved efficiently in $O(n^3)$, where n is the dimension of the state variable (see (Van Dooren, 1981) for instance).

□

J. Experiments

We detail in this section the experiments that compare the performance of DS-OFU to an ϵ -greedy approach.

LQR system. The experiments are conducted for following LQR system, with $n = d = 2$:

$$A_* = \begin{pmatrix} 1.01 & 0.01 \\ 0.01 & 0.5 \end{pmatrix}; \quad B_* = I; \quad Q = I; \quad R = I. \quad (115)$$

The LQR system is similar to the one used in (Dean et al., 2018) but with smaller dimension ($n = 2$ vs $n = 3$) to simplify the experiment. As in (Dean et al., 2018), the system is marginally unstable as $\lambda(A_*) \approx (1.01, 0.5)$ i.e., has one unstable and one stable mode.

Baseline. Our experiment is not intended to provide an extensive comparison between adaptive LQR algorithms and their heuristic variations. We rather focused on developing an experiment that could complement the theoretical understanding we currently have. As a result, we considered comparing DS-OFU with algorithms that **1**) offer a similar $\tilde{O}(\sqrt{T})$ frequentist regret guarantee **2**) are provably efficient (i.e., tractable with theoretical guarantees on the fact that the chosen controller is indeed the one that the algorithm should use). We reviewed the following algorithms:

- We do not compare to TS algorithm because either the guarantees are provided for Bayesian regret (Ouyang et al., 2017), or for frequentist regret but limited to 1 dimensional systems (Abeille & Lazaric, 2018). We also do not compare to (Dean et al., 2018) as the regret guarantee is worse $\tilde{O}(T^{2/3})$.
- We do not compare to other OFULQ algorithms as they are intractable (Abbasi-Yadkori & Szepesvári, 2011; Faradonbeh et al., 2017)¹⁹.
- ϵ -greedy approaches that has been studied in (Faradonbeh et al., 2018b; Mania et al., 2019; Simchowitz & Foster, 2020). (Mania et al., 2019) do not provide explicit algorithm, but a proof that ϵ -greedy (properly tuned) is optimal (\sqrt{T} regret). (Faradonbeh et al., 2018b; Simchowitz & Foster, 2020) propose the same scheme, but the latter provide an algorithm that is fully explicit (in the way ϵ should be tuned) and for which they claim to have optimal dependencies in the dimension n, d . We thus focus on (Simchowitz & Foster, 2020) (CECCE).
- We do not compare to OSLO (Cohen et al., 2019), since it requires a long initial phase (\sqrt{T}). As a result, OSLO has \sqrt{T} regret which is not due to the optimistic SPD resolution, but follows from (Mania et al., 2019) (actually, OSLO is closer to an Explore-Then-Commit scheme).²⁰

As a result, we only compare DS-OFU and CECCE.

Initialization phase. Both DS-OFU and CECCE require an initialization procedure before actually addressing the exploration-exploitation dilemma, to ensure that the state is bounded. LAGLQ takes a stabilizing set Θ_0 as input, while

¹⁹Heuristic gradient descent approaches to the non-convex problem could be used, but again, we wanted to focus on algorithms that provably solve their optimization problem

²⁰Furthermore, if we implement the algorithm rigorously, the length of the initial phase implies that the policy is also updated very rarely. In particular, for the experiment conducted here, and given the time horizon, OSLO does not perform any update (the first update occurs for $T \approx 10^8$).

CECCE provides a procedure to compute such set, given that a stabilizing controller K_0 is known. To allow for fair comparison, we use the stabilizing controller approach of CECCE both to perform the initialization procedure of CECCE and to construct a stabilizing set for LAGLQ that is given as input.

More in detail, we use a initial phase of length T_0 for all algorithms where we apply a stable controller K_0 and add random Gaussian noise $\mathcal{N}(0, I)$ to the control. This early phase is not counted towards the regret since it is common to the 2 algorithms. After this phase, we are given a stabilizing set and we set $t = 0$ and $x_0 = 0$. Then, the safe set is provided as input to CECCE and DS-OFU which follow respectively a ϵ -greedy with decaying variance and optimistic strategies.

The length T_0 of the system identification phase we use in the experiments is not the one recommended by CECCE. The theoretical lower bound on T_0 is very worst case and would lead to a initialization phase of length $T_0 \approx 10^8$ in our case. On the other hand, a shorter length $T_0 \approx 2 \times 10^4$ is enough to guarantee the stabilizability of the state for the problem at hand (easy to verify by looking at the value of $\|x_t\|$ over the trajectory).

Experiments parameters. We run the DS-OFU and CECCE algorithms on trajectories of length $T = 10^6$ for the LQR problem parametrized by (115). We report in Fig. 3 the average and the 90th percentile of the regret over 100 trajectories. We observe that LAGLQ is better than CECCE when both are implemented as recommended by theory. The worse perf of CECCE is directly due to the amount of noise injected in the system, which seems to make CECCE over-explorative. This may be caused by two issues:

1. The injected noise is large because the exploration scheme does not adapt to the uncertainty (the confidence sets). As a result, it has to be conservative and set w.r.t. the most difficult direction (in terms of exploration).
2. The analysis is loose (the constants not the rate) which leads to bad design and hence bad perfs (notice that we may have the same effect for LAGLQ, coming from the fact that the confidence set is not tight).

In the attempt of addressing the second possible cause of “bad” performance, we also test a “tuned” instance of ϵ -greedy, where the variance of the injected noise scales with $\|P_*\|^4$ (instead of $\|P_*\|^{5/2}$), and set all the others constants to 1. We do not tune the rate as it is already optimal. In this case, ϵ -greedy seems to be comparable to LAGLQ, but the overall trend still seems worse than LAGLQ.

Runtime of LAGLQ. Overall the implementation of LAGLQ on this problem is very efficient. Each call to DS-OFU takes around 50 – 80m.s. on a standard computer and it takes around 35 – 40 iterations to find an optimistic ϵ -feasible optimal policy for an accuracy $\epsilon = 10^{-12}$. Furthermore, we noticed that empirically the backup procedure BACKUPPROC is never really triggered. We conjecture that the corner cases considered by the backup procedure may actually correspond to a set of systems of zero Lebesgue measure.