

Federated LQR: Learning through Sharing

Zhaolin Ren^{*}, Aoxiao Zhong^{*}, Zhengyuan Zhou[†], and Na Li^{*}

^{*}School of Engineering and Applied Sciences, Harvard University

[†]Stern School of Business, New York University

September 29, 2020

Contents

1	Introduction	2
1.1	Our contributions	3
1.2	Related work	4
1.3	Notations	4
2	Problem Setup	5
3	Properties of distributed LQR tracking problem	6
4	Algorithm design	7
4.1	Background on zeroth-order model-free LQR learning	7
4.2	Proposed algorithm: Federated LQR	8
5	Main results	10
6	Numerical results	13
7	Conclusion and future work	14
A	Optimal controller for LQR tracking problem	17
B	Proofs for properties of LQR tracking	19
B.1	Lemma 1: Existence of global minimizers	19
B.2	Lemma 2: Non-convexity	19
B.3	Proposition 1: Gradient dominance	19
B.4	Proposition 2: Local Lipschitzness and smoothness	30
C	Algorithms	44

D	Analysis of main results	45
D.1	Theorem 1	46
D.2	Theorem 2: Benchmark algorithm convergence	53
D.3	Theorem 3: Comparison of federated and benchmark algorithm	56
D.3.1	Properties of the one-point and two-point estimators	56
D.3.2	Convergence for one-point estimator	58
D.3.3	Convergence for two-point estimator	59
D.4	Discussion	60
E	More details about numerical results	61
E.1	Linear system	61
E.2	Nonlinear system	62

Abstract

In many multi-agent reinforcement learning applications such as flocking, multi-robot applications and smart manufacturing, distinct agents share similar dynamics but face different objectives. In these applications, an important question is how the similarities amongst the agents can accelerate learning in spite of the agents’ differing goals. We study a distributed LQR (Linear Quadratic Regulator) tracking problem which models this setting, where the agents, acting independently, share identical (unknown) dynamics and cost structure but need to track different targets. In this paper, we propose a communication-efficient, federated model-free zeroth-order algorithm that provably achieves a convergence speedup linear in the number of agents compared with the communication-free setup where each agent’s problem is treated independently. We support our arguments with numerical simulations of both linear and nonlinear systems.

1 Introduction

Sample complexity is an important factor in determining the broad applicability of a Reinforcement Learning (RL) algorithm. To reduce sample complexity and hence accelerate learning, various approaches have been proposed in the field of parallel RL ([1], [2]), where multiple agents with the same unknown dynamics and objective share information in their learning process. To broaden the scope of this setting, a natural extension is to consider the case when distinct agents share similar dynamics but seek different objectives. This is motivated by applications such as flocking [3], multi-robot applications [4], and smart manufacturing [5], etc. For instance, in formation control of multi-robot systems, a group of similar robots might need to form a pre-specified shape by tracking different target positions. We note that our setting shares many similarities with federated learning [6], such as (i) distributed learning via averaging of local updates, and (ii) communication constraints. Therefore, for ease of exposition, we will use the term *federated* when referring to our RL problem later.¹

We are motivated by the following questions:

¹We recognize that unlike many papers in the federated learning literature, our work does not consider privacy. Our use of the term “federated” is made primarily to differentiate our

1. Can the agents leverage the similarities in their problem structures in order to accelerate learning, despite their differing targets?
2. To what extent can they do so in a communication-efficient way?

To make concrete progress, we center our study on a federated LQR tracking problem, where different agents follow the same (unknown) dynamics but seek to track distinct targets. For each agent i , let $x_t^i \in \mathbb{R}^n$ and $u_t^i \in \mathbb{R}^k$ denote its state and action respectively at time t , with the state dynamics evolving as $x_{t+1}^i = Ax_t^i + Bu_t^i$. Agent i 's goal is to drive the state to a target location x^{i*} , which we model via the following discounted infinite-horizon cost:

$$\begin{aligned} \min_{\substack{u_t^i, \\ t \in \mathbb{N}}} J^i &:= \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \left((x_t^i - x^{i*})^\top Q (x_t^i - x^{i*}) + (u_t^i)^\top R u_t^i \right) \right] \\ \text{s.t.} \quad &x_{t+1}^i = Ax_t^i + Bu_t^i, \quad x_0^i \sim N(0, \Sigma), \end{aligned} \quad (1)$$

where $0 < \gamma < 1$ is a discount factor and Σ is a positive-definite matrix. To ease technical analysis, we assume only the initial state x_0^i is stochastic. While the above federated LQR tracking problem enjoys greater modelling flexibility, it brings new challenges in both algorithm design and convergence analysis. For instance, recent work [7] has shown that the classical LQR (single agent and $x^{i*} = 0$) enjoys the *gradient domination* property, known also as the Polyak-Lojasiewicz inequality ([8], [9]), a crucial condition to achieve global convergence for policy gradient. However, it is unclear if gradient domination still holds for our more general LQR tracking problem (even in the single agent setting). In addition, for this problem, an optimal policy for each agent i is $u_t^i = Kx_t^i + g^i$ [10], where K is identical across the agents, due to the common (A, B, Q, R) matrices, but the (g^i) 's are distinct. While the common K suggests possible merits from pooling information, it is unclear how the agents should efficiently aggregate/communicate their information (without sharing their data) in joint policy learning and what the quantitative speedup may be. Further, the distinct (g^i) 's make it difficult to ascertain if communications among the agents would be truly beneficial.

1.1 Our contributions

First, we study the LQR tracking problem (1) and establish its gradient domination property for both the single-agent and the federated setting. While previous works [7] have established the gradient domination property of LQR when $x^{i*} = 0$, that setting is simpler as the optimal policy takes the linear form $u^i = Kx^i$. To the best of our knowledge, this paper is the first to establish the gradient domination property for the more general LQR tracking problem, when an extra (and distinct) g^i is present in each agent's optimal policy. Second, we exploit the fact that the same K matrix is shared across agents (despite distinct g^i terms) and propose a model-free federated zeroth-order policy gradient algorithm where we interweave local update steps and averaging steps (when agents

setting from other works in distributed LQR learning where agents' dynamics or cost functions could be more general.

pool information on the matrix K), balancing between communication efficiency (allowing for independent updates) and benefits from shared learning (averaging over the K matrix). Under mild conditions, we show that our algorithm provides a speedup linear in the number of agents over a communication-free algorithm where each agent learns by itself. This linear speedup, even in the important special case where tracking is absent (i.e. all $x^{i*} = 0$), is previously unknown and contributes to the broad landscape of reinforcement learning that policy gradient can be scaled up gracefully (in terms of both sample complexity and computational efficiency) for the multi-agent LQR problem. Further, and perhaps even more surprisingly, such linear speedup still remains intact even when each agent has a heterogeneous component built in their objective. Finally, we provide simulation results to demonstrate the effectiveness of our proposed federated policy gradient algorithm in both the linear system setting (where our convergence theory applies) and the nonlinear system setting (that goes beyond our convergence guarantees).

1.2 Related work

First, as our setting is model-free, gradient estimators in our algorithms are based on zeroth-order, i.e. function value information, situating our work in the zeroth-order optimization literature ([11]–[16]). Second, our work reposes on the LQR literature. LQR is a classical reinforcement learning problem [10], and in recent years it has been studied with renewed vigor from both model-based ([17]–[19]) and model-free perspectives ([7], [14], [20]). Another line of LQR work focuses on online learning of LQR subject to possibly adversarial conditions ([21], [22]). Our work takes a model-free approach and is closest in spirit to ([7], [20]). Third, we note that our problem is a special instance of parallel or concurrent RL ([1], [2], [23]). Fourth, our work is related to federated learning, an approach proposed to train centralized models across agents with heterogeneous data distributions subject to communication and privacy constraints ([6], [24]). Recent works have extended federated learning to the RL setting ([25], [26]). Our proposed algorithm uses a local computation and then aggregation procedure, a hallmark of federated learning algorithms based on local SGD [27]. Fifth, the fact that we seek to learn individualized policies for different agents sharing similar problem structures relates our work to federated meta-learning ([2], [28]).

1.3 Notations

$\|\cdot\|$ refers to the Euclidean norm for vectors and Frobenius norm for matrices. In addition, for any algorithm, \mathcal{F}^t denotes its filtration up to time t and $\mathbb{E}^t := \mathbb{E}[\cdot | \mathcal{F}^t]$ denotes the conditional expectation on \mathcal{F}^t . When appropriate, for any vector v , we use $(v)_K$ to refer to its subvector corresponding to the K matrix and $(v)_g$ to refer to its subvector corresponding to the constant term g . For a positive integer m , $[m]$ refers to the set $\{1, 2, \dots, m\}$.

2 Problem Setup

We now formally introduce the federated LQR tracking problem. Consider m agents each with the same dynamics described in (1). The overall goal is to minimize the averaged cost of all the agents (here we denote $\tilde{x}_t^i := x_t^i - x^{i*}$):

$$\begin{aligned} \min_{\substack{u_t^i \\ i \in [m] \\ t \in \mathbb{N}}} J_{avg} &:= \frac{1}{m} \sum_{i=1}^m \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t ((\tilde{x}_t^i)^\top Q \tilde{x}_t^i + (u_t^i)^\top R u_t^i) \right] \\ \text{s.t.} \quad \tilde{x}_t^i &:= x_t^i - x^{i*}, x_{t+1}^i = A x_t^i + B u_t^i, x_0^i \sim N(0, \Sigma). \end{aligned} \quad (2)$$

Because the dynamics and objectives are totally decoupled among agents, the optimal controllers for Problem (2) are equivalent to optimal controllers for the single agent Problem (1). As discussed earlier, agent i 's optimal policy is $u_t^i = K^i x_t^i + g^i$ at time t [10]. Note that the target position, x^{i*} , can vary amongst the agents. Despite this, due to the common (A, B, Q, R) matrices, an optimal controller set takes the form $\{(K, g^i)\}_{i=1}^m$ where every agent share the same $K^i = K$ but with distinct g^i . We demonstrate this property in the appendix.

Due to the above property, the following two formulations both allow us to find an optimal controller set $\{(K^i, g^i)\}_{i=1}^m$ that minimizes J_{avg} for Problem (2). The first formulation, which we refer to as the *federated* formulation, uses the fact that we can use the same controller K for each of the agents.

$$\begin{aligned} \min_{K, G} \bar{J}(K, G) &= \frac{1}{m} \sum_{i=1}^m \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t ((\tilde{x}_t^i)^\top Q \tilde{x}_t^i + (u_t^i)^\top R u_t^i) \right] \\ \text{s.t.} \quad \tilde{x}_t^i &:= x_t^i - x^{i*}, x_{t+1}^i = A x_t^i + B u_t^i, u_t^i = K x_t^i + g^i, \\ x_0^i &\sim N(0, \Sigma), G = [g_1^\top \quad g_2^\top \quad \cdots \quad g_m^\top]^\top. \end{aligned} \quad (3)$$

The second formulation, which we call the *independent* formulation, treats the m problems independently, giving a distinct variable K^i for each agent $i \in [m]$.

$$\begin{aligned} \min_{K^i, g^i} J^i(K^i, g^i) &= \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t ((\tilde{x}_t^i)^\top Q (\tilde{x}_t^i) + (u_t^i)^\top R u_t^i) \right] \\ \text{s.t.} \quad \tilde{x}_t^i &:= x_t^i - x^{i*}, x_{t+1}^i = A x_t^i + B u_t^i, \\ u_t^i &= K^i x_t^i + g^i, x_0^i \sim N(0, \Sigma). \end{aligned} \quad (4)$$

The federated formulation allows the m agents to learn together based on shared information about K , while the independent formulation removes communication burdens since the agents can learn independently. In the sequel, we will describe and develop an algorithm to minimize J_{avg} that in some sense interpolates between the two approaches, enjoying the communication savings of independent updates as well as the information pooling from having a single K . During our learning process, we assume that we are unaware of the system matrices A and B , but know the cost matrices Q and R , as well as the target

x^{i*} for each i ; assuming the knowledge of stage cost/reward functions is a reasonable assumption in many RL applications [29]. We also assume that we have full state information. Given a policy (K^i, g^i) , and a randomly drawn initial state x_0^i , this allows us to measure the cost $J^i(K^i, g^i; x_0^i)$ for an agent i , which is a noisy realization of the (expected) local cost $J^i(K^i, g^i)$. When the context is clear, we may drop the index i and refer to $J(K, g)$ as the single-agent cost for LQR tracking problem. Throughout our analysis, we will also impose the following assumptions.

Assumption 1. The dynamical system governed by the matrices (A, B) is controllable [10].

Thus, there exists $K \in \mathbb{R}^{k \times n}$ such that $\rho(A + BK) < 1$, where $\rho(\cdot)$ measures the spectral radius of a matrix.

Assumption 2. We have access to a stable controller K_0 such that $\rho(A + BK) < 1$.

This is an assumption commonly made in the LQR literature ([7], [19], [20]).

Assumption 3. Cost matrices Q, R and the covariance Σ for the initial state are all positive definite.

3 Properties of distributed LQR tracking problem

We first study the properties of \bar{J} and J^i which will facilitate our algorithm design and analysis.

Lemma 1 (Existence of global minimizer). *For each J^i , $i \in [m]$, there exists $(J^i)^* \geq 0$ such that*

$$\min_{K, g} J^i(K, g) = (J^i)^*.$$

By a similar token, there exists $J_{avg}^ \geq 0$ such that*

$$\min_{\{K^i, g^i\}_{i=1}^m} \frac{1}{m} \sum_{i=1}^m J^i(K^i, g^i) = \min_{K, G} \bar{J}(K, G) = \min J_{avg} = J_{avg}^*.$$

We note in particular that the global minimum values of $\frac{1}{m} \sum_{i \in [m]} J^i$ (independent LQR tracking) and \bar{J} (federated LQR tracking) are identical to each other and to the global minimum of J_{avg} .

Lemma 2 (Non-convexity). *If $n \geq 2$, the LQR tracking problem (1) is in general non-convex.*

This is a consequence of the fact that when $n \geq 2$, there exists matrices K and K' such that $J(K, 0)$ and $J(K', 0)$ are both finite but $J((K + K')/2, 0)$ is not finite. We provide an example in the appendix. Without additional assumptions, gradient descent on a non-convex problem can only reach a stationary point, not necessarily a global minimum. However, it is known that

under gradient domination and appropriate smoothness properties, gradient descent methods, even for a non-convex problem, can find a global minimizer at a rate comparable to that of strongly convex functions [30]. Pioneering work in [7] established the gradient domination and local smoothness of the classical LQR problem (where the target $x^* = 0$), when the optimization variable only involves K . We establish that both properties hold in the more general tracking case when the policy involves K as well as g .

Proposition 1 (Gradient domination of LQR tracking problem). *Suppose K is stable for the system (A, B) , and that Σ is positive definite. Then, for any agent i , there exists $\mu^i > 0$ such that*

$$J_i(K, g^i) \leq \mu^i \|\nabla J^i(K, g^i)\|^2.$$

Moreover, for the federated cost $\bar{J}(K, G)$, letting $\mu := \max_{i \in [m]} \mu^i$, we have

$$\bar{J}(K, G) \leq \mu \|\nabla \bar{J}(K, G)\|^2.$$

While building on the techniques in [7], the addition of a constant term g changes the complexion of the problem, and requires new analysis. We defer the full technical proof of Proposition 1 to the appendix. However, gradient domination alone cannot ensure convergence to a global minimum. As noted in ([7], [20]), a function needs to also exhibit local smoothness properties.

Proposition 2 (Local smoothness of LQR tracking problem). *Consider any $i \in [m]$. Let $\mathcal{G}_C = \{(K, g) \mid J^i(K, g) \leq C\}$ be a sublevel set of J^i , where $C > 0$. Then, there exists a local radius $\rho^i > 0$, Lipschitz parameter $\lambda^i > 0$ and Lipschitz smoothness parameter $L^i > 0$ such that for any (K, g) in \mathcal{G} , whenever $\|(K', g') - (K, g)\| \leq \rho^i$,*

$$\begin{aligned} \|J^i(K', g') - J^i(K, g)\| &\leq \lambda^i \|(K', g') - (K, g)\|, \\ \|\nabla J^i(K', g') - \nabla J^i(K, g)\| &\leq L^i \|(K', g') - (K, g)\|. \end{aligned}$$

Similarly, for the federated cost \bar{J} , let $\bar{\mathcal{G}}_C = \{(K, G) \mid \bar{J}^i(K, G) \leq C\}$ be a sublevel set of \bar{J} , where $C > 0$. Then, there exists a local radius $\bar{\rho} > 0$, Lipschitz parameter $\bar{\lambda} > 0$ and Lipschitz smoothness parameter such that for any (K, G) in $\bar{\mathcal{G}}_C$, whenever $\|(K', G') - (K, G)\| \leq \bar{\rho}$,

$$\begin{aligned} \|\bar{J}(K', G') - \bar{J}(K, G)\| &\leq \bar{\lambda} \|(K', G') - (K, G)\|, \\ \|\nabla \bar{J}(K', G') - \nabla \bar{J}(K, G)\| &\leq \bar{L} \|(K', G') - (K, G)\|. \end{aligned}$$

We provide a proof of this result in the appendix.

4 Algorithm design

4.1 Background on zeroth-order model-free LQR learning

In model-free LQR learning, a common choice is the use of policy gradient [7]. To illustrate this, consider the single agent case, when the tracking target x^* is zero.² Then, an optimal policy to minimize the infinite-horizon cost is given

²We drop the index i in this subsection for ease of notation.

by $u_t = Kx_t$, i.e. a linear time-invariant policy. Policy gradient on K then involves iterating the update $K_{t+1} = K_t - \eta \hat{z}_t$, where \hat{z}_t is an estimate of the policy gradient $\nabla J(K_t)$, $\eta > 0$ is a step-size. Exactly computing the gradient $\nabla J(K_t)$ requires knowledge of the system matrices (A, B) , and recent works ([7], [20]) have proposed a zeroth-order mechanism to estimate the gradient. In this line of work, the estimate \hat{z}_t is calculated based on function value information about the cost $J(\cdot)$ near the current policy K_t . One example is the symmetric two-point estimator, where

$$\hat{z}_t = (n \cdot k) \frac{J(K + rU; x_0) - J(K - rU; x_0)}{2r} U.$$

Above, $r > 0$ can be viewed as a smoothing radius, and U_t is a random perturbation vector drawn at random from the unit sphere $\text{Unif}(\mathbb{S}^{k \cdot n - 1})$, where we note that $n \cdot k$ is the dimension of the ambient space of K . The factor $n \cdot k$ is necessary so that the estimator has the same scaling as the true gradient $\nabla J(K_t)$. The symmetric two-point estimator resembles a finite-difference approximation of the gradient, and as such is clearly motivated. We point out here a technical caveat — the notation $J(K \pm rU; x_0)$ means that the two measurements are based on the policies $K \pm rU$ for a common random initial state x_0 . In cases when we cannot obtain two neighboring function evaluations corresponding to the same noisy initial state x_0 , a one-point estimator based on a single function evaluation may be required ([14], [31])³. In our work, we adapt existing zeroth-order estimators proposed in the literature to our setup, where the policy now comprises not just the matrix K but also the constant term g .

4.2 Proposed algorithm: Federated LQR

Motivation. Our overall goal is to find a controller set $\{(K^i, g^i)\}$ such that $\frac{1}{m} \sum_{i=1}^m J^i(K^i, g^i)$ is close to the optimal value J_{avg}^* , using zeroth-order information to form policy gradient updates. When the m agents learn independently, they are unable to share information on K . Conversely, in the federated formulation (Problem 3), since agents share the same K , intuitively this enables them to learn the optimal K faster, accelerating the learning process. However, the benefits of information sharing comes at the price of increasing the amount of communication and coordination among agents. To see this, we observe that the gradient of the federated cost, $\bar{J}(K, G)$ where $G = [g_1^\top \quad g_2^\top \quad \cdots \quad g_m^\top]^\top$, takes the following form,

$$\nabla_K \bar{J}(K, G) = \frac{1}{m} \sum_{i=1}^m \nabla_K J^i(K, g^i), \quad (5)$$

$$\nabla_{g^i} \bar{J}(K, G) = \frac{1}{m} \nabla_{g^i} J^i(K, g^i) \quad \forall i \in [m]. \quad (6)$$

when $\rho(A + BK) < 1$. While computing the gradient of \bar{J} requires only local computations of $\nabla J^i(K, g^i)$ and then aggregating these computations to form

³While we focus on the two-point estimator for clarity in the main paper, we do include analysis of the one-point estimator in our appendix.

∇_K , this imposes a high synchronous requirement on the agents, since communication is required for each (zeroth-order) gradient update. The strength of the independent formulation where each K^i evolves freely is the absence of such communication burdens.

Our federated approach. To balance between communication efficiency and benefits from sharing the same K , we propose a federated algorithm where agents interweave steps of gradient updates using local information and steps where they pool information on the matrix K . Our federated LQR learning algorithm is provided in Algorithm 2. Below we highlight its key features.

Algorithm 1: Federated LQR Learning algorithm

- 1 Given: iteration number $T \geq 1$, communication interval $H \in \mathbb{N}$, initial stable K_0 , initial g_0^i for each $i \in [m]$, step size $\eta > 0$, and smoothing radius $r > 0$
- 2 **for** epoch $e \in \{0, \dots, \frac{T}{H} - 1\}$ **do**
- 3 Set $K_{eH}^i \leftarrow K_{eH}$ for each $i \in [m]$
- 4 **for** iteration $t = eH, eH + 1, \dots, (e + 1)H - 1$ **do**
- 5 **for** (simultaneously) each agent $i \in [m]$ **do**
- 6 Sample $(x_0)_t^i \sim \mathcal{D}$, $u_t^i \sim \text{Unif}(\mathbb{S}^{n \cdot k + k - 1})$
- 7 Set the zeroth-order estimator z_t^i as follows (note $d := n \cdot k + k$) :

$$(z_t^i) \leftarrow \frac{d(J^i((K_t^i, g_t^i) + ru_t^i(x_0)_t^i) - J^i((K_t^i, g_t^i) - ru_t^i(x_0)_t^i))}{2r} u_t^i$$

Update

$$K_{t+1}^i \leftarrow K_t^i - \eta(z_t^i)_K,$$

$$g_{t+1}^i \leftarrow g_t^i - \frac{\eta}{m}(z_t^i)_g, \quad t \leftarrow t + 1$$

- 8 **end**
 - 9 **end**
 - 10 Set $K_{(e+1)H} \leftarrow \frac{1}{m} \sum_{i=1}^m K_{(e+1)H}^i$ (model averaging of K at end of epoch e)
 - 11 **end**
 - 12 **return** $K_T, g_T^i \forall i \in [m]$
-

- *Initial conditions (Line 1).* We assume knowledge of a stable $K_0 \in \mathbb{R}^{k \times n}$. The choice of initial constant term $g_0^i \in \mathbb{R}^k$ for each i is flexible, and can be picked to be the zero vector without prior information. We assume a constant step-size η and smoothing radius r .
- *Independent updates during each epoch (Lines 6 through 8).* During the e -th epoch, starting from time $t = eH$ and a common initial matrix K_{eH} , each agent $i \in [m]$ simultaneously runs independent zeroth-order gradient descent

dynamics for H iterations, as follows:

$$K_{t+1}^i = K_t^i - \eta(z_t^i)_K, \quad g_{t+1}^i = g_t^i - \frac{\eta}{m}(z_t^i)_g,$$

where z_t^i is the zeroth-order estimator of $\nabla J^i(K_t^i, g_t^i)$ computed by agent i .

- *Zeroth-order gradient estimator (Line 7)*. We adopt a two-point zeroth-order gradient estimator, as seen in Line 7 of the algorithm.
- *Averaging at the end of each epoch over K (Line 10)*. At the end of each epoch, we perform model averaging over the K through

$$K_{(e+1)H} = \frac{1}{m} \sum_{i \in [m]} K_{(e+1)H}^i.$$

Benchmark algorithm. We evaluate the performance of our algorithm against one which is communication-free, where the m agents evolve their policy gradient dynamics independently. We note that the zeroth-order estimator $z_t^i(\cdot)$ is the same across both algorithms. We place description of the benchmark algorithm in the appendix due to space limitations.

5 Main results

Preliminaries. We first introduce several quantities which will appear in the main results. The reader might wish to skim this part during a first reading, returning for a closer look later.

Consider the terms

$$Z_\infty^i := \max\{\|z_t^i(\cdot)\|\}, \quad Z_\infty := \max_{i \in [m]} Z_\infty^i$$

$$Z_{2,K} := \max_t \frac{1}{m} \sum_{i=1}^m \mathbb{E}^t \|(z_t^i)_K\|^2, \quad Z_{2,G} := \max_t \frac{1}{m} \sum_{i=1}^m \mathbb{E}^t \|(z_t^i)_{g^i}\|^2,$$

which collectively form bounds on the maximum size and variance of the zeroth-order gradient update. We next define the stability region \mathcal{G}_0 and \mathcal{G}'_0 which we show in the appendix that the iterates of Algorithm 2 and the benchmark algorithm respectively stay within. Let

$$\mathcal{G}_0 := \{(K, G) : \bar{J}(K, G) \leq 10\bar{J}(K_0, G_0)\},$$

$$\mathcal{G}'_0 := \left\{ \{K^i, g^i\}_{i \in [m]} : \frac{1}{m} \sum_{i=1}^m J^i(K^i, g^i) \leq 10 \frac{1}{m} \sum_{i=1}^m J^i(K_0^i, g_0^i) \right\}.$$

Next, by Proposition 2, note that there exists a local radius $\rho > 0$, local Lipschitz parameter $\lambda > 0$ and local smoothness parameter $L > 0$ such that the following holds.

1. If $(K, G) \in \mathcal{G}_0$, for any $i \in [m]$, and $\|(K', (g^i)') - (K, g^i)\| \leq \rho$, then

$$\begin{aligned}\|\nabla J^i(K', (g^i)') - \nabla J^i(K, g^i)\| &\leq L\|(K', (g^i)') - (K, g^i)\|, \\ \|J^i(K', (g^i)') - J^i(K, g^i)\| &\leq \lambda\|(K', (g^i)') - (K, g^i)\|.\end{aligned}$$

2. If $(\{K_i, g^i\}_{i=1}^m) \in \mathcal{G}'_0$, for any $i \in [m]$, and $\|((K^i)', (g^i)') - (K^i, g^i)\| \leq \rho$,

$$\|\nabla J^i((K^i)', (g^i)') - \nabla J^i(K^i, g^i)\| \leq L\|((K^i)', (g^i)') - (K^i, g^i)\|.$$

3. If $(K, G) \in \mathcal{G}_0$, and $\|(K', G') - (K, G)\| \leq \rho$, then

$$\|\nabla \bar{J}(K', G') - \nabla \bar{J}(K, G)\| \leq L\|(K', G') - (K, G)\|.$$

These preliminary definitions pave the way for our main results. We begin with a statement of the convergence of the federated LQR tracking algorithm, Algorithm 2.

Theorem 1 (Convergence of federated LQR tracking). *Suppose the step-size η , smoothing radius r and communication interval H are chosen to satisfy*

$$\begin{aligned}\eta &\leq \min \left\{ \frac{m\epsilon}{240\mu L(Z_{2,G} + Z_{2,K} + 6m\eta LH^2 Z_\infty^2)}, \frac{L}{8}, \frac{\rho}{2HZ_\infty} \right\}, \\ r &\leq \min \left\{ \sqrt{\frac{\epsilon\mu}{720L}}, \rho \right\}.\end{aligned}$$

Then, if the error tolerance ϵ satisfies $\epsilon \log(120\Delta_0/\epsilon) \leq 5\Delta_0$, the iterate (K_t, G_t) produced by Algorithm 2 satisfies

$$\frac{1}{m} \sum_{i=1}^m J^i(K_T, g_T^i) - J_{avg}^* \leq \epsilon,$$

when $T = \frac{4\mu}{\eta} \log(120\Delta_0/\epsilon)$ steps, with probability at least $3/4$, where $\Delta_0 = \frac{1}{m} \sum_{i=1}^m J^i(K_0, g_0^i) - J_{avg}^$.*

We make a few brief remarks about Theorem 1, deferring extended discussion to the appendix. Inherently, maintaining stability of the matrix K , i.e. $\rho(A + BK) < 1$, is a critical issue for LQR learning, since when K is unstable, the infinite-horizon LQR cost can diverge [7]. This stability requirement is in tension with the constant accumulation of noise in the learning procedure due to the stochastic zeroth-order updates. For this reason, the convergence result holds with a constant probability, namely 0.75 (cf. Theorem 1). Nonetheless, we wish to clarify that the number 0.75 itself can be increased by carefully tightening our analysis in certain places. Second, to further increase the convergence probability of the overall learning process to $1 - \delta$ for *any* $\delta > 0$, we can either (i) reduce the learning rate by using a more conservative stepsize η , or (ii) evaluate a short list of solutions generated by several independent runs of the

algorithm (cf. [32]). We believe both these approaches will come at the cost of an additional dependence on $\tilde{O}\left(\frac{1}{\delta}\right)$ in the sample complexity – we leave rigorous analysis of this to future work. Finally, we note that our analysis relies heavily on the gradient domination and local smoothness properties of the federated cost \bar{J} .

We have a corresponding convergence result for the benchmark distributed independent algorithm.

Theorem 2 (Convergence for independent distributed LQR tracking). *Suppose*

$$\eta \leq \min \left\{ \frac{\epsilon}{240\mu L(Z_{2,G} + Z_{2,K})}, \frac{\rho}{2Z_\infty}, \frac{L}{8} \right\},$$

$$r \leq \min \left\{ \sqrt{\frac{\epsilon\mu}{720L}}, \rho \right\}.$$

Then, if the error tolerance ϵ satisfies $\epsilon \log(120\Delta_0/\epsilon) \leq 5\Delta_0$, the iterates $\{(K_t, g_t^i)_{i=1}^m\}$ produced by the benchmark algorithm satisfy

$$\frac{1}{m} \sum_{i=1}^m J^i(K_T^i, g_T^i) - J_{avg}^* \leq \epsilon,$$

when $T = \frac{4\mu}{\eta} \log(120\Delta_0/\epsilon)$ steps, with probability at least $3/4$, where $\Delta_0 = \frac{1}{m} \sum_{i=1}^m J^i(K_0^i, g_0^i) - J_{avg}^$.*

The analysis technique is similar to that for Theorem 1, and we defer further discussion to the appendix. Building on Theorems 1 and 2, we next compare the convergence rate of the proposed and benchmark algorithms.

Theorem 3 (Speedup from the federated approach). *Let $H = \tilde{O}\left(\frac{\sqrt{T}}{\sqrt{d\lambda m}}\right)$ be the communication interval. Then, for the federated algorithm, to reach an ϵ -optimality gap with probability at least $3/4$, subject to the assumptions in Theorem 1, we will need $T = \tilde{O}\left(\frac{d}{m\epsilon}\right)$ steps.*

Meanwhile, for the benchmark, communication-free algorithm, subject to the assumptions in Theorem 2, reaching an ϵ -optimality gap with probability at least $3/4$ requires $T = \tilde{O}\left(\frac{d}{\epsilon}\right)$ steps.

Above, \tilde{O} hides logarithmic factors as well as smoothness and strong convexity parameters, and $d = n \cdot k + k$. Thus, under the assumptions set out in Theorem 1 and Theorem 2, the federated, communication-based, approach yields a speedup linear in the number of agents.

Under appropriate settings of H (scaling as $\tilde{O}\left(\frac{\sqrt{T}}{\sqrt{d\lambda m}}\right)$), Theorem 3 says that convergence for the federated approach is $\tilde{O}(m)$ times faster. This follows primarily from the periodic averaging of the K matrices, which reduces the variance in the optimization problem up to a factor of m in a similar vein to the results in the federated averaging/local SGD literature ([24], [27]). However, while ideally we can allow the step-size η to increase as m increases, thus accelerating convergence, in practice, due to stability concerns, there exists some

maximum step-size which permits convergence. This echoes the observation about mini-batching in [20]. A more detailed discussion is provided in the appendix.

6 Numerical results

Linear system. We compare our proposed algorithm with the benchmark algorithm for a federated LQR tracking problem. We use a LQR problem with A,B, Q, R matrices each in $\mathbb{R}^{3 \times 3}$, which can be found in the appendix. Each initial state x_0^i is sampled uniformly at random from the canonical basis vectors and the discount factor γ is set to 0.9. The tracking targets x^{i*} for each agent are randomly sampled from a zero mean Gaussian with covariance matrix $\frac{1}{10} \times I$. A two-point estimator is used.

In Figure 1a, we first show the learning process of our federated algorithm when $m = 8$ (errors denote optimality gap), comparing to the learning process of single agent. We see that the federated algorithm indeed converges, supporting the convergence result in Theorem 1, with a convergence speed faster than that for a single agent. Next, in Figure 1b, we evaluate the performance of the two algorithms by comparing the largest stepsize which permits convergence to a fixed error threshold of $\epsilon = 0.05$ with probability at least 0.7. This is an important metric since Theorems 1 and 2 show that the number of iterations required to learn a policy whose cost is ϵ -close to the optimal cost scales inversely with the stepsize, i.e. larger stepsize means fewer iterations needed for convergence. For the independent algorithm, we see that maximum stepsize permitting convergence does not increase with the number of agents, which makes sense since the different agents learn independently. Conversely, when $H = 1$, we see that the federated approach does yield an increase in maximum stable stepsize that is linear in the number of agents m as m increases from 1 to 16, which supports the conclusion in Theorem 3, but that this improvement plateaus as the number of agents increases further, confirming our earlier discussion about stability considerations precluding arbitrarily large stepsize. As H increases, i.e. less frequent communication, the increase in the maximum step-size with the number of agents remains up to 16 agents, beyond which the algorithm seems to perform the same as the independent approach. This reaffirms our observation in Theorem 3 that the maximum communication interval permitting convergence at a rate similar to when $H = 1$ scales inversely with \sqrt{m} .

Non-linear system. To demonstrate broad applicability of our algorithm, we also perform simulations of a nonlinear cartpole system [33]. In this system, a pole, which starts upright with an initial degree $\theta_0 \in [-\pi/2, \pi/2]$ and initial angular velocity $\theta'_0 \in \mathbb{R}$, is attached by an unactuated joint to a cart located at position $x \in \mathbb{R}$ along a frictionless track with velocity $v \in \mathbb{R}$. While the underlying dynamics are nonlinear [33], due to the model-free nature of our algorithm, we are nonetheless able to learn a policy (K^i, g^i) for different agents i sharing similar dynamics that all seek to balance the pole ($\theta^{i*}, \theta'^{i*} = 0$) whilst tracking




Figure 1: (a): Errors at each iteration of proposed algorithm and single agent baseline on a linear system.(mean and standard deviation for 20 runs) (b): Maximum step-size that allows for convergence to an error tolerance of $\epsilon = 0.05$. (c): Costs at each iteration of proposed algorithm on a non-linear cartpole system.

a distinct target location x^{i*} . For our implementation, we used the Cartpole package in OpenAI Gym [34]. In Figure 1c, we show the cost trajectory of the federated algorithm for this nonlinear system, and see convergent behavior, suggesting that our algorithm might be applicable even for nonlinear systems. More details are provided in the appendix.

7 Conclusion and future work

In this paper, we proved that gradient domination holds for the general LQR tracking problem. This paved the way for the development of a federated algorithm that accelerates the LQR learning process for agents sharing common dynamics but having distinct tracking targets. We highlight a few directions for future work: (1) analyzing heterogeneous dynamics, (2) studying more general cost functions, (3) considering adversarial noise.

References

- [1] M. Dimakopoulou and B. V. Roy, “Coordinated exploration in concurrent reinforcement learning,” in *Proceedings of the 35th International Conference on Machine Learning*, ser. PMLR, vol. 80, 2018, pp. 1270–1278.
- [2] T. Chen, K. Zhang, G. B. Giannakis, and T. Başar, “Communication-efficient distributed reinforcement learning,” *arXiv preprint arXiv:1812.03239*, 2018.
- [3] R. Olfati-Saber, “Flocking for multi-agent dynamic systems: Algorithms and theory,” *IEEE Transactions on automatic control*, vol. 51, no. 3, pp. 401–420, 2006.
- [4] K.-K. Oh, M.-C. Park, and H.-S. Ahn, “A survey of multi-agent formation control,” *Automatica*, vol. 53, pp. 424–440, 2015.

- [5] H. S. Kang, J. Y. Lee, S. Choi, H. Kim, J. H. Park, J. Y. Son, B. H. Kim, and S. Do Noh, “Smart manufacturing: Past research, present findings, and future directions,” *International journal of precision engineering and manufacturing-green technology*, vol. 3, no. 1, pp. 111–128, 2016.
- [6] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, “Federated learning: Strategies for improving communication efficiency,” *CoRR*, vol. abs/1610.05492, 2016. arXiv: 1610.05492.
- [7] M. Fazel, R. Ge, S. M. Kakade, and M. Mesbahi, “Global convergence of policy gradient methods for the linear quadratic regulator,” in *Proceedings of the 35th International Conference on Machine Learning*, ser. PMLR, vol. 80, 2018, pp. 1466–1475.
- [8] B. T. Polyak, “Gradient methods for the minimisation of functionals,” *USSR Computational Mathematics and Mathematical Physics*, vol. 3, no. 4, pp. 864–878, 1963.
- [9] S. Lojasiewicz, “A topological property of real analytic subsets,” *Coll. du CNRS, Les equations aux derivees partielles*, pp. 87–89, 1963.
- [10] B. D. Anderson and J. B. Moore, *Optimal control: linear quadratic methods*. Courier Corporation, 2007.
- [11] Y. Nesterov and V. Spokoiny, “Random gradient-free minimization of convex functions,” *Foundations of Computational Mathematics*, vol. 17, no. 2, pp. 527–566, 2017.
- [12] O. Shamir, “An optimal algorithm for bandit and zero-order convex optimization with two-point feedback,” *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 1703–1713, 2017.
- [13] D. Hajinezhad, M. Hong, and A. Garcia, “ZONE: Zeroth-order nonconvex multiagent optimization over networks,” *IEEE Transactions on Automatic Control*, vol. 64, no. 10, pp. 3995–4010, 2019.
- [14] Y. Li, Y. Tang, R. Zhang, and N. Li, “Distributed reinforcement learning for decentralized linear quadratic control: A derivative-free policy optimization approach,” ser. PMLR, vol. 120, Jun. 2020, pp. 814–814.
- [15] A. K. Sahu, D. Jakovetic, D. Bajovic, and S. Kar, “Distributed zeroth order optimization over random networks: A Kiefer-Wolfowitz stochastic approximation approach,” in *Proceedings of the 57th IEEE Conference on Decision and Control*, 2018, pp. 4951–4958.
- [16] Y. Tang and N. Li, “Distributed zero-order algorithms for nonconvex multi-agent optimization,” in *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, IEEE, 2019, pp. 781–786.
- [17] Y. Abbasi-Yadkori, N. Lazic, and C. Szepesvári, “Model-free linear quadratic control via reduction to expert prediction,” in *The 22nd International Conference on Artificial Intelligence and Statistics*, ser. PMLR, vol. 89, 2019, pp. 3108–3117.

- [18] S. Oymak and N. Ozay, “Non-asymptotic identification of lti systems from a single trajectory,” in *2019 American Control Conference (ACC)*, IEEE, 2019, pp. 5655–5661.
- [19] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu, “On the sample complexity of the linear quadratic regulator,” *Foundations of Computational Mathematics*, pp. 1–47, 2019.
- [20] D. Malik, A. Pananjady, K. Bhatia, K. Khamaru, P. Bartlett, and M. Wainwright, “Derivative-free methods for policy optimization: Guarantees for linear quadratic systems,” in *PMLR*, vol. 89, 2019, pp. 2916–2925.
- [21] N. Agarwal, E. Hazan, and K. Singh, “Logarithmic regret for online control,” in *Advances in Neural Information Processing Systems*, 2019, pp. 10 175–10 184.
- [22] M. Simchowitz, K. Singh, and E. Hazan, “Improper learning for non-stochastic control,” in *Conference on Learning Theory*, ser. PMLR, vol. 125, 2020, pp. 3320–3436.
- [23] K. Zhang, Z. Yang, H. Liu, T. Zhang, and T. Basar, “Fully decentralized multi-agent reinforcement learning with networked agents,” ser. PMLR, vol. 80, Oct. 2018, pp. 5872–5881.
- [24] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, “On the convergence of fedavg on non-iid data,” in *8th International Conference on Learning Representations*, 2020.
- [25] H. H. Zhuo, W. Feng, Q. Xu, Q. Yang, and Y. Lin, “Federated reinforcement learning,” *arXiv preprint arXiv:1901.08277*, 2019.
- [26] B. Liu, L. Wang, and M. Liu, “Lifelong federated reinforcement learning: A learning architecture for navigation in cloud robotic systems,” *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4555–4562, 2019.
- [27] S. U. Stich, “Local SGD converges fast and communicates little,” in *7th International Conference on Learning Representations*, 2019.
- [28] A. Fallah, A. Mokhtari, and A. Ozdaglar, “Personalized federated learning: A meta-learning approach,” *arXiv preprint arXiv:2002.07948*, 2020.
- [29] B. Recht, “A tour of reinforcement learning: The view from continuous control,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 2, pp. 253–279, 2019.
- [30] H. Karimi, J. Nutini, and M. Schmidt, “Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition,” in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, Springer, 2016, pp. 795–811.
- [31] A. D. Flaxman, A. T. Kalai, and H. B. McMahan, “Online convex optimization in the bandit setting: Gradient descent without a gradient,” in *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, 2005, pp. 385–394.

- [32] S. Ghadimi and G. Lan, “Stochastic first-and zeroth-order methods for nonconvex stochastic programming,” *SIAM Journal on Optimization*, vol. 23, no. 4, pp. 2341–2368, 2013.
- [33] S. Geva and J. Sitte, “A cartpole experiment benchmark for trainable controllers,” *IEEE Control Systems Magazine*, vol. 13, no. 5, pp. 40–51, 1993.
- [34] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, *Openai gym*, 2016. arXiv: 1606.01540 [cs.LG].
- [35] K. J. Åström and R. M. Murray, *Feedback systems: an introduction for scientists and engineers*. Princeton university press, 2010.

Appendix A Optimal controller for LQR tracking problem

We derive in this section the optimal controller for the cost J^i in Equation 1. We will show that the optimal controller takes the form $u_t^i = Kx_t^i + g^i$, where K is independent of x^{i*} . For convenience, we will drop the i index in the superscript in this section.

Proposition 3 (Optimal controller for J). *Suppose the system given by $(\sqrt{\gamma}A, \sqrt{\gamma}B)$ is controllable. The cost-to-go value function, $V(x_0)$, for the discounted infinite-horizon problem 1, takes the form*

$$V(x_0) = x_0^\top P x_0 + 2x_0^\top q + r,$$

where

$$\begin{aligned} P &= Q - (\gamma B^\top P A)^\top (R + \gamma B^\top P B)^{-1} (\gamma B^\top P A) + \gamma A^\top P A, \\ q &= -Qx^* - (\gamma B^\top P A)^\top (R + \gamma B^\top P B)^{-1} (\gamma B^\top q) + \gamma A^\top q, \\ r &= (x^*)^\top Q x^* - (\gamma B^\top q)^\top (R + \gamma B^\top P B)^{-1} (\gamma B^\top q) + \gamma r. \end{aligned}$$

Moreover, the optimal controller is time-invariant and of the form

$$u = Kx + g, \text{ where}$$

$$K = -(R + \gamma B^\top P B)^{-1} (\gamma B^\top P A), \quad g = -(R + \gamma B^\top P B)^{-1} \gamma B^\top q.$$

In particular, note that K is independent of x^* , but g is not. Then, the optimal expected cost (recall expectation is taken over the initial condition) takes the following form

$$\begin{aligned} C(u^*) &= \mathbb{E}[x_0^\top P x_0 + 2x_0^\top q + r] \\ &= \text{tr}(P\Sigma) + 2\mathbb{E}[x_0]^\top q + r, \end{aligned}$$

where $\Sigma = \mathbb{E}[x_0 x_0^\top]$.

Proof. We note that the cost-to-go value function starting from x_0 takes the form

$$V(x_0) = x_0^\top P x_0 + 2x_0^\top q + r,$$

which is motivated by appealing to the value function in the finite horizon [10] and taking the limit as t goes to infinity.

Then,

$$V(x) = \min_u (x - x^*)^\top Q (x - x^*) + u^\top R u + \gamma V(Ax + Bu).$$

Taking the gradient with respect to u , we find that

$$\begin{aligned} \nabla_u &= 2Ru + \nabla_u (\gamma ((Ax + Bu)^\top P (Ax + Bu) + 2(Ax + Bu)^\top q + r)) \\ &= 2Ru + 2\gamma B^\top P (Ax + Bu) + 2\gamma B^\top q. \end{aligned}$$

Setting ∇_u to be 0, we find that $2(R + \gamma B^\top P B)u + 2\gamma B^\top P Ax + 2\gamma B^\top q = 0$, which implies that

$$u^* = -(R + \gamma B^\top P B)^{-1} (\gamma B^\top P Ax + \gamma B^\top q).$$

Then, plugging this back into $V(x)$,

$$\begin{aligned} V(x) &= x^\top Q x + 2x^\top (-Qx^*) + (x^*)^\top Q x^* \\ &\quad + ((R + \gamma B^\top P B)^{-1} (\gamma B^\top P Ax + \gamma B^\top q))^\top R ((R + \gamma B^\top P B)^{-1} (\gamma B^\top P Ax + \gamma B^\top q)) \\ &\quad + \gamma [(Ax - B ((R + \gamma B^\top P B)^{-1} (\gamma B^\top P Ax + \gamma B^\top q)))^\top P \\ &\quad \quad (Ax - B ((R + \gamma B^\top P B)^{-1} (\gamma B^\top P Ax + \gamma B^\top q)))] \\ &\quad + \gamma (2q^\top (Ax - B ((R + \gamma B^\top P B)^{-1} (\gamma B^\top P Ax + \gamma B^\top q)))) + \gamma r \\ &= x^\top Q x + 2x^\top (-Qx^*) + (x^*)^\top Q x^* \\ &\quad + (\gamma B^\top P Ax + \gamma B^\top q)^\top (R + \gamma B^\top P B)^{-1} (\gamma B^\top P Ax + \gamma B^\top q) \\ &\quad - \gamma [2(Ax)^\top P (B ((R + \gamma B^\top P B)^{-1} (\gamma B^\top P Ax + \gamma B^\top q)))] \\ &\quad + \gamma x^\top A^\top P Ax \\ &\quad + 2\gamma [q^\top Ax - q^\top (B ((R + \gamma B^\top P B)^{-1} (\gamma B^\top P Ax + \gamma B^\top q)))] + \gamma r \\ &= x^\top Q x + 2x^\top (-Qx^*) + (x^*)^\top Q x^* \\ &\quad - (\gamma B^\top P Ax + \gamma B^\top q)^\top (R + \gamma B^\top P B)^{-1} (\gamma B^\top P Ax + \gamma B^\top q) \\ &\quad + \gamma x^\top A^\top P Ax + 2\gamma q^\top Ax + \gamma r \end{aligned}$$

Collecting terms, we find that

$$\begin{aligned} P &= Q - (\gamma B^\top P A)^\top (R + \gamma B^\top P B)^{-1} (\gamma B^\top P A) + \gamma A^\top P A, \\ q &= -Qx^* - (\gamma B^\top P A)^\top (R + \gamma B^\top P B)^{-1} (\gamma B^\top q) + \gamma A^\top q, \\ r &= (x^*)^\top Q x^* - (\gamma B^\top q)^\top (R + \gamma B^\top P B)^{-1} (\gamma B^\top q) + \gamma r. \end{aligned}$$

□

Appendix B Proofs for properties of LQR tracking

B.1 Lemma 1: Existence of global minimizers

Proof of Lemma 1. That a minimizer exists for each $J^i(K, g)$, $i \in [m]$ is a consequence of Proposition 3. That the optimal K^i for each J^i can be the same for each $i \in [m]$ is consequence of the observation in Proposition 3 that the optimal controller for J^i does not depend on x^{i*} . \square

B.2 Lemma 2: Non-convexity

We provide here a simple example, $n = 2$ (here n is state dimension), showing that $J(K)$, the centralized LQR cost, is non-convex. By extension $J(K, g)$ cannot be convex in general as well.

Proposition 4. *Consider the centralized LQR cost*

$$J(K) = \sum_{t=0}^{\infty} x_t^\top Q x_t + u_t^\top R u_t,$$

where $u_t = Kx_t$, and $x_{t+1} = Ax_t + Bu_t$, where $x_t \in \mathbb{R}^n$ and $u_t \in \mathbb{R}^k$. Then $J(K)$ is in general non-convex for $n \geq 2$.

Proof. We provide a counterexample where the set of stable controllers K is not convex. Pick

$$A = \begin{bmatrix} -0.5 & 0 \\ 0 & -0.5 \end{bmatrix}, \quad B = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad K_1 = \begin{bmatrix} 0 & 4040 \\ 0 & 0 \end{bmatrix}, \quad K_2 = \begin{bmatrix} 0 & 0 \\ 4040 & 0 \end{bmatrix}.$$

Then, (A, B) is a controllable system. In addition, observe that

$$\rho(A+BK_1) = \rho\left(\begin{bmatrix} -0.5 & 4040 \\ 0 & -0.5 \end{bmatrix}\right) < 1, \quad \rho(A+BK_2) = \rho\left(\begin{bmatrix} -0.5 & 0 \\ 4040 & -0.5 \end{bmatrix}\right) < 1.$$

However,

$$\rho\left(A + B\left(\frac{K_1 + K_2}{2}\right)\right) = \rho\left(\begin{bmatrix} -0.5 & 2020 \\ 2020 & -0.5 \end{bmatrix}\right) > 1,$$

which can be shown via direct computation. This proves that LQR is in general non-convex for $n \geq 2$. \square

B.3 Proposition 1: Gradient dominance

We first prove gradient dominance for the single-agent LQR tracking, before showing that the federated formulation with the same K variable for every agent is also gradient dominant.

Let $0 < \gamma < 1$ be a discount factor, and Σ be a covariance matrix. Consider the following problem:

$$\begin{aligned} \min_{K,g} \quad & J(K, g) = \mathbb{E} [\sum_t^\infty \gamma^t c(x_t, u_t)], \\ \text{such that} \quad & c(x_t, u_t) = (x_t - x^*)^\top Q(x_t - x^*) + u_t^\top R u_t \\ & x_{t+1} = A x_t + B u_t \\ & u_t = K x_t + g, \\ & x_0 \sim N(0, \Sigma). \end{aligned} \tag{7}$$

For a policy (K, g) , define the following cost-go-function $V_{K,g} : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$, where for any $z \in \mathbb{R}^n$,

$$V_{K,g}(z) = \sum_{t=0}^{\infty} \gamma^t c(x_t, K x_t + g), \quad \text{where } x_0 = z.$$

Lemma 3. *Suppose $\sqrt{\gamma} \rho((A + BK)) < 1$. Then,*

$$V_{K,g}(z) = z^\top P_K z + 2z^\top q_{K,g} + r_{K,g},$$

where

$$\begin{aligned} P_K &= Q + K^\top R K + \gamma(A + BK)^\top P_K (A + BK), \\ q_{K,g} &= (I - \gamma(A + BK)^\top)^{-1} (-Q x^* + K^\top R g + \gamma(A + BK)^\top P_K B g), \\ r_{K,g} &= \frac{1}{1 - \gamma} ((x^*)^\top Q x^* + g^\top R g + \gamma(g^\top B^\top P_K B g + 2g^\top B^\top q_{K,g})). \end{aligned}$$

Proof. We note that direct calculations show that $V_{K,g}(z) = z^\top P_K z + 2z^\top q_{K,g} + r_{K,g}$ for some $P_K, q_{K,g}$, and $r_{K,g}$. We will seek to find a recursive formula for $P_K, q_{K,g}$ and $r_{K,g}$. Since $V_{K,g}(z) = z^\top P_K z + 2z^\top q_{K,g} + r_{K,g}$,

$$\begin{aligned} V_{K,g}(z) &= z^\top P_K z + 2z^\top q_{K,g} + r_{K,g} \\ &= c(z, Kz + g) + \gamma V_{K,g}(Az + B(Kz + g)) \\ &= (z - x^*)^\top Q(z - x^*) + (Kz + g)^\top R(Kz + g) \\ &\quad + \gamma(A + BK)z + Bg)^\top P_K((A + BK)z + Bg) + 2((A + BK)z + Bg)^\top q_{K,g} + r_{K,g} \\ &= z^\top (Q + K^\top R K + (A + BK)^\top P_K (A + BK))z \\ &\quad + 2z^\top (-Q x^* + K^\top R g + (A + BK)^\top P_K B g + (A + BK)^\top q_{K,g}) \\ &\quad + (x^*)^\top Q x^* + g^\top R g + \gamma g^\top B^\top P_K B g + 2g^\top B^\top q_{K,g} + r_{K,g}. \end{aligned}$$

Matching coefficients, we find that

$$\begin{aligned} P_K &= Q + K^\top R K + \gamma(A + BK)^\top P_K (A + BK), \\ q_{K,g} &= -Q x^* + K^\top R g + \gamma(A + BK)^\top P_K B g + \gamma(A + BK)^\top q_{K,g}, \\ r_{K,g} &= (x^*)^\top Q x^* + g^\top R g + \gamma[g^\top B^\top P_K B g + 2g^\top B^\top q_{K,g} + r_{K,g}]. \end{aligned}$$

Some algebraic simplifications for the equations involving $q_{K,g}$ and $r_{K,g}$ then yield the result. \square

We next compute the gradient of J .

Lemma 4. *Define*

$$C_K = RK + \gamma B^\top P_K(BK + A), \quad d_{K,g} = Rg + \gamma B^\top P_K Bg + \gamma B^\top q_{K,g}.$$

The gradients of J with respect to K and g are the following:

$$\begin{aligned} \nabla_K J &= 2C_K \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t x_t^\top \right] + 2d_{K,g} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t^\top \right], \\ \nabla_g J &= 2C_K \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t \right] + 2d_{K,g} \sum_{t=0}^{\infty} \gamma^t. \end{aligned}$$

Proof. Observe that for any $z \in \mathbb{R}^n$,

$$\begin{aligned} V_{K,g}(z) &= c(z, Kz + g) + \gamma V((A + BK)z + Bg) \\ &= (z - x^*)^\top Q(z - x^*) + (Kz + g)^\top R(Kz + g) \\ &\quad + \gamma((A + BK)z + Bg)^\top P_K((A + BK)z + Bg) + 2((A + BK)z + Bg)^\top q_{K,g} + r_{K,g}. \end{aligned}$$

Taking the gradient first with respect to K , we find that

$$\begin{aligned} \nabla_K V_{K,g}(z) &= 2R(Kz + g)z^\top + 2\gamma(B^\top P_K((A + BK)z + Bg)z^\top + B^\top q_{K,g}z^\top) \\ &\quad + \gamma \nabla_K V_{K,g}(x_1) \Big|_{x_1=(A+BK)z+Bg}. \end{aligned}$$

Using recursion, and taking expectations, we find that

$$\nabla_K J = \mathbb{E} \left[\sum_{t=0}^{\infty} 2\gamma^t \left((RK + \gamma B^\top P_K(A + BK))x_t x_t^\top + (Rg + \gamma B^\top P_K Bg + B^\top q_{K,g})x_t^\top \right) \right]$$

We next compute the gradient with respect to g . Observe that

$$\nabla_g V_{K,g}(z) = 2R(Kz + g) + 2\gamma B^\top P_K((A + BK)z + Bg) + 2\gamma B^\top q_{K,g} + \gamma \nabla_g V(x_1) \Big|_{x_1=(A+BK)z+Bg}.$$

Using recursion, and taking expectations, we find that

$$\nabla_g J = 2\mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \left((RK + \gamma B^\top P(A + BK))x_t + Rg + \gamma B^\top P_K Bg + \gamma B^\top q_{K,g} \right) \right].$$

□

We define the advantage function $A_{K,g}(x, u)$ as

$$A_{K,g}(x, u) = c(x, u) + \gamma V_{K,g}(Ax + Bu) - V_{K,g}(x).$$

For any two policies (K, g) and (K', g') , the next lemma provides an expression for the difference between $J(K, g)$ and $J(K', g')$ in terms of the advantage function.

Lemma 5. For any two policies (K, g) and (K', g') , let $\{(x_t, u_t)\}$ and $\{(x'_t, u'_t)\}$ denote the state-action trajectories corresponding to (K, g) and (K', g') respectively, where $u_t = Kx_t + g$ and $u'_t = K'x'_t + g'$. Then,

$$J(K, g) - J(K', g') = -\mathbb{E} \sum_{t=0}^{\infty} A_{K,g}(x'_t, K'x'_t + g).$$

Moreover, defining

$$\begin{aligned} C_K &= RK + \gamma B^\top P(BK + A), \quad d_{K,g} = Rg + \gamma B^\top PBg + \gamma B^\top q, \\ \Sigma^* &= \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t^* (x_t^*)^\top \right], \quad \rho^* = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t^* \right], \quad \beta_\gamma = \sum_{t=0}^{\infty} \gamma^t, \end{aligned}$$

for an optimal policy (K^*, g^*) , we can show that

$$\begin{aligned} &J(K, g) - J(K^*, g^*) \\ &\leq \frac{1}{\sigma_{\min}(R)} \left(\text{tr}(C_K^\top C_K \Sigma^*) + 2 \text{tr}((\rho^*)^\top C_K^\top d_{K,g}) + d_{K,g}^\top d_{K,g} \beta_\gamma \right). \end{aligned}$$

Proof. First consider the difference in cost-to-go, starting from the initial state x_0 .

$$\begin{aligned} V_{K',g'}(x_0) - V_{K,g}(x_0) &= \sum_{t=0}^{\infty} \gamma^t c(x'_t, u'_t) - V_{K,g}(x_0) \\ &= \sum_{t=0}^{\infty} \gamma^t (c(x'_t, u'_t) + V_{K,g}(x'_t) - V_{K,g}(x'_t)) - V_{K,g}(x_0) \\ &= \sum_{t=0}^{\infty} \gamma^t (c(x'_t, u'_t) + \gamma V_{K,g}(x'_{t+1}) - V_{K,g}(x'_t)). \end{aligned}$$

For the last line, we utilized the fact that $x_0 = x'_0$.

Then, we get that

$$V_{K',g'}(x_0) - V_{K,g}(x_0) = \sum_{t=0}^{\infty} \gamma^t A_{K,g}(x'_t, u'_t)$$

We next compute $A_{K,g}(x, K'x + g')$. Recall that

$$V_{K,g}(x) = (x - x^*)^\top Q(x - x^*) + (Kx + g)^\top R(Kx + g) + \gamma V_{K,g}((A + BK)x + Bg).$$

Then, observe that

$$\begin{aligned}
& A_{K,g}(x, K'x + g') \\
&= c(x, K'x + g') + \gamma V_{K,g}(Ax + B(K'x + g')) - V_{K,g}(x) \\
&= (x - x^*)^\top Q(x - x^*) + (K'x + g')^\top R(K'x + g') \\
&\quad + \gamma V_{K,g}(Ax' + B(K'x + g')) - V_{K,g}(x) \\
&= x^\top (Q + (K')^\top RK')x + 2x^\top (-Qx^* + (K')^\top Rg') + g'^\top Rg' \\
&\quad + \gamma ((A + BK')x + Bg')^\top P_K ((A + BK')x + Bg') \\
&\quad + 2\gamma ((A + BK')x + Bg')^\top q_{K,g} + \gamma r_{K,g} - V_{K,g}(x) \\
&= x^\top (Q + (K + K' - K)^\top R(K + K' - K))x \\
&\quad + 2x^\top (-Qx^* + (K + K' - K)^\top R(g + g' - g)) + (g + g' - g)^\top R(g + g' - g) \\
&\quad + \gamma ((A + B(K + K' - K))x + B(g + g' - g))^\top P_K ((A + B(K + K' - K))x + B(g + g' - g)) \\
&\quad + 2\gamma ((A + B(K + K' - K))x + B(g + g' - g))^\top q_{K,g} + \gamma r_{K,g} - V_{K,g}(x) \\
&= x^\top (2(K' - K)^\top RK + (K' - K)^\top R(K' - K))x \\
&\quad + 2x^\top ((K' - K)^\top Rg' + K^\top R(g' - g)) \\
&\quad + 2((g' - g)^\top Rg) + (g' - g)^\top R(g' - g) \\
&\quad + \gamma \left[2x^\top (B(K' - K))^\top P_K (A + BK)x + x^\top (B(K' - K))^\top P_K (B(K' - K))x \right] \\
&\quad + \gamma \left[2x^\top (B(K' - K))^\top P_K (Bg') + 2x^\top (A + BK)^\top P_K (B(g' - g)) \right] \\
&\quad + \gamma \left[2x^\top (B(K' - K))^\top q_{K,g} \right] \\
&\quad + \gamma \left[2(B(g' - g))^\top P_K (Bg) + (B(g' - g))^\top P_K (B(g' - g)) + 2(B(g' - g))^\top q_{K,g} \right] \\
&= 2 \operatorname{tr}(xx^\top (K' - K)^\top RK) + 2 \operatorname{tr}(x^\top (K' - K)^\top R(g' - g)) + 2 \operatorname{tr}(x^\top (K' - K)^\top Rg) \\
&\quad + 2\gamma \operatorname{tr}(x^\top (K' - K)^\top B^\top P_K (BK + A)x) + 2\gamma \operatorname{tr}(x^\top (K' - K)^\top B^\top P_K B(g' - g)) \\
&\quad + 2\gamma \operatorname{tr}(x^\top (K' - K)^\top B^\top P_K Bg) + 2\gamma \operatorname{tr}(x^\top (K' - K)^\top B^\top q_{K,g}) \\
&\quad + \operatorname{tr}(xx^\top (K' - K)^\top (R + \gamma B^\top P_K B)(K' - K)) \\
&\quad + 2 \operatorname{tr}(x^\top K^\top R(g' - g)) + 2 \operatorname{tr}(g^\top R(g' - g)) + 2\gamma \operatorname{tr}((g' - g)^\top B^\top P_K (A + BK)x) \\
&\quad + 2\gamma \operatorname{tr}((g' - g)^\top B^\top P_K Bg) + 2\gamma \operatorname{tr}((g' - g)^\top B^\top q_{K,g}) \\
&\quad + (g' - g)^\top (R + \gamma B^\top P_K B)(g' - g)
\end{aligned}$$

Via completing the square, defining

$$C_K = RK + \gamma B^\top P_K (BK + A), \quad d_{K,g} = Rg + \gamma B^\top P_K Bg + \gamma B^\top q_{K,g},$$

using the fact that $(R + \gamma B^\top P_K B)$ is positive definite, we find that the advantage

function satisfies the following inequality:

$$\begin{aligned}
& A_{K,g}(x, K'x + g') \\
&= \left((K' - K)x + (g' - g) + (R + \gamma B^\top P_K B)^{-1} (C_K x + d_{K,g}) \right)^\top (R + \gamma B^\top P_K B g) \\
&\quad \left((K' - K)x + (g' - g) + (R + \gamma B^\top P_K B)^{-1} (C_K x + d_{K,g}) \right) \\
&\quad - (C_K x + d_{K,g})^\top (R + \gamma B^\top P_K B)^{-1} (C_K x + d_{K,g}) \\
&\geq - (C_K x + d_{K,g})^\top (R + \gamma B^\top P_K B)^{-1} (C_K x + d_{K,g})
\end{aligned}$$

Using the inequality for the advantage, and taking expectations, we find that

$$J(K', g') - J(K, g) \geq -\mathbb{E} \sum_{t=0}^{\infty} \gamma^t \left[(C_K x'_t + d_{K,g})^\top (R + \gamma B^\top P_K B)^{-1} (C_K x'_t + d_{K,g}) \right].$$

Applying this to an optimal policy (K^*, g^*) , we find that

$$\begin{aligned}
& J(K, g) - J(K^*, g^*) \\
&\leq \mathbb{E} \sum_{t=0}^{\infty} \gamma^t \left[(C_K x_t^* + d_{K,g})^\top (R + \gamma B^\top P_K B)^{-1} (C_K x_t^* + d_{K,g}) \right] \\
&\leq \frac{1}{\sigma_{\min}(R)} \mathbb{E} \sum_{t=0}^{\infty} \gamma^t \left[(C_K x_t^* + d_{K,g})^\top (C_K x_t^* + d_{K,g}) \right] \\
&= \frac{1}{\sigma_{\min}(R)} \left(\text{tr}(C_K^\top C_K \mathbb{E} \sum_{t=0}^{\infty} \gamma^t x_t^* (x_t^*)^\top) + 2 \text{tr}(\mathbb{E} \sum_{t=0}^{\infty} \gamma^t x_t^* (x_t^*)^\top C_K^\top d_{K,g}) + d_{K,g}^\top d_{K,g} \sum_{t=0}^{\infty} \gamma^t \right)
\end{aligned}$$

□

Before finally establishing gradient dominance, we first need the following elementary linear algebra fact.

Lemma 6. *Suppose $g \in \mathbb{R}^n$ for some $n \in \mathbb{Z}_+$. Then, for any positive semi-definite (PSD) n by n matrix Σ , the $(n+1)$ by $(n+1)$ matrix*

$$M := \begin{bmatrix} gg^\top + \Sigma & g \\ g^\top & 1 \end{bmatrix}$$

is PSD.

Proof. Observe that

$$M = \begin{bmatrix} \Sigma & 0 \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} g \\ 1 \end{bmatrix} \begin{bmatrix} g^\top & 1 \end{bmatrix},$$

where the first summand matrix is PSD, and the second summand is a rank-one matrix outer product and hence also PSD. Thus, M is PSD since the sum of two PSD matrices is PSD. □

Proposition 5 (Gradient dominance of LQR with tracking). *Define*

$$C^* = \max \left\{ \|\Sigma^*\|_2 + \|\rho^*\|_2, \sum_{t=0}^{\infty} \gamma^t + \|\rho^*\|_2 \right\},$$

where

$$\Sigma^* := \mathbb{E} \sum_{t=0}^{\infty} \gamma^t x_t^* (x_t^*)^\top, \quad \rho^* := \mathbb{E} \sum_{t=0}^{\infty} \gamma^t x_t^*$$

are the discounted state covariance matrix and mean vector respectively corresponding to an optimal policy. Suppose

$$\mathbb{E}[x_0] = 0, \quad \mathbb{E}[x_0 x_0^\top] = \Sigma \succeq \alpha I \text{ for some } \alpha > 0, \quad |\rho(A + BK)| < 1.$$

Then,

$$J(K, g) - J(K^*, g^*) \leq \frac{C^*}{4 \min\{\alpha, 1\}^2 \sigma_{\min}(R)} \|\nabla J\|_F^2.$$

We also have a lower bound:

$$J(K, g) - J(K^*, g^*) \geq \frac{\min\{\alpha, 1\}}{\|R + B^\top P_K B\|_2} \text{tr}(E_{K,g}^\top E_{K,g})$$

Proof. From Lemma 5, using our definition of C^* , we get that

$$\begin{aligned} & J(K, g) - J(K^*, g^*) \\ & \leq \frac{1}{\sigma_{\min}(R)} \left(\text{tr}(C_K^\top C_K \mathbb{E} \sum_{t=0}^{\infty} \gamma^t x_t^* (x_t^*)^\top) + 2 \text{tr}((\mathbb{E} \sum_{t=0}^{\infty} \gamma^t x_t^*)^\top C_K^\top d_{K,g}) + d_{K,g}^\top d_{K,g} \sum_{t=0}^{\infty} \gamma^t \right) \\ & \leq \frac{C^*}{\sigma_{\min}(R)} \left(\|C_K\|_F^2 + \|d_{K,g}\|_F^2 \right). \end{aligned}$$

Define the quantities

$$\Sigma_{K,g} = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t x_t^\top \right], \quad \rho_{K,g} = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t \right], \quad \beta_\gamma = \sum_{t=0}^{\infty} \gamma^t,$$

where the state trajectory $\{x_t\}$ is generated by the policy (K, g) . Then, based on Lemma 4, we have the following expression for $\|\nabla J(K, g)\|_F^2$,

$$\begin{aligned} \|\nabla J(K, g)\|_F^2 &= 4 \left\| \begin{bmatrix} C_K \Sigma_{K,g} + d_{K,g} \rho_{K,g}^\top & C_K \rho_{K,g} + \beta_\gamma d_{K,g} \end{bmatrix} \right\|_F^2 \\ &= 4 \left\| \begin{bmatrix} C_K & d_{K,g} \end{bmatrix} \begin{bmatrix} \Sigma_{K,g} & \rho_{K,g} \\ \rho_{K,g}^\top & \beta_\gamma \end{bmatrix} \right\|_F^2 \\ &= 4 \text{tr}(E_{K,g}^\top E_{K,g} M_{K,g} M_{K,g}^\top), \end{aligned}$$

where

$$E_{K,g} := [C_K \quad d_{K,g}], \quad M_{K,g} := \begin{bmatrix} \Sigma_{K,g}^\top & \rho_{K,g} \\ \rho_{K,g} & \beta_\gamma \end{bmatrix}.$$

We note that

$$\begin{aligned} M_{K,g} &:= \begin{bmatrix} \Sigma_{K,g}^\top & \rho_{K,g} \\ \rho_{K,g} & \beta_\gamma \end{bmatrix} \\ &= \sum_{t=0}^{\infty} \gamma^t \begin{bmatrix} \mathbb{E}[x_t x_t^\top] & \mathbb{E}[x_t] \\ \mathbb{E}[x_t]^\top & 1 \end{bmatrix} \\ &= \sum_{t=0}^{\infty} \gamma^t \begin{bmatrix} \mathbb{E}[x_t] \mathbb{E}[x_t]^\top + \mathbb{E}[(A+BK)^t x_0 x_0^\top ((A+BK)^t)^\top] & \mathbb{E}[x_t] \\ \mathbb{E}[x_t]^\top & 1 \end{bmatrix} \\ &\succeq \begin{bmatrix} \Sigma & 0 \\ 0 & 1 \end{bmatrix} \\ &\succeq \alpha \begin{bmatrix} I_{n \times n} & 0 \\ 0 & 1 \end{bmatrix}. \end{aligned}$$

where we used Lemma 6 for the penultimate matrix inequality, and the assumptions $\mathbb{E}[x_0] = 0$ and $\mathbb{E}[x_0 x_0^\top] \succeq \alpha I$ for the final matrix inequality. This implies then that

$$\begin{aligned} \|\nabla J(K, g)\|_F^2 &= 4 \operatorname{tr}(E_{K,g}^\top E_{K,g} M_{K,g} M_{K,g}^\top) \\ &\geq 4 \min\{\alpha, 1\}^2 \operatorname{tr}(E_{K,g}^\top E_{K,g}). \end{aligned}$$

Therefore,

$$\begin{aligned} J(K, g) - J(K^*, g^*) &\leq \frac{C^*}{\sigma_{\min}(R)} \left(\|C_K\|_F^2 + \|d_{K,g}\|_F^2 \right) \\ &= \frac{C^*}{\sigma_{\min}(R)} \operatorname{tr}(E_{K,g}^\top E_{K,g}) \\ &\leq \frac{C^*}{4 \min\{\alpha, 1\}^2 \sigma_{\min}(R)} \|\nabla J(K, g)\|_F^2 \end{aligned}$$

For the lower bound, pick $K' = K - (R + \gamma B^\top P_K B)^{-1} C_K$, and $g' = g' - (R + \gamma B^\top P_K B)^{-1} d_{K,g}$. Then, following the analysis of the advantage function in Lemma 5, we have

$$\begin{aligned} J(K, g) - J(K^*, g^*) &\geq J(K, g) - J(K', g') \\ &= -\mathbb{E} \sum_{t=0}^{\infty} A_{K,g}(x'_t, K' x'_t + g') \\ &= \mathbb{E} \sum_{t=0}^{\infty} \gamma^t [(C_K x'_t + d_{K,g})^\top (R + \gamma B^\top P_K B)^{-1} (C_K x'_t + d_{K,g})] \\ &= \operatorname{tr}((E_{K,g} M_{K',g'})^\top (R + B^\top P_K B)^{-1} E_{K,g} M_{K',g'}) \\ &\geq \frac{\min\{\alpha, 1\}}{\|R + B^\top P_K B\|_2} \operatorname{tr}(E_{K,g}^\top E_{K,g}) \end{aligned}$$

□

We next show that the distributed LQR tracking, considered in Problem 3, is also gradient dominant. We hence focus our attention on $\bar{J}(K, g)$.

Lemma 7 (Gradient of coupled distributed LQR tracking cost). *For a policy $(K, G) = (K, g^1, g^2, \dots, g^m)$, we have*

$$\begin{aligned}\nabla_K \bar{J}(K, G) &= \sum_{i=1}^m 2 \left(C_K \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t^i (x_t^i)^\top \right] + d_{K, g^i} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t (x_t^i)^\top \right] \right). \\ \nabla_{g^i} \bar{J}(K, G) &= 2 \left(C_K \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t^i \right] + d_{K, g^i} \sum_{t=0}^{\infty} \gamma^t \right) \quad \forall i \in [m].\end{aligned}$$

Proof. For any $x = (x^1, x^2, \dots, x^m) \in \mathbb{R}^{mn}$, where for each $i \in [m]$, $x^i \in \mathbb{R}^n$, define the average value function

$$\begin{aligned}V_{K, G}(x) &:= \sum_{i=1}^m V_{K, g^i}(x^i) \\ &= \sum_{i=1}^m (x^i)^\top P_K x^i + 2(x^i)^\top q_{K, g^i} + r_{K, g^i}.\end{aligned}$$

Then, supposing that for each $i \in [m]$, $x^i \stackrel{iid}{\sim} N(0, \Sigma)$, we find that

$$\bar{J}(K, G) = \mathbb{E}_x V_{K, G}(x).$$

Then, differentiating, we get

$$\begin{aligned}\nabla_K V_{K, G}(x) &= \sum_{i=1}^m \nabla_K V_{K, g^i}(x^i), \\ \nabla_{g^i} V_{K, G}(x) &= \nabla_{g^i} V_{K, g^i}(x^i) \quad \forall i \in [m].\end{aligned}$$

Taking expectations, using our result for the gradient of a single-agent LQR tracking problem in Lemma 4, we find that

$$\begin{aligned}\nabla_K \bar{J}(K, G) &= \sum_{i=1}^m 2 \left(C_K \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t^i (x_t^i)^\top \right] + d_{K, g^i} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t (x_t^i)^\top \right] \right). \\ \nabla_{g^i} \bar{J}(K, G) &= 2 \left(C_K \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t^i \right] + d_{K, g^i} \sum_{t=0}^{\infty} \gamma^t \right) \quad \forall i \in [m].\end{aligned}$$

□

Our next result shows that gradient dominance is preserved for the coupled LQR tracking problem.

Proposition 6 (Gradient dominance for coupled LQR tracking problem). *Define*

$$C^{i*} = \max \left\{ \|\Sigma^*\|_2 + \|\rho^{i*}\|_2, \sum_{t=0}^{\infty} \gamma^t + \|\rho^{i*}\|_2 \right\},$$

where

$$\Sigma^{i*} := \mathbb{E} \sum_{t=0}^{\infty} \gamma^t x_t^{i*} (x_t^{i*})^\top, \quad \rho^{i*} := \mathbb{E} \sum_{t=0}^{\infty} \gamma^t x_t^{i*}$$

are the discounted state covariance matrix and mean vector respectively corresponding to an optimal policy for agent i . Define then

$$C^* = \max_{i \in m} C^{i*}.$$

Suppose for each $i \in [m]$ that

$$\mathbb{E}[x_0^i] = 0, \quad \mathbb{E}[x_0^i (x_0^i)^\top] = \Sigma \succeq \alpha I \text{ for some } \alpha > 0, \quad |\rho(A + BK)| < 1.$$

Then,

$$\bar{J}(K, G) - \bar{J}(K^*, G^*) \leq \frac{C^*}{4 \min\{\alpha, 1\}^2 \sigma_{\min}(R)} \|\nabla \bar{J}\|_F^2.$$

Proof. Building on our gradient dominance analysis for single-agent LQR tracking in Proposition 5, we find that

$$\begin{aligned} \bar{J}(K, G) - \bar{J}(K^*, G^*) &= \sum_{i=1}^m J(K, g^i) - J(K^*, (g^i)^*) \\ &\leq \sum_{i=1}^m \frac{C^*}{\sigma_{\min}(R)} \left(\|C_K\|_F^2 + \|d_{K, g^i}\|_F^2 \right). \end{aligned}$$

Meanwhile, defining for convenience

$$\Sigma_{K, g^i}^i := \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t^i (x_t^i)^\top \right] \quad \text{and} \quad \rho_{K, g^i}^i := \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t (x_t^i)^\top \right] \quad \forall i \in [m], \quad \beta_\gamma = \sum_{t=0}^{\infty} \gamma^t,$$

using the result in Lemma 7, we have

$$\begin{aligned} &\|\nabla J(K, G)\|_F^2 \\ &= \|\nabla_K \bar{J}(K, G)\|_F^2 + \sum_{i=1}^m \|\nabla_{g^i} \bar{J}(K, G)\|_F^2 \\ &= 4 \left\| \begin{bmatrix} C_K & d_{K, g^1} & d_{K, g^2} & \dots & d_{K, g^m} \end{bmatrix} \begin{bmatrix} \sum_{i=1}^m \Sigma_{K, g^i}^i & \rho_{K, g^1}^1 & \rho_{K, g^2}^2 & \dots & \rho_{K, g^m}^m \\ (\rho_{K, g^1}^1)^\top & 1 & 0 & \dots & 0 \\ (\rho_{K, g^2}^2)^\top & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ (\rho_{K, g^m}^m)^\top & 0 & 0 & \dots & 1 \end{bmatrix} \right\|_F^2 \\ &= 4 \text{tr}(E_{K, G}^\top E_{K, G} M_{K, G} M_{K, G}^\top), \end{aligned}$$

where

$$E_{K,G} := [C_K \quad d_{K,g^1} \quad d_{K,g^2} \quad \cdots \quad d_{K,g^m}],$$

$$M_{K,G} := \begin{bmatrix} \sum_{i=1}^m \Sigma_{K,g^i}^i & \rho_{K,g^1}^1 & \rho_{K,g^2}^2 & \cdots & \rho_{K,g^m}^m \\ (\rho_{K,g^1}^1)^\top & 1 & 0 & \cdots & 0 \\ (\rho_{K,g^2}^2)^\top & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ (\rho_{K,g^m}^m)^\top & 0 & 0 & \cdots & 1 \end{bmatrix}.$$

Using Lemma 6, and the assumptions $\mathbb{E}[x_0^i] = 0$ and $\mathbb{E}[x_0^i(x_0^i)^\top] = \Sigma \geq \alpha I$ for each i , we can show that

$$M_{K,G} \succeq \begin{bmatrix} m\alpha I_{n \times n} & 0_{n \times m} \\ 0_{m \times n} & I_{m \times m} \end{bmatrix}.$$

This implies then that

$$\begin{aligned} \|\nabla \bar{J}(K, G)\|_F^2 &= 4 \operatorname{tr}(E_{K,G}^\top E_{K,G} M_{K,G} M_{K,G}^\top) \\ &\geq 4 \left(m^2 \alpha^2 \|C_K\|_F^2 + \sum_{i=1}^m \|d_{K,g^i}\|_F^2 \right) \\ &\geq 4 \min\{\alpha, 1\}^2 \left(\sum_{i=1}^m \|C_K\|_F^2 + \|d_{K,g^i}\|_F^2 \right) \end{aligned}$$

Therefore,

$$\begin{aligned} \bar{J}(K, G) - \bar{J}(K^*, G^*) &\leq \sum_{i=1}^m \frac{C^*}{\sigma_{\min}(R)} \left(\|C_K\|_F^2 + \|d_{K,g^i}\|_F^2 \right) \\ &\leq \frac{C^*}{4 \min\{\alpha, 1\}^2 \sigma_{\min}(R)} \|\nabla \bar{J}(K, G)\|_F^2. \end{aligned}$$

□

The preceding work allows us to prove Proposition 6.

Proof of Proposition 6. By Proposition 5, we have that

$$J^i(K, g^i) \leq \mu^i \|\nabla J^i(K, g^i)\|^2,$$

where

$$\mu^i = \frac{C^{i*}}{4 \min\{\alpha, 1\}^2 \sigma_{\min}(R)},$$

and C^{i*} is as defined in Proposition 6. In addition, by Proposition 6, we see that

$$\bar{J}(K, G) - \bar{J}(K^*, G^*) \leq \max_{i \in [m]} \mu^i \|\nabla \bar{J}(K, G)\|^2$$

□

B.4 Proposition 2: Local Lipschitzness and smoothness

In the sequel, without loss of generality, we assume that the covariance of the initial noise distribution, Σ , is the identity matrix $I_{n \times n}$.

We first define $C(K)$ to be the following cost, a standard LQR with no tracking based on the system matrices, cost matrices and discount factor γ .

$$\begin{aligned} C(K) &= \sum_{t=0}^{\infty} x_t^\top Q x_t + u_t^\top R u_t, \\ \text{such that } x_{t+1} &= \sqrt{\gamma} A x_t + \sqrt{\gamma} B u_t, \\ u_t &= K x_t. \end{aligned} \tag{8}$$

This lemma shows that if $J(K, g) \leq C$ for some $C > 0$, then $C(K) \leq C$.

Lemma 8. *Suppose $J(K, g) \leq C_0$ for some $C_0 > 0$, and that $\mathbb{E}[x_0] = 0$. Then, $C(K) \leq C_0$.*

Proof. Based on lemma 3, we see that

$$\begin{aligned} J(K, g) &= \text{tr}(\Sigma P_K) + \mathbb{E}[x_0]^\top q_{K,g} + r_{K,g}, \\ &= \text{tr}(\Sigma P_K) + r_{K,g} = C(K) + r_{K,g}, \end{aligned}$$

where we used $\mathbb{E}[x_0] = 0$ as well as the easy-to-check fact that $C(K) = \text{tr}(\Sigma P_K)$. Since $r_{K,g} \geq 0$ has to hold — otherwise the value function $V(K, g; 0)$, *i.e.* value function of (K, g) when $x_0 = 0$, is negative —, it follows that $J(K, g) = C(K) + r_{K,g} \leq C_0 \implies C(K) \leq C_0$. \square

Next we show that if $J(K, g)$ is bounded by C_0 , then P_K , K and g are all bounded by constants involving C_0 and system parameters.

Lemma 9. *Suppose $J(K, g) \leq C_0$ for some $C_0 > 0$. Then, $P_K \leq C_0$, and moreover, there exists constants $c_1, c_2 > 0$ depending on C_0 and A, B, Q, R, γ such that*

$$\|K\| \leq c_1, \|g\| \leq c_2.$$

Proof. That $P_K \leq C_0$ holds follows from $C(K) \leq C_0$ and $\Sigma = I$. Next, by Lemma 25 of [7], since $C(K) \leq C_0$, there exists c_1 depending on C_0 and system parameters such that $\|K\| \leq c_1$.

Finally we show that g is bounded. To this end, note that the cost at $t = 0$ is bounded by C_0 , namely

$$\begin{aligned} \mathbb{E}[(x_0 - x^*)^\top Q(x_0 - x^*) + (Kx_0 + g)^\top R(Kx_0 + g)] &\leq C_0 \\ \implies \text{tr}(Q) + (x^*)^\top Q x^* + x_0^\top (K^\top R K) x_0 + g^\top R g &\leq C_0 \\ \implies g^\top R g &\leq C_0 \\ \implies \|g\| &\leq \sqrt{C_0 / \sigma_{\min}(R)} := c_2. \end{aligned}$$

\square

Next, define $S_K := (I - \gamma(A + BK)^\top)^{-1}$. We show that S_K is bounded, and that the difference $\|S_{K'} - S_K\|$ is bounded by $\|K' - K\|$ times some constant.

Lemma 10. *Define $S_K := (I - \gamma(A + BK)^\top)^{-1}$. Suppose $J(K, g) \leq C_0$. Then,*

$$\|S_K\|_2 \leq \frac{2C_0}{\sigma_{\min}(Q)} := c_3.$$

Proof. From the argument in Lemma 13 of [7], it holds that

$$\text{tr}(\Sigma_K) \leq \frac{C(K)}{\sigma_{\min}(Q)},$$

where

$$\Sigma_K = \mathbb{E} \sum_{t=0}^{\infty} x_t x_t^\top,$$

where $\{x_t\}$ is generated by the policy $u_t = Kx_t$, $x_{t+1} = \sqrt{\gamma}Ax_t + \sqrt{\gamma}Bu_t$, and $x_0 \sim N(0, \Sigma)$. We first show that $\|S_K\|_2 \leq \frac{2C(K)}{\sigma_{\min}(Q)}$. To this end, since $\sqrt{\gamma}\rho(A + BK) < 1$, by the same argument as in Lemma 23 of [7], we find that

$$\begin{aligned} \|S_K\|_2 &\leq \frac{1}{1 - \sqrt{\gamma}\rho(\sqrt{\gamma}A + \sqrt{\gamma}BK)} \\ &\leq \frac{1}{1 - \rho(\sqrt{\gamma}A + \sqrt{\gamma}BK)} \\ &\leq 2 \text{tr}(\Sigma_K) \\ &\leq \frac{2C(K)}{\sigma_{\min}(Q)}. \end{aligned}$$

□

Next, we seek to bound $(S_{K'} - S_K)$ — to be precise, we show when $K - K'$ is small enough, $\|S_{K'} - S_K\|_2$ can be bounded in terms of $\|K' - K\|$.

Lemma 11. *Suppose $\|S_K\|_2(\gamma\|B\|_2\|K - K'\|) \leq \frac{1}{2}$, and $\gamma\rho(A + BK') < 1$ (so $S_{K'}$ is well-defined). Then,*

$$\|S_{K'} - S_K\|_2 \leq 2\|S_K\|_2\gamma\|B\|_2\|K' - K\|.$$

Proof. For the purposes of this proof, we define

$$C := I - \gamma(A + BK), \quad D := \gamma B(K - K').$$

Then, we have

$$\begin{aligned} S_{K'} - S_K &= (C + D)^{-1} - C^{-1} \\ &= C^{-1} (I + C^{-1}D)^{-1} - C^{-1} \\ &= C^{-1} ((I + C^{-1}D)^{-1} - I) \end{aligned}$$

Next, consider $I + C^{-1}D$. Observe that

$$\begin{aligned} (I + C^{-1}D)^{-1} &= I - C^{-1}D(I + C^{-1}D)^{-1} \\ \implies \|I + C^{-1}D\| &\leq 2, \end{aligned}$$

where the first can be verified by multiplying both sides with $I + C^{-1}D$, and the second is a consequence of $\|C^{-1}D\|_2 \leq \|S_K\|_2(\gamma\|B\|_2\|K' - K\|) \leq 1/2$, our assumption. Therefore, using some algebraic manipulation, we have that

$$\begin{aligned} S_{K'} - S_K &= C^{-1}((I + C^{-1}D)^{-1} - I) \\ &= C^{-1}(C^{-1}D(I + C^{-1}D)^{-1}), \end{aligned}$$

which implies that

$$\begin{aligned} \|S_{K'} - S_K\|_2 &= \|C^{-1}(C^{-1}D(I + C^{-1}D)^{-1})\|_2 \\ &\leq 2\gamma\|S_K\|_2^2\|B\|\|K' - K\| \end{aligned}$$

□

Next, we show that if K' is within a ball around K , then we must have $\gamma\rho(A + BK') < 1$.

Lemma 12. *Suppose*

$$\|K' - K\| \leq \frac{\sigma_{\min}(Q)}{4C(K)\|\sqrt{\gamma}B\|_2(\|\sqrt{\gamma}A + \sqrt{\gamma}BK\|_2 + 1)}.$$

Then, $\sqrt{\gamma}(A + BK') < 1$.

Proof. This follows from the same argument in Lemma 22 of [7], but with μ in that proof replaced with 1, and the system (A, B) replace by $(\sqrt{\gamma}A, \sqrt{\gamma}B)$. □

Finally, we explicitly state the condition on $(K' - K)$ such that $\gamma\|S_K\|_2\|B\|_2\|K' - K\| \leq \frac{1}{2}$.

Lemma 13. *Suppose $\|K' - K\| \leq \frac{\sigma_{\min}(Q)}{4\gamma\|B\|C(K)}$. Then,*

$$\gamma\|S_K\|_2\|B\|_2\|K' - K\| \leq \frac{1}{2}.$$

At this point, we take a chance to state a few results establishing $C(K)$ is locally Lipschitz. These results are based on [7] and [20]. Similar to [20], we assume boundedness of the initial distribution of x_0 in the proofs, noting that an extension to sub-Gaussian random distribution is possible by appealing to high-probability bounds and standard truncation arguments. Thus, in the sequel, we will work with the assumption $\|x_0\|^2 \leq C_n$ for some $C_n > 0$.

Lemma 14. *Suppose $C(K) \leq C_0$. Suppose $\|K' - K\| \leq c_4$. Then,*

$$\begin{aligned} |C(K') - C(K)| &\leq \frac{n}{C_n} c_{11} \|K' - K\|, \\ |C(K', x_0) - C(K, x_0)| &\leq c_{11} \|K' - K\|, \\ \|P_{K'} - P_K\|_2 &\leq (c_8 + c_9 + c_{10}) \|K' - K\|. \end{aligned}$$

By Lemma 13 of [7], it follows that $\|P_K\| \leq C(K)$.

Next, we prove a bound for $\|q_{K,g}\|$.

Lemma 15. *Suppose $\sqrt{\gamma}(A + BK) < 1$, and $J(K, g) < C_0$. Suppose also that $\|K' - K\| \leq c_4$. Then,*

$$\|q_{K,g}\| \leq c_{12}$$

for an explicit constant $c_{12} > 0$ involving C_0 and the system parameters.

Proof. We have

$$q_{K,g} = S_K(-Qx^* + K^\top Rg + \gamma(A + BK)^\top P_K Bg)$$

Observe that $\|-Qx^*\| \leq \|Q\|_2 \|x^*\|$. Next, note that

$$\begin{aligned} \|K^\top Rg\| &\leq \|K\| \|R\| \|g\| \\ &\leq c_1 c_2 \|R\|. \end{aligned}$$

Next, observe that

$$\begin{aligned} \gamma \|(A + BK)^\top P_K Bg\| &\leq \sqrt{\gamma} \|P_K\| \|B\|_2 \|g\| \\ &\leq \sqrt{\gamma} C_0 \|B\|_2 c_2. \end{aligned}$$

This implies that

$$\begin{aligned} \|q_{K,g}\| &= \|S_K(-Qx^* + K^\top Rg + \gamma(A + BK)^\top P_K Bg)\| \\ &\leq \|S_K\|_2 \|-Qx^* + K^\top Rg + \gamma(A + BK)^\top P_K Bg\| \\ &\leq c_3 (c_1 c_2 \|R\| + \sqrt{\gamma} C_0 \|B\|_2 c_2 + \sqrt{\gamma} C_0 \|B\|_2 c_2) := c_{12}. \end{aligned}$$

□

Next, we provide a bound for the quantity $(q_{K',g'} - q_{K,g})$.

Lemma 16. *Suppose $\sqrt{\gamma}\rho(A + BK) < 1$, and $J(K, g) < C_0$. Suppose also that $\|K' - K\| \leq c_4$, $\|g' - g\| \leq 1$. Then,*

$$\|q_{K',g'} - q_{K,g}\| \leq c_{15} \|K' - K\| + c_{16} \|g' - g\|,$$

where

$$c_{15} := 2\gamma c_3 \|B\|_2 \|Qx^*\| + c_{13} + c_{14}, \quad c_{16} := 2c_1 c_3 \|R\|_2 + c_3 C_0 \|B\|_2$$

and $c_{13}, c_{14} > 0$ are constants depending on C_0 and the system parameters, introduced in the proof later.

Proof. We have

$$\begin{aligned}
& q_{K',g'} - q_{K,g} \\
&= S_{K'}(-Qx^* + (K')^\top Rg' + \gamma(A + BK')^\top P_{K'}Bg') \\
&\quad - S_K(-Qx^* + K^\top Rg + \gamma(A + BK)^\top P_KBg) \\
&= (S_{K'} - S_K)(-Qx^*) + (S_{K'}(K')^\top Rg' - S_KK^\top Rg) + (S_{K'}\gamma(A + BK')^\top P_{K'}Bg' - S_K\gamma(A + BK)^\top P_KBg)
\end{aligned}$$

We treat the three terms separately. First, observe that

$$\begin{aligned}
\|(S_{K'} - S_K)(-Qx^*)\| &\leq (2\gamma\|S_K\|_2\|B\|_2\|K' - K\|)\|Q\|_2\|x^*\| \\
&\leq 2\gamma c_3\|B\|_2\|Q\|_2\|x^*\|\|K' - K\|.
\end{aligned}$$

Next, observe that

$$\begin{aligned}
& \|S_{K'}(K')^\top Rg' - S_KK^\top Rg\| \\
&= \|(S_{K'} - S_K)((K')^\top Rg') + S_K((K')^\top Rg') - S_KK^\top Rg\| \\
&= \|(S_{K'} - S_K)((K')^\top Rg') + S_K((K')^\top Rg' - K^\top Rg) + S_KK^\top Rg - S_K^\top K^\top Rg\| \\
&\leq 2\gamma c_3\|B\|_2\|K' - K\|(\|K\| + 1)\|R\|_2\|g\| + 1 \\
&\quad + 2c_3(\|K^\top R(g' - g) + (K' - K)^\top Rg + (K' - K)^\top R(g' - g)\|) \\
&\leq 2\gamma c_3\|B\|_2(c_1 + 1)\|R\|_2(c_2 + 1)\|K' - K\| \\
&\quad + 2c_3(\|K\|\|R\|_2\|g' - g\| + \|K' - K\|\|R\|_2\|g\| + \|K' - K\|\|R\|_2\|g' - g\|) \\
&\leq 2\gamma c_3\|B\|_2(c_1 + 1)\|R\|_2(c_2 + 1)\|K' - K\| \\
&\quad + 2c_3(c_1\|R\|_2\|g' - g\| + \|K' - K\|\|R\|_2c_2 + \|R\|_2\|K' - K\|) \\
&= (2\gamma c_3\|B\|_2(c_1 + 1)\|R\|_2(c_2 + 1) + \|R\|_2c_2 + \|R\|_2)\|K' - K\| + 2c_1c_3\|R\|_2\|g' - g\| \\
&= (2\gamma c_3\|B\|_2(c_1 + 1)\|R\|_2(c_2 + 1) + (c_2 + 1)\|R\|_2)\|K' - K\| + 2c_1c_3\|R\|_2\|g' - g\|
\end{aligned}$$

We next handle $(S_{K'}\gamma(A + BK')^\top P_{K'}Bg' - S_K\gamma(A + BK)^\top P_KBg)$. By a similar algebraic manipulation to the immediately preceding set of derivations, we find that

$$\begin{aligned}
& \|S_{K'}\gamma(A + BK')^\top P_{K'}Bg' - S_K\gamma(A + BK)^\top P_KBg\| \\
&= \|(S_{K'} - S_K)\gamma(A + BK')^\top P_{K'}Bg' + S_K(\gamma(A + BK')^\top P_{K'}Bg' - \gamma(A + BK)^\top P_KBg)\| \\
&\leq \gamma\|S_{K'} - S_K\|\|(A + BK')^\top P_{K'}Bg'\| + \|S_K\|_2\|\gamma(A + BK')^\top P_{K'}Bg' - \gamma(A + BK)^\top P_KBg\|.
\end{aligned}$$

We first handle $\gamma\|S_{K'} - S_K\|\|(A + BK')^\top P_{K'}Bg'\|$. This involves bounding $\sqrt{\gamma}\|(A + BK')^\top P_{K'}Bg'\|$. To this end, observe that

$$\begin{aligned}
& \sqrt{\gamma}\|(A + BK')^\top P_{K'}Bg'\| \\
&\leq \|P_{K'}\|_2\|B\|_2\|g'\| \\
&= \|P_K + P_{K'} - P_K\|_2\|B\|_2\|g'\| \\
&\leq (\|P_K\|_2 + (c_8 + c_9 + c_{10})\|K' - K\|)\|B\|_2(c_2 + 1) \\
&\leq (C_0 + c_8 + c_9 + c_{10})\|B\|_2(c_2 + 1)
\end{aligned}$$

This gives us

$$\begin{aligned}\gamma\|S_{K'} - S_K\| \|(A + BK')^\top P_{K'} Bg'\| &\leq \sqrt{\gamma}(C_0 + c_8 + c_9 + c_{10})\|B\|_2(c_2 + 1)\|S_{K'} - S_K\| \\ &\leq \sqrt{\gamma}(C_0 + c_8 + c_9 + c_{10})\|B\|_2(c_2 + 1)2\gamma\|S_K\|_2\|B\|_2\|K' - K\| \\ &\leq 2\gamma^{3/2}(C_0 + c_8 + c_9 + c_{10})\|B\|_2^2(c_2 + 1)\|K' - K\|\end{aligned}$$

Next, we handle $\|\gamma(A + BK')^\top P_{K'} Bg' - \gamma(A + BK)^\top P_K Bg\|$. Observe that

$$\begin{aligned}&\|\gamma(A + BK')^\top P_{K'} Bg' - \gamma(A + BK)^\top P_K Bg\| \\ &= \|\gamma((A + BK')^\top - (A + BK)^\top)P_{K'} Bg' + \gamma(A + BK)^\top(P_{K'} Bg' - P_K Bg)\| \\ &= \|\gamma((A + BK')^\top - (A + BK)^\top)\| \|P_{K'} Bg'\| + \gamma\|(A + BK)^\top((P_{K'} - P_K)Bg' + P_K B(g' - g))\| \\ &\leq \gamma\|B\|_2^2\|K' - K\|(C_0 + c_8 + c_9 + c_{10})\|B\|_2(c_2 + 1) \\ &\quad + \sqrt{\gamma}(\|P_{K'} - P_K\|\|B\|_2(c_2 + 1) + \|P_K\|_2\|B\|_2\|g' - g\|) \\ &\leq \gamma\|B\|_2^2(C_0 + c_8 + c_9 + c_{10})\|B\|_2(c_2 + 1)\|K' - K\| \\ &\quad + \sqrt{\gamma}\|B\|_2(c_2 + 1)(c_8 + c_9 + c_{10})\|K' - K\| + C_0\|B\|_2\|g' - g\| \\ &\leq (\gamma\|B\|_2^2(C_0 + c_8 + c_9 + c_{10})\|B\|_2(c_2 + 1) + \sqrt{\gamma}\|B\|_2(c_2 + 1)(c_8 + c_9 + c_{10}))\|K' - K\| + C_0\|B\|_2\|g' - g\|.\end{aligned}$$

Thus,

$$\begin{aligned}&\|S_K\|_2\|\gamma(A + BK')^\top P_{K'} Bg' - \gamma(A + BK)^\top P_K Bg\| \\ &\leq c_3(\gamma\|B\|_2^2(C_0 + c_8 + c_9 + c_{10})\|B\|_2(c_2 + 1) + \sqrt{\gamma}\|B\|_2(c_2 + 1)(c_8 + c_9 + c_{10}))\|K' - K\| \\ &\quad + c_3C_0\|B\|_2\|g' - g\|.\end{aligned}$$

Define now

$$\begin{aligned}c_{13} &:= 2\gamma c_3\|B\|_2(c_1 + 1)\|R\|_2(c_2 + 1) + (c_2 + 1)\|R\|_2, \\ c_{14} &= 2\gamma^{3/2}(C_0 + c_8 + c_9 + c_{10})\|B\|_2^2(c_2 + 1) \\ &\quad + c_3(\gamma\|B\|_2^2(C_0 + c_8 + c_9 + c_{10})\|B\|_2(c_2 + 1) + \sqrt{\gamma}\|B\|_2(c_2 + 1)(c_8 + c_9 + c_{10}))\end{aligned}$$

Then, combining everything, we find that

$$\|q_{K',g'} - q_{K,g}\| \leq (2\gamma c_3\|B\|_2\|Qx^*\| + c_{13} + c_{14})\|K' - K\| + (2c_1c_3\|R\|_2 + c_3C_0\|B\|_2)\|g' - g\|$$

□

We are now finally ready to bound $(r_{K',g'} - r_{K,g})$, after which we can wrap up our local Lipschitz argument for $J(K, g)$.

Lemma 17. *Suppose $\sqrt{\gamma}\rho(A + BK) < 1$, and $J(K, g) < C_0$. Suppose also that $\|K' - K\| \leq c_4$, $\|g' - g\| \leq 1$. Then,*

$$\|r_{K',g'} - r_{K,g}\| \leq c_{19}\|K' - K\| + c_{20}\|g' - g\|$$

where

$$c_{19} := c_{17} + c_{18},$$

and

$$\begin{aligned} c_{17} &:= \frac{1}{1-\gamma} [2c_2 + 1 + \gamma(\|B\|_2^2(C_0 + c_8 + c_9 + c_{10})(c_2 + 1) + \|B\|_2^2 c_2 C_0)], \\ c_{18} &:= \frac{1}{1-\gamma} [2\gamma(\|B\|_2(c_{12} + c_{15} + c_{16}) + \|B\|_2 c_2 c_{16})], \\ c_{20} &:= \frac{1}{1-\gamma} [\gamma\|B\|_2^2 c_2^2 (c_8 + c_9 + c_{10}) + 2\gamma\|B\|_2 c_2 (c_{15})]. \end{aligned}$$

Proof. Recall that

$$r_{K,g} = \frac{1}{1-\gamma} ((x^*)^\top Q x^* + g^\top R g + \gamma(g^\top B^\top P_K B g + 2g^\top B^\top q_{K,g})).$$

Then,

$$\begin{aligned} r_{K',g'} - r_{K,g} &= \frac{1}{1-\gamma} ((g')^\top R g' + \gamma((g')^\top B^\top P_{K'} B g' + 2(g')^\top B^\top q_{K',g'})) \\ &\quad - \frac{1}{1-\gamma} (g^\top R g + \gamma(g^\top B^\top P_K B g + 2g^\top B^\top q_{K,g})). \end{aligned}$$

We omit the factor $\frac{1}{1-\gamma}$ for now. We first consider the term $(g')^\top R g' - g^\top R g$. Note that

$$\begin{aligned} \|(g')^\top R g' - g^\top R g\| &= \|(g' - g)^\top g' + g^\top (g' - g)\| \\ &\leq (c_2 + 1)\|g' - g\| + c_2\|g' - g\| \\ &= (2c_2 + 1)\|g' - g\|. \end{aligned}$$

Next, we consider $\gamma[(g')^\top B^\top P_{K'} B g' - g^\top B^\top P_K B g]$. Observe that

$$\begin{aligned} &((g')^\top B^\top P_{K'} B g' - g^\top B^\top P_K B g) \\ &= (g' - g)^\top B^\top P_{K'} B g' + g^\top B^\top (P_{K'} B g' - P_K B g) \\ &= (g' - g)^\top B^\top P_{K'} B g' + g^\top B^\top ((P_{K'} - P_K) B g' + P_K B (g' - g)). \end{aligned}$$

Then,

$$\begin{aligned} &\gamma\|(g')^\top B^\top P_{K'} B g' - g^\top B^\top P_K B g\| \\ &\leq \gamma\|B\|_2(C_0 + c_8 + c_9 + c_{10})\|B\|(c_2 + 1)\|g' - g\| \\ &\quad + \gamma(\|B\|_2^2 c_2^2 \|P_{K'} - P_K\| + \|B\|_2^2 c_2 C_0 \|g' - g\|) \\ &\leq \gamma\|B\|_2(C_0 + c_8 + c_9 + c_{10})\|B\|(c_2 + 1)\|g' - g\| + \gamma\|B\|_2^2 c_2^2 (c_8 + c_9 + c_{10})\|K' - K\| + \gamma\|B\|_2^2 c_2 C_0 \|g' - g\| \\ &= \gamma(\|B\|_2^2(C_0 + c_8 + c_9 + c_{10})(c_2 + 1) + \|B\|_2^2 c_2 C_0)\|g' - g\| + \gamma\|B\|_2^2 c_2^2 (c_8 + c_9 + c_{10})\|K' - K\|. \end{aligned}$$

We finally turn to $2\gamma((g')^\top B^\top q_{K',g'}) - 2\gamma(g^\top B^\top q_{K,g})$. We have

$$\begin{aligned} &\|(g')^\top B^\top q_{K',g'} - g^\top B^\top q_{K,g}\| \\ &= \|(g' - g)^\top B^\top q_{K',g'} + g^\top B^\top (q_{K',g'} - q_{K,g})\| \\ &\leq \|g' - g\|\|B\|_2(\|q_{K,g}\| + \|q_{K',g'} - q_{K,g}\|) + \|g\|\|B\|_2\|q_{K',g'} - q_{K,g}\| \\ &\leq \|g' - g\|\|B\|_2(c_{12} + c_{15} + c_{16}) + \|B\|_2 c_2 (c_{15}\|K' - K\| + c_{16}\|g' - g\|) \\ &= (\|B\|_2(c_{12} + c_{15} + c_{16}) + \|B\|_2 c_2 c_{16})\|g' - g\| + \|B\|_2 c_2 (c_{15}\|K' - K\|). \end{aligned}$$

Combining everything, we find that

$$\begin{aligned}
& |r_{K',g'} - r_{K,g}| \\
& \leq \frac{1}{1-\gamma} [\|(g')^\top Rg' - g^\top Rg\| + \gamma\|((g')^\top B^\top P_{K'}Bg' - g^\top B^\top P_KBg)\| + 2\gamma\|(g')^\top B^\top q_{K',g'} - g^\top B^\top q_{K,g}\|] \\
& \leq \frac{1}{1-\gamma} [2c_2 + 1 + \gamma(\|B\|_2^2(C_0 + c_8 + c_9 + c_{10})(c_2 + 1) + \|B\|_2^2c_2C_0)] \|g' - g\| \\
& \quad + \frac{1}{1-\gamma} [2\gamma(\|B\|_2(c_{12} + c_{15} + c_{16}) + \|B\|_2c_2c_{16})] \|g' - g\| \\
& \quad + \frac{1}{1-\gamma} [\gamma\|B\|_2^2c_2^2(c_8 + c_9 + c_{10}) + 2\gamma\|B\|_2c_2(c_{15})] \|K' - K\|
\end{aligned}$$

□

We are now ready to state the local Lipschitz property of $J(K, g)$ and its sample cost variant $J(K, g; x_0)$.

Proposition 7 (Local Lipschitz of sample and expected cost of $J(K, g)$). *Suppose $\sqrt{\gamma}\rho(A + BK) < 1$, and $J(K, g) < C_0$. Suppose also that $\|K' - K\| \leq c_4$, $\|g' - g\| \leq 1$. Then,*

$$J(K', g') - J(K, g) \leq c_{21}\|(K' - K, g' - g)\|,$$

where

$$c_{21} = \frac{\max\{\frac{n}{C_n}c_{11} + c_{19}, c_{20}\}}{\sqrt{2}}.$$

Similarly, for the initial state x_0 of any sample,

$$J(K', g'; x_0) - J(K, g; x_0) \leq c_{22}\|(K' - K, g' - g)\|,$$

where

$$c_{22} = \frac{\max\{C_n(c_8 + c_9 + c_{10}) + \sqrt{C_n}c_{15} + c_{19}, c_{16} + c_{20}\}}{\sqrt{2}}.$$

Proof. We first show the local Lipschitz property of the expected cost. Recall from Lemma 14 that $|C(K') - C(K)| \leq \frac{n}{C_n}c_{11}\|K' - K\|$. Then, using the assumption that $\mathbb{E}[x_0] = 0$ and $\mathbb{E}[x_0] = I$, we find that

$$\begin{aligned}
|J(K', g') - J(K, g)| &= |C_{K'} + r_{K',g'} - (C_K + r_{K,g})| \\
&\leq |C(K') - C(K)| + |r_{K',g'} - r_{K,g}| \\
&\leq \frac{n}{C_n}c_{11}\|K' - K\| + c_{19}\|K' - K\| + c_{20}\|g' - g\| \\
&= \left(\frac{n}{C_n}c_{11} + c_{19}\right)\|K' - K\| + c_{20}\|g' - g\|
\end{aligned}$$

where we used Lemma 17 for the second last line. Then the result for the expected cost holds by the observation that $(a+b)^2 \leq 2(a^2+b^2)$ for any $0 \leq a, b$.

We next examine the sample cost for a starting state x_0 , which we can write as $V_{K,g}(x_0)$, using notation from Lemma 3. We have

$$\begin{aligned}
& |V_{K',g'}(x_0) - V_{K,g}(x_0)| \\
&= |x_0^\top (P_{K'} - P_K)x_0 + 2x_0^\top (q_{K',g'} - q_{K,g}) + r_{K',g'} - r_{K,g}| \\
&\leq \|P_{K'} - P_K\|_2 \|x_0\|^2 + 2\|x_0\|^\top \|q_{K',g'} - q_{K,g}\| + |r_{K',g'} - r_{K,g}| \\
&\leq C_n(c_8 + c_9 + c_{10})\|K' - K\| + 2\sqrt{C_n}(c_{15}\|K' - K\| + c_{16}\|g' - g\|) + c_{19}\|K' - K\| + c_{20}\|g' - g\| \\
&= (C_n(c_8 + c_9 + c_{10}) + \sqrt{C_n}c_{15} + c_{19})\|K' - K\| + (c_{16} + c_{20})\|g' - g\|.
\end{aligned}$$

□

We are now ready to show that $J(K, g)$ is also locally Lipschitz smooth, *i.e.* its gradient is Lipschitz. We first recall from Lemma 4 that the gradients of J with respect to K and g are the following:

$$\begin{aligned}
\nabla_K J &= 2C_K \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t x_t^\top \right] + 2d_{K,g} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t^\top \right], \\
\nabla_g J &= 2C_K \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t \right] + 2d_{K,g} \sum_{t=0}^{\infty} \gamma^t,
\end{aligned}$$

where

$$C_K = RK + \gamma B^\top P_K (BK + A), \quad d_{K,g} = Rg + \gamma B^\top P_K Bg + \gamma B^\top q_{K,g}.$$

We first prove the following lemma, which will simplify the analysis of $\mathbb{E}[\sum_{t=0}^{\infty} \gamma^t x_t x_t^\top]$. To this end we first define a sequence $\{y_t\}$, $y_0 = x_0$, that is obtained by following the linear policy $u_t = Kx_t$ for the system $(\sqrt{\gamma}A, \sqrt{\gamma}B)$ such that $y_{t+1} = \sqrt{\gamma}A + \sqrt{\gamma}BK y_t$. Then, define

$$\Sigma_K := \mathbb{E} \left[\sum_{t=0}^{\infty} y_t y_t^\top \right].$$

Then,

Lemma 18. *Recall the sequence $\{y_t\}$ and the corresponding state sum-of-covariance matrix Σ_K that we defined before this lemma. We have*

$$\mathbb{E} \left[\sum_t \gamma^t x_t x_t^\top \right] = \Sigma_K + (S_K^\top Bg)(Bg)^\top S_K.$$

Proof. Observe that

$$\mathbb{E}[x_t x_t^\top] = \mathbb{E}[(x_t - \mathbb{E}[x_t])(x_t - \mathbb{E}[x_t])^\top] + \mathbb{E}[x_t] \mathbb{E}[x_t]^\top.$$

Recall that $x_{t+1} = (A + BK)x_t + Bg$, so $\mathbb{E}[x_{t+1}] = (A + BK)^t \mathbb{E}[x_t] + Bg$. A direct calculation shows that (since $\mathbb{E}[x_0] = 0$)

$$\begin{aligned}\mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t x_t\right] &= \sum_{t=0}^{\infty} \gamma^t (A + BK)^t Bg \\ &= S_K^\top Bg.\end{aligned}$$

Next, via similar algebraic calculation, we find that

$$\mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t (x_t - \mathbb{E}[x_t])(x_t - \mathbb{E}[x_t])^\top\right] = \mathbb{E}\left[\sum_{t=0}^{\infty} y_t y_t^\top\right],$$

since for each summand the terms involving g cancel out. \square

Next we show that $\Sigma_{K'} - \Sigma_K$ can also be bounded.

Lemma 19. *Suppose $\sqrt{\gamma}\rho(A + BK) < 1$, and $J(K, g) < C_0$. Suppose also that $\|K' - K\| \leq c_4$. Then,*

$$\|\Sigma_{K'} - \Sigma_K\| \leq c_{23}\|K' - K\|,$$

where $c_{23} = 8\left(\frac{C_0}{\sigma_{\min}(Q)}\right)^2 \sqrt{\gamma}\|B\|_2$. In addition, due to the assumption $\|K' - K\| \leq c_4$, we have

$$\|\Sigma_{K'} - \Sigma_K\| \leq 2c_3.$$

Proof. This follows from Lemma 16 of [7]. \square

Next, we show that $C_{K,g}$ and $d_{K,g}$ are both bounded. Recall that $E_{K,g} = [C_{K,g} \quad d_{K,g}]$. From the lower bound in Lemma 5, we see that (using our simplifying assumption that $\mathbb{E}[x_0 x_0^\top] = I$)

$$\begin{aligned}\text{tr}(E_{K,g}^\top E_{K,g}) &\leq (J(K, g) - J(K^*, g^*))\|R + B^\top P_K B\|_2 \\ \implies \|C_{K,g}\|^2 + \|d_{K,g}\|^2 &\leq C_0(\|R\|_2 + \|B\|_2^2 C_0) := c_{24}^2.\end{aligned}$$

We next bound the quantities $C_{K'} - C_K$ and $d_{K',g'} - d_{K,g}$.

Lemma 20. *Suppose $\sqrt{\gamma}\rho(A + BK) < 1$, and $J(K, g) < C_0$. Suppose also that $\|K' - K\| \leq c_4$, $\|g' - g\| \leq 1$. Then,*

$$\begin{aligned}\|C_{K'} - C_K\| &\leq c_{25}\|K' - K\| \\ \|d_{K',g'} - d_{K,g}\| &\leq c_{26}\|K' - K\| + c_{27}\|g' - g\|,\end{aligned}$$

where

$$\begin{aligned}c_{25} &= \|R\|_2 + \gamma\|B\|_2[(c_8 + c_9 + c_{10}) + \|B\|_2 C_0] \\ c_{26} &= \gamma\|B\|_2^2(c_2 + 1)(c_8 + c_9 + c_{10}) + \gamma\|B\|_2 c_{15} \\ c_{27} &= \|R\|_2 + \gamma\|B\|_2^2 C_0 + \gamma\|B\|_2 c_{16}.\end{aligned}$$

Proof. We first handle $C_{K'} - C_K$. We have

$$\begin{aligned}
& \|C_{K'} - C_K\| \\
& \leq \|RK' + \gamma B^\top P_{K'}(BK' + A) - (RK + \gamma B^\top P_K(BK + A))\| \\
& \leq \|R\|_2 \|K' - K\| + \gamma \|B\|_2 \|(P_{K'} - P_K)(A + BK') + P_K B(K' - K)\| \\
& \leq \|R\|_2 \|K' - K\| + \gamma \|B\|_2 [(c_8 + c_9 + c_{10}) \|K' - K\| + \|B\|_2 C_0 \|K' - K\|] \\
& = (\|R\|_2 + \gamma \|B\|_2 [(c_8 + c_9 + c_{10}) + \|B\|_2 C_0]) \|K' - K\|.
\end{aligned}$$

Meanwhile,

$$\begin{aligned}
& \|d_{K',g'} - d_{K,g}\| \\
& = \|Rg' + \gamma B^\top P_{K'} Bg' + \gamma B^\top q_{K',g'} - (Rg + \gamma B^\top P_K Bg + \gamma B^\top q_{K,g})\| \\
& \leq \|R\|_2 \|g' - g\| + \gamma \|B\|_2 \|(P_{K'} - P_K)Bg' + P_K B(g' - g)\| + \gamma \|B\|_2 \|q_{K',g'} - q_{K,g}\| \\
& \leq \|R\|_2 \|g' - g\| + \gamma \|B\|_2 (\|B\|_2 (c_2 + 1)(c_8 + c_9 + c_{10}) \|K' - K\| + C_0 \|B\|_2 \|g' - g\|) \\
& \quad + \gamma \|B\|_2 (c_{15} \|K' - K\| + c_{16} \|g' - g\|) \\
& \leq (\|R\|_2 + \gamma \|B\|_2^2 C_0 + \gamma \|B\|_2 c_{16}) \|g' - g\| \\
& \quad + (\gamma \|B\|_2^2 (c_2 + 1)(c_8 + c_9 + c_{10}) + \gamma \|B\|_2 c_{15}) \|K' - K\|
\end{aligned}$$

□

We are finally ready to bound $\|\nabla J(K', g') - \nabla J(K, g)\|$.

Proposition 8 ($J(K, g)$ is locally smooth). *Suppose $\sqrt{\gamma}\rho(A + BK) < 1$, and $J(K, g) < C_0$. Suppose also that $\|K' - K\| \leq c_4$, $\|g' - g\| \leq 1$. Then,*

$$\begin{aligned}
& \|\nabla_K J(K', g') - \nabla_K J(K, g)\| \leq c_{28} \|K' - K\| + c_{29} \|K' - K\|, \\
& \|\nabla_g J(K', g') - \nabla_g J(K, g)\| \leq c_{30} \|K' - K\| + c_{31} \|K' - K\|.
\end{aligned}$$

Together, we find that

$$\begin{aligned}
\|\nabla J(K', g') - \nabla J(K, g)\| & \leq \|\nabla_K J(K', g') - \nabla_K J(K, g)\| + \|\nabla_g J(K', g') - \nabla_g J(K, g)\| \\
& \leq \frac{\max\{c_{28} + c_{30}, c_{29} + c_{31}\}}{\sqrt{2}} \|(K' - K, g' - g)\|,
\end{aligned}$$

where

$$\begin{aligned}
c_{28} & = 2c_{25} [\|B\|_2^2 (c_2 + 1)^2 (c_3 + 2\gamma c_3^2 \|B\|_2)^2 + (3c_3)] + 2c_{24} 2\gamma c_3^2 \|B\|_2 (\|B\|_2^2 (c_2 + 1)^2 (c_3 + 2\gamma c_3^2 \|B\|_2)) \\
& \quad + c_2 \|B\|_2 (2\gamma c_3^2 \|B\|_2 + 2c_{23}) \\
& \quad + 2c_{26} \|B\|_2 (c_2 + 1)(c_3 + 2\gamma c_3^2 \|B\|_2) + 2c_{24} c_2 \|B\|_2 2\gamma c_3^2 \|B\|_2 \\
c_{29} & = 2c_{24} c_3 \|B\|_2 (c_2 + 1) \|B\|_2 (c_3 + 2\gamma c_3^2 \|B\|_2) + 2c_2 \|B\|_2 (c_3 + 2\gamma c_3^2 \|B\|_2) \\
& \quad + 2c_{27} \|B\|_2 (c_2 + 1)(c_3 + 2\gamma c_3^2 \|B\|_2) + 2c_{24} \|B\|_2 (c_3 + 2\gamma c_3^2 \|B\|_2) \\
c_{30} & = 2c_{25} (c_3 + 2\gamma c_3^2 \|B\|_2) \|B\|_2 (c_2 + 1) + 2c_{24} (2\gamma c_3^2 \|B\|_2) \|B\|_2 (c_2 + 1) + 2c_{26} \frac{1}{1 - \gamma}, \\
c_{31} & = 2c_{24} c_3 \|B\|_2 + 2c_{27} \frac{1}{1 - \gamma}.
\end{aligned}$$

Proof. We first tackle $\|\nabla_K J(K', g') - \nabla_K J(K, g)\|$. Using Lemma 18, we have

$$\begin{aligned}\nabla_K J(K, g) &= 2C_K \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t x_t^\top \right] + 2d_{K,g} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t x_t^\top \right] \\ &= 2C_K (S_K^\top Bg(Bg)^\top S_K + \Sigma_K) + 2d_{K,g} (Bg)^\top S_K.\end{aligned}$$

Thus

$$\begin{aligned}&\|\nabla_K J(K', g') - \nabla_K J(K, g)\| \\ &\leq 2\|(C_{K'} - C_K)(S_{K'}^\top Bg'(Bg')^\top S_{K'} + \Sigma_{K'}) + C_K[S_{K'}^\top Bg'(Bg')^\top S_{K'} + \Sigma_{K'} - (S_K^\top Bg(Bg)^\top S_K + \Sigma_K)]\| \\ &\quad + 2\|(d_{K',g'} - d_{K,g})(Bg')^\top S_{K'} + d_{K,g}((Bg')^\top S_{K'} - (Bg)^\top S_K)\|\end{aligned}$$

We handle the first summand. Observe that

$$\begin{aligned}&2\|(C_{K'} - C_K)(S_{K'}^\top Bg'(Bg')^\top S_{K'} + \Sigma_{K'}) + C_K[S_{K'}^\top Bg'(Bg')^\top S_{K'} + \Sigma_{K'} - (S_K^\top Bg(Bg)^\top S_K + \Sigma_K)]\| \\ &\leq 2\|C_{K'} - C_K\| \|S_{K'}^\top Bg'(Bg')^\top S_{K'} + \Sigma_{K'}\| \\ &\quad + 2\|C_K\| \|(S_{K'} - S_K)(Bg'(g')^\top B^\top S_{K'}) + S_K(Bg'(g')^\top B^\top S_{K'} - Bg(g)^\top B^\top S_K) + \Sigma_{K'} - \Sigma_K\| \\ &\leq 2c_{25} \|K' - K\| [\|B\|_2^2 (c_2 + 1)^2 (c_3 + 2\gamma c_3^2 \|B\|_2)^2 + (3c_3)] \\ &\quad + 2c_{24} 2\gamma c_3^2 \|B\|_2 \|K' - K\| (\|B\|_2^2 (c_2 + 1)^2 (c_3 + 2\gamma c_3^2 \|B\|_2)) \\ &\quad + 2c_{24} c_3 \|B\|_2 \|(g' - g)(g')^\top B^\top S_{K'} + g((g')^\top B^\top S_{K'} - g^\top B^\top S_K)\| \\ &\quad + 2c_{23} \|K' - K\| \\ &\leq 2c_{25} \|K' - K\| [\|B\|_2^2 (c_2 + 1)^2 (c_3 + 2\gamma c_3^2 \|B\|_2)^2 + (3c_3)] \\ &\quad + 2c_{24} 2\gamma c_3^2 \|B\|_2 \|K' - K\| (\|B\|_2^2 (c_2 + 1)^2 (c_3 + 2\gamma c_3^2 \|B\|_2)) \\ &\quad + 2c_{24} c_3 (\|B\|_2 \|g' - g\| (c_2 + 1) \|B\|_2 (c_3 + 2\gamma c_3^2 \|B\|_2)) \\ &\quad + 2c_2 \|g' - g\| \|B\|_2 (c_3 + 2\gamma c_3^2 \|B\|_2) + c_2 \|B\|_2 (2\gamma c_3^2 \|B\|_2 \|K' - K\|) \\ &\quad + 2c_{23} \|K' - K\|\end{aligned}$$

For the next summand, observe that

$$\begin{aligned}&2\|(d_{K',g'} - d_{K,g})(Bg')^\top S_{K'} + d_{K,g}((Bg')^\top S_{K'} - (Bg)^\top S_K)\| \\ &\leq 2(c_{26} \|K' - K\| + c_{27} \|g' - g\|) \|B\|_2 (c_2 + 1) (c_3 + 2\gamma c_3^2 \|B\|_2) \\ &\quad + 2c_{24} \|g' - g\| \|B\|_2 (c_3 + 2\gamma c_3^2 \|B\|_2) + 2c_{24} c_2 \|B\|_2 2\gamma c_3^2 \|B\|_2 \|K' - K\|.\end{aligned}$$

Therefore,

$$\|\nabla_K J(K', g') - \nabla_K J(K, g)\| \leq c_{28} \|K' - K\| + c_{29} \|g' - g\|,$$

where

$$\begin{aligned}c_{28} &= 2c_{25} [\|B\|_2^2 (c_2 + 1)^2 (c_3 + 2\gamma c_3^2 \|B\|_2)^2 + (3c_3)] + 2c_{24} 2\gamma c_3^2 \|B\|_2 (\|B\|_2^2 (c_2 + 1)^2 (c_3 + 2\gamma c_3^2 \|B\|_2)) \\ &\quad + c_2 \|B\|_2 (2\gamma c_3^2 \|B\|_2) + 2c_{23} \\ &\quad + 2c_{26} \|B\|_2 (c_2 + 1) (c_3 + 2\gamma c_3^2 \|B\|_2) + 2c_{24} c_2 \|B\|_2 2\gamma c_3^2 \|B\|_2 \\ c_{29} &= 2c_{24} c_3 \|B\|_2 (c_2 + 1) \|B\|_2 (c_3 + 2\gamma c_3^2 \|B\|_2) + 2c_2 \|B\|_2 (c_3 + 2\gamma c_3^2 \|B\|_2) \\ &\quad + 2c_{27} \|B\|_2 (c_2 + 1) (c_3 + 2\gamma c_3^2 \|B\|_2) + 2c_{24} \|B\|_2 (c_3 + 2\gamma c_3^2 \|B\|_2).\end{aligned}$$

Next, observe that

$$\begin{aligned}
& \|\nabla_g J(K'g') - \nabla_g J(K, g)\| \\
&= \|2C_{K'} S_{K'}^\top Bg' + 2d_{K',g'} \sum_t \gamma^t - 2C_K S_K^\top Bg + 2d_{K,g} \sum_t \gamma^t\| \\
&\leq 2\|C_{K'} - C_K\| \|S_{K'}\|_2 \|B\|_2 (c_2 + 1) + 2\|C_K\| \|(S_{K'}^\top - S_K^\top)Bg' + S_K^\top B(g' - g)\| \\
&\quad + 2\|d_{K',g'} - d_{K,g}\| \frac{1}{1-\gamma} \\
&\leq 2c_{25} \|K' - K\| (c_3 + 2\gamma c_3^2 \|B\|_2) \|B\|_2 (c_2 + 1) + 2c_{24} (2\gamma c_3^2 \|B\|_2) \|K' - K\| \|B\|_2 (c_2 + 1) + 2c_{24} c_3 \|B\|_2 \|g' - g\| \\
&\quad + 2(c_{26} \|K' - K\| + c_{27} \|g' - g\|) \frac{1}{1-\gamma} \\
&= \left(2c_{25} (c_3 + 2\gamma c_3^2 \|B\|_2) \|B\|_2 (c_2 + 1) + 2c_{24} (2\gamma c_3^2 \|B\|_2) \|B\|_2 (c_2 + 1) + 2c_{26} \frac{1}{1-\gamma} \right) \|K' - K\| \\
&\quad + \left(2c_{24} c_3 \|B\|_2 + 2c_{27} \frac{1}{1-\gamma} \right) \|g' - g\|
\end{aligned}$$

□

Proof of Proposition 2. The local Lipschitz and smoothness results in Proposition 2 are a consequence of Proposition 7 and Proposition 2. □

Constants used in local Lipschitz and smoothness arguments. Below is a list of constants used in the local Lipschitz and smoothness argument, where $C_0 > 0$ is an upper bound on $J(K, g)$.

1. $c_1 = \frac{\sqrt{(C_0 - C(K^*))(\rho(R) + \|B\|^2 C_0)}}{\sigma_{\min}(R)} + \frac{\|B\| C_0 \|A\|}{\sigma_{\min}(R)}$
2. $c_2 = \sqrt{\frac{C_0}{\sigma_{\min}(R)}}$
3. $c_3 = \frac{2C_0}{\sigma_{\min}(Q)}$
4. $c_4 = \min \left\{ \frac{\sigma_{\min}(Q)}{4C_0 \|\sqrt{\gamma} B\| (\|\sqrt{\gamma} A + \sqrt{\gamma} B K\| + 1)}, 1 \right\}$
5. $c_5 = \frac{\sigma_{\min}(Q)}{4\gamma \|B\| C_0}$
6. $c_6 = \frac{1}{\sigma_{\min}(R)} \left(\sqrt{(\|R\|_2 + \|B\|_2^2 C_0)(C_0 - C^*)} + \|B\|_2 \|A\|_2 C_0 \right)$
7. $c_7 = \max \left\{ \frac{C_0}{\sigma_{\min}(Q)} \sqrt{(\|R\|_2 + \|B\|_2^2 C_0)(C_0 - C^*)}, c_6 \right\}$
8. $c_8 = 4 \left(\frac{C_0}{\sigma_{\min}(Q)} \right)^2 \|Q\|_2 \|B\|_2 (\|A\|_2 + \|B\|_2 c_7 + 1)$
9. $c_9 = 8 \left(\frac{C_0}{\sigma_{\min}(Q)} \right)^2 (c_7^2) \|R\|_2 \|B\|_2 (\|A\|_2 + \|B\|_2 c_7 + 1)$

10. $c_{10} = 2 \left(\frac{C_0}{\sigma_{\min}(Q)} \right)^2 (c_7 + 1) \|R\|_2$
11. $c_{11} = C_n(c_8 + c_9 + c_{10})$
12. $c_{12} = c_3 (c_1 c_2 \|R\| + 2\sqrt{\gamma} C_0 \|B\|_2 c_2)$
13. $c_{13} = 2\gamma c_3 \|B\|_2 (c_1 + 1) \|R\|_2 (c_2 + 1) + (c_2 + 1) \|R\|_2$
14. $c_{14} = 2\gamma^{3/2} (C_0 + c_8 + c_9 + c_{10}) \|B\|_2^2 (c_2 + 1) + c_3 (\gamma \|B\|_2^2 (C_0 + c_8 + c_9 + c_{10}) \|B\|_2 (c_2 + 1) + \sqrt{\gamma} \|B\|_2 (c_2 + 1) (c_8 + c_9 + c_{10}))$
15. $c_{15} = 2\gamma c_3 \|B\|_2 \|Qx^*\| + c_{13} + c_{14}$
16. $c_{16} = 2c_1 c_3 \|R\|_2 + c_3 C_0 \|B\|_2$
17. $c_{17} = \frac{1}{1-\gamma} [2c_2 + 1 + \gamma (\|B\|_2^2 (C_0 + c_8 + c_9 + c_{10}) (c_2 + 1) + \|B\|_2^2 c_2 C_0)]$
18. $c_{18} = \frac{1}{1-\gamma} [2\gamma (\|B\|_2 (c_{12} + c_{15} + c_{16}) + \|B\|_2 c_2 c_{16})]$
19. $c_{19} = c_{17} + c_{18}$
20. $c_{20} = \frac{1}{1-\gamma} [\gamma \|B\|_2^2 c_2^2 (c_8 + c_9 + c_{10}) + 2\gamma \|B\|_2 c_2 (c_{15})]$
21. $c_{21} = \frac{\max\{\frac{\gamma}{C_n} c_{11} + c_{19}, c_{20}\}}{\sqrt{2}}$
22. $c_{22} = \frac{\max\{C_n (c_8 + c_9 + c_{10}) + \sqrt{C_n} c_{15} + c_{19}, c_{16} + c_{20}\}}{\sqrt{2}}$
23. $c_{23} = 8 \left(\frac{C_0}{\sigma_{\min}(Q)} \right)^2 \sqrt{\gamma} \|B\|_2$
24. $c_{24} = \sqrt{C_0 (\|R\|_2 + \|B\|_2^2 C_0)}$
25. $c_{25} = \|R\|_2 + \gamma \|B\|_2 [(c_8 + c_9 + c_{10}) + \|B\|_2 C_0]$
26. $c_{26} = \gamma \|B\|_2^2 (c_2 + 1) (c_8 + c_9 + c_{10}) + \gamma \|B\|_2 c_{15}$
27. $c_{27} = \|R\|_2 + \gamma \|B\|_2^2 C_0 + \gamma \|B\|_2 c_{16}$
28. $c_{28} = 2c_{25} [\|B\|_2^2 (c_2 + 1)^2 (c_3 + 2\gamma c_3^2 \|B\|_2)^2 + (3c_3)] + 2c_{24} 2\gamma c_3^2 \|B\|_2 (\|B\|_2^2 (c_2 + 1)^2 (c_3 + 2\gamma c_3^2 \|B\|_2) + c_2 \|B\|_2 (2\gamma c_3^2 \|B\|_2 + 2c_{23} + 2c_{26} \|B\|_2 (c_2 + 1) (c_3 + 2\gamma c_3^2 \|B\|_2) + 2c_{24} c_2 \|B\|_2 2\gamma c_3^2 \|B\|_2)$
29. $c_{29} = 2c_{24} c_3 \|B\|_2 (c_2 + 1) \|B\|_2 (c_3 + 2\gamma c_3^2 \|B\|_2) + 2c_2 \|B\|_2 (c_3 + 2\gamma c_3^2 \|B\|_2) + 2c_{27} \|B\|_2 (c_2 + 1) (c_3 + 2\gamma c_3^2 \|B\|_2) + 2c_{24} \|B\|_2 (c_3 + 2\gamma c_3^2 \|B\|_2)$
30. $c_{30} = 2c_{25} (c_3 + 2\gamma c_3^2 \|B\|_2) \|B\|_2 (c_2 + 1) + 2c_{24} (2\gamma c_3^2 \|B\|_2) \|B\|_2 (c_2 + 1) + 2c_{26} \frac{1}{1-\gamma}$
31. $c_{31} = 2c_{24} c_3 \|B\|_2 + 2c_{27} \frac{1}{1-\gamma}$

Appendix C Algorithms

We restate here Algorithm 1 from the main paper, as well as Algorithm 2 (benchmark algorithm) which we omitted. In particular, we consider the more general setting where one-point or two-point estimator can be used.

Algorithm 2: Federated LQR Learning algorithm

- 1 Given: iteration number $T \geq 1$, communication interval $H \in \mathbb{N}$, initial stable K_0 , initial g_0^i for each $i \in [m]$, step size $\eta > 0$, and smoothing radius $r > 0$
- 2 **for** *epoch* $e \in \{0, \dots, \frac{T}{H} - 1\}$ **do**
- 3 Set $K_{eH}^i \leftarrow K_{eH}$ for each $i \in [m]$
- 4 **for** *iteration* $t = eH, eH + 1, \dots, (e + 1)H - 1$ **do**
- 5 **for** (*simultaneously*) *each agent* $i \in [m]$ **do**
- 6 Sample $(x_0)_t^i \sim \mathcal{D}$, $u_t^i \sim \text{Unif}(\mathbb{S}^{n \cdot k + k - 1})$
- 7 Set the zeroth-order estimator z_t^i as follows :

$$(z_t^i) \leftarrow \begin{cases} z_r^1((K_t^i, g_t^i), u_t^i, (x_0)_t^i) & \text{if in one-point setting,} \\ z_r^2((K_t^i, g_t^i), u_t^i, (x_0)_t^i) & \text{if in two-point setting.} \end{cases}$$

- 8 **end**
 - 9 **end**
 - 10 Set $K_{(e+1)H} \leftarrow \frac{1}{m} \sum_{i=1}^m K_{(e+1)H}^i$ (model averaging of K at end of epoch e)
 - 11 **end**
 - 12 **return** $K_T, g_T^i \forall i \in [m]$
-

Remark 1. The one-point and two-point zeroth order estimators are defined as follows for the LQR problem with random initialization:

$$z_r^1((K_t^i, g_t^i), u_t^i, (x_0)_t^i) := \frac{d}{r} \hat{J}^i((K_t^i, g_t^i) + ru_t^i; (x_0)_t^i) u_t^i,$$

$$z_r^2((K_t^i, g_t^i), u_t^i, (x_0)_t^i) := \frac{d}{2r} (J^i((K_t^i, g_t^i) + ru_t^i; (x_0)_t^i) - J^i((K_t^i, g_t^i) - ru_t^i; (x_0)_t^i)) u_t^i,$$

where $d = n \cdot k + k$ (i.e. dimension of K plus dimension of g). In addition, we note that $\hat{J}^i(K, g) := \min\{J^i(K, g), 10J_0^i(K_0, g_0^i)\}$ (where $J_0^i := J^i(K_0, g_0^i)$), which incurs a bias in the gradient estimation procedure but simplifies the maximum bound on z_r^1 ⁴.

The next algorithm, Algorithm 3, is the benchmark algorithm which we will compare our federated algorithm against. We note that it basically runs m

⁴This truncation technique was also used in [14]

independent zeroth-order gradient descent independently.

Algorithm 3: Benchmark distributed LQR algorithm

- 1 Given: iteration number $T \geq 1$, communication interval $H \in \mathbb{N}$, initial stable K_0 , initial g_0^i for each $i \in [m]$, step size $\eta > 0$, and smoothing radius $r > 0$
- 2 **for** iteration $t = 0, 1, \dots, T - 1$ **do**
- 3 **for** (simultaneously) each agent $i \in [m]$ **do**
- 4 Sample $(x_0)_t^i \sim \mathcal{D}$, $u_t^i \sim \text{Unif}(\mathbb{S}^{n \cdot k - 1})$
- 5 Set the zeroth-order estimator z_t^i as follows:

$$(z_t^i) \leftarrow \begin{cases} z_r^1((K_t^i, g_t^i), u_t^i, (x_0)_t^i) & \text{if in one-point setting,} \\ z_r^2((K_t^i, g_t^i), u_t^i, (x_0)_t^i) & \text{if in two-point setting.} \end{cases}$$

- 6 Update $K_{t+1}^i \leftarrow K_t^i - \eta(z_t^i)_K$, $t \leftarrow t + 1$
 - 7 Update $g_{t+1}^i \leftarrow g_t^i - \eta(z_t^i)_g$
 - 8 **end**
 - 9 **end**
 - 10 **return** $K_T^i, g_T^i \forall i \in [m]$
-

Appendix D Analysis of main results

We first state the following terms which will be used in the proofs to follow. Consider the terms

$$Z_\infty^i := \max\{\|z_t^i(\cdot)\|\}, \quad Z_\infty := \max_{i \in [m]} Z_\infty^i$$

$$Z_{2,K} := \max_t \frac{1}{m} \sum_{i=1}^m \mathbb{E}^t \|(z_t^i)_K\|^2, \quad Z_{2,G} = \max_t \frac{1}{m} \sum_{i=1}^m \mathbb{E}^t \|(z_t^i)_{g^i}\|^2,$$

which collectively form bounds on the maximum size and variance of the zeroth-order gradient update. We next define the stability region \mathcal{G}_0 and \mathcal{G}'_0 which we show in the appendix that the iterates of Algorithm 2 and the benchmark algorithm respectively stay within. Let

$$\mathcal{G}_0 := \{(K, G) : \bar{J}(K, G) \leq 10\bar{J}(K_0, G_0)\},$$

$$\mathcal{G}'_0 := \left\{ \{K^i, g^i\}_{i \in [m]} : \frac{1}{m} \sum_{i=1}^m J^i(K^i, g^i) \leq 10 \frac{1}{m} \sum_{i=1}^m J^i(K_0^i, g_0^i) \right\}.$$

Next, by Proposition 2, note that there exists a local radius $\rho > 0$, local Lipschitz parameter $\lambda > 0$ and local smoothness parameter $L > 0$ such that the following holds.

1. If $(K, G) \in \mathcal{G}_0$, for any $i \in [m]$, and $\|(K', (g^i)') - (K, g^i)\| \leq \rho$, then

$$\begin{aligned} \|\nabla J^i(K', (g^i)') - \nabla J^i(K, g^i)\| &\leq L\|(K', (g^i)') - (K, g^i)\|, \\ \|J^i(K', (g^i)') - J^i(K, g^i)\| &\leq \lambda\|(K', (g^i)') - (K, g^i)\|. \end{aligned}$$
2. If $(\{K_i, g^i\}_{i=1}^m) \in \mathcal{G}'_0$, for any $i \in [m]$, and $\|((K^i)', (g^i)') - (K^i, g^i)\| \leq \rho$,

$$\|\nabla J^i((K^i)', (g^i)') - \nabla J^i(K^i, g^i)\| \leq L\|((K^i)', (g^i)') - (K^i, g^i)\|.$$
3. If $(K, G) \in \mathcal{G}_0$, and $\|(K', G') - (K, G)\| \leq \rho$, then

$$\|\nabla \bar{J}(K', G') - \nabla \bar{J}(K, G)\| \leq L\|(K', G') - (K, G)\|.$$

D.1 Theorem 1

In this subsection, we provide analysis and proof of convergence for Theorem 1, which provides analysis for Algorithm 2.

At a high level, proving Theorem 1 requires us to show the following two things.

1. Showing an expected decrease in cumulative cost every epoch — we will do so by designing a suitable virtual sequence that matches the true iterates at the start of each epoch, and show a per-iteration decrease in cost for this virtual sequence.
2. Showing that the iterates remain stable throughout the algorithm with large probability — we will do so by a martingale argument.

Let $\mathcal{I}_H = \{H, 2H, \dots, T\}$ denote the iterations where model averaging over K is performed. Observe then that the iterates of K evolve in the following way.

$$K_{t+1}^i = \begin{cases} K_t^i - \eta(z_t^i)_K & \text{if } t \notin \mathcal{I}_H \\ \frac{1}{m} \sum_{i=1}^m K_t^i & \text{if } t \in \mathcal{I}_H. \end{cases}$$

To facilitate analysis, consider the following virtual sequence

$$\bar{K}_t = \frac{1}{m} \sum_{i=1}^m K_t^i, \quad \bar{K}_0 = K_0.$$

We see that the virtual sequence of K and the true sequence of G evolve as follows:

$$(\bar{K}_{t+1}, G_{t+1}) = (\bar{K}_t, G_t) - \eta \frac{1}{m} \sum_{i=1}^m \nabla J^i(K_t^i, g_t^i).$$

We will work extensively with this joint sequence. We first define a stability region of the algorithm, which we will show the joint sequence (\bar{K}_t, G_t) stays in throughout the algorithm. Let

$$\mathcal{G}_0 := \{(K, G) : \bar{J}(K, G) \leq 10\bar{J}(K_0, G_0)\}.$$

Then, by Proposition 2, there exists $\rho > 0$ such that the following holds.

1. If $(K, G) \in \mathcal{G}_0$, for any $i \in [m]$, if $\|(K', g^i) - (K, g^i)\| \leq \rho$, then $\|\nabla J^i(K', g^i) - \nabla J^i(K, g^i)\| \leq L\|(K', g^i) - (K, g^i)\|$
2. If $(K, G) \in \mathcal{G}_0$, if $\|(K', G') - (K, G)\| \leq \rho$, then $\|\nabla \bar{J}(K', G') - \nabla \bar{J}(K, G)\| \leq L\|(K', G') - (K, G)\|$.

We first show a per-iteration change in optimality gap for the joint sequence, which will be essential in the convergence proofs. For convenience, define the optimality gap

$$\Delta_t = \bar{J}(\bar{K}_t, G_t) - \bar{J}^*.$$

Lemma 21 (Per-iteration change in optimality gap). *Suppose $(\bar{K}_t, G_t) \in \mathcal{G}^0$. Then, if we choose step-size $\eta > 0$, smoothing radius $r > 0$ and communication interval $H \geq 1$ such that*

$$\eta \leq \min \left\{ \frac{\rho}{2HZ_\infty}, \frac{L}{8} \right\}, \quad r \leq \min \left\{ \sqrt{\frac{\epsilon}{720L\mu}}, \rho \right\},$$

the optimality gap satisfies the following bound

$$\mathbb{E}^t[\Delta_{t+1}] \leq \left(1 - \frac{\eta}{4\mu}\right)\Delta_t + \frac{\eta^2 L}{2} \left(\frac{Z_{2,K} + Z_{2,G}}{m} + 6\eta LH^2 Z_{\infty,K}^2 \right) + \frac{\eta}{120\mu}\epsilon.$$

Proof. By choosing $\eta \leq \frac{\rho}{2Z_\infty}$, we find that

$$\begin{aligned} \|(\bar{K}_{t+1}, G_{t+1}) - (\bar{K}_t, G_t)\|^2 &= \left\| \bar{K}_t - \eta \frac{1}{m} \sum_{i=1}^m (z_t^i)_K(K_t^i, g_t^i) - \bar{K}_t \right\|^2 + \sum_{i=1}^m \left\| g_t^i - \frac{\eta}{m} (z_t^i)_{g^i}(K_t^i, g_t^i) - g_t^i \right\|^2 \\ &\leq \eta^2 (2Z_\infty^2) \leq \rho^2. \end{aligned}$$

Since $(\bar{K}_t, G_t) \in \mathcal{G}_0$, by local smoothness, and taking conditional expectation on \mathcal{F}_t (which we denote by \mathbb{E}^t), it follows that

$$\begin{aligned} \mathbb{E}^t \bar{J}(\bar{K}_{t+1}, G_{t+1}) &\leq \bar{J}(\bar{K}_t, G_t) - \eta \mathbb{E}^t \left\langle \nabla_K \bar{J}(\bar{K}_t, G_t), \frac{1}{m} \sum_{i=1}^m (z_t^i)_K(K_t^i, g_t^i) \right\rangle + \frac{L}{2} \mathbb{E}^t \left\| \eta \frac{1}{m} \sum_{i=1}^m z_t^i(K_t^i, g_t^i) \right\|^2 \\ &\quad - \frac{\eta}{m^2} \sum_{i=1}^m \mathbb{E}^t \langle \nabla_{g^i} J^i(\bar{K}_t, g_t^i), (z_t^i)_{g^i}(K_t^i, g_t^i) \rangle + \frac{L}{2} \frac{1}{m^2} \sum_{i=1}^m \mathbb{E}^t \left\| \eta (z_t^i)_{g^i}(K_t^i, g_t^i) \right\|^2 \\ &= \bar{J}(\bar{K}_t, G_t) - \eta \left\langle \nabla_K \bar{J}(\bar{K}_t, G_t), \frac{1}{m} \sum_{i=1}^m \nabla_K J_r^i(K_t^i, g_t^i) \right\rangle + \frac{L\eta^2}{2} \mathbb{E}^t \left\| \frac{1}{m} \sum_{i=1}^m (z_t^i)_K(K_t^i, g_t^i) \right\|^2 \\ &\quad - \frac{\eta}{m^2} \sum_{i=1}^m \langle \nabla_{g^i} J^i(\bar{K}_t, g_t^i), \nabla_{g^i} J_r^i(K_t^i, g_t^i) \rangle + \frac{L\eta^2}{2} \frac{1}{m^2} \sum_{i=1}^m \mathbb{E}^t \left\| (z_t^i)_{g^i}(K_t^i, g_t^i) \right\|^2 \end{aligned}$$

We first handle the term $-\eta \langle \nabla_K \bar{J}(\bar{K}_t, G_t), \frac{1}{m} \sum_{i=1}^m \nabla_K J_r^i(K_t^i, g_t^i) \rangle$. Observe that

$$\begin{aligned}
& -\eta \left\langle \nabla_K \bar{J}(\bar{K}_t, G_t), \frac{1}{m} \sum_{i=1}^m \nabla_K J_r^i(K_t^i, g_t^i) \right\rangle \\
&= -\eta \langle \nabla_K \bar{J}(\bar{K}_t, G_t), \nabla_K \bar{J}(\bar{K}_t, G_t) \rangle - \eta \left\langle \nabla_K \bar{J}(\bar{K}_t, G_t), \frac{1}{m} \sum_{i=1}^m \nabla_K J_r^i(K_t^i, g_t^i) - \nabla_K \bar{J}(\bar{K}_t, G_t) \right\rangle \\
&= -\eta \|\nabla_K \bar{J}(\bar{K}_t, G_t)\|^2 + \eta \left(\frac{\|\nabla_K \bar{J}(\bar{K}_t, G_t)\|^2}{2} + \frac{\|\frac{1}{m} \sum_{i=1}^m \nabla_K J_r^i(K_t^i, g_t^i) - \nabla_K \bar{J}(\bar{K}_t, G_t)\|^2}{2} \right) \\
&\leq -\frac{\eta}{2} \|\nabla_K \bar{J}(\bar{K}_t, G_t)\|^2 + \frac{\eta}{2m} \sum_{i=1}^m \|\nabla_K J_r^i(K_t^i, g_t^i) - \nabla_K J^i(\bar{K}_t, g_t^i)\|^2
\end{aligned}$$

Note that $eH \leq t < (e+1)H$ for some epoch e . Then,

$$\begin{aligned}
& \frac{\eta}{2m} \sum_{i=1}^m \|\nabla_K J_r^i(K_t^i, g_t^i) - \nabla_K J^i(\bar{K}_t, g_t^i)\|^2 \\
&\leq \frac{\eta}{2m} \sum_{i=1}^m 2\|\nabla_K J_r^i(K_t^i, g_t^i) - \nabla_K J_i(K_t^i, g_t^i)\|^2 + 2\|\nabla_K J_i(K_t^i, g_t^i) - \nabla_K J_i(\bar{K}_t, g_t^i)\|^2 \\
&= \frac{\eta}{m} \sum_{i=1}^m L^2 r^2 + L^2 \|K_t^i - \bar{K}_t\|^2 \\
&= L^2 \left(r^2 + \frac{1}{m} \sum_{i=1}^m \left\| K_{eH} - \eta \sum_{\tau=eH}^{t-1} (z_\tau^i)_K - \left(K_{eH} - \eta \sum_{\tau=eH}^{t-1} \frac{1}{m} \sum_{j=1}^m (z_\tau^j)_K \right) \right\|^2 \right) \\
&\leq 2L^2 \eta (r^2 + \eta^2 H^2 Z_\infty^2).
\end{aligned}$$

To move from the second to the third line, we used the (ρ, L) local Lipschitz condition for each ∇J^i , and our choice of $r \leq \rho$, $2\eta H Z_\infty \leq \rho$.

Next, we consider the term $-\eta \sum_{i=1}^m \langle \nabla_{g^i} J^i(\bar{K}_t, g_t^i), \nabla_{g^i} J_r^i(K_t^i, g_t^i) \rangle$. Observe that

$$\begin{aligned}
& -\frac{\eta}{m^2} \sum_{i=1}^m \langle \nabla_{g^i} J^i(\bar{K}_t, g_t^i), \nabla_{g^i} J_r^i(K_t^i, g_t^i) \rangle \\
&= -\frac{\eta}{m^2} \sum_{i=1}^m \frac{1}{2} \left[\|\nabla_{g^i} J^i(\bar{K}_t, g_t^i)\|^2 + \|\nabla_{g^i} J^i(\bar{K}_t, g_t^i) - \nabla_{g^i} J_r^i(K_t^i, g_t^i)\|^2 \right] \\
&\leq -\frac{\eta}{m^2} \sum_{i=1}^m \frac{1}{2} \|\nabla_{g^i} J^i(\bar{K}_t, g_t^i)\|^2 + \\
&\quad + \frac{\eta}{m^2} \left(\sum_{i=1}^m \|\nabla_{g^i} J^i(\bar{K}_t, g_t^i) - \nabla_{g^i} J^i(K_t^i, g_t^i)\|^2 + \|\nabla_{g^i} J^i(K_t^i, g_t^i) - \nabla_{g^i} J_r^i(K_t^i, g_t^i)\|^2 \right).
\end{aligned}$$

A similar analysis to the one for $\nabla_K(\cdot)$ finds that

$$\begin{aligned} & \frac{\eta}{m^2} \sum_{i=1}^m \left\| \nabla_{g^i} J^i(\bar{K}_t, g_t^i) - \nabla_{g^i} J^i(K_t^i, g_t^i) \right\|^2 + \left\| \nabla_{g^i} J^i(K_t^i, g_t^i) - \nabla_{g^i} J_r^i(K_t^i, g_t^i) \right\|^2 \\ & \leq 2L^2 \eta \left(\frac{r^2}{m} + \eta^2 H^2 \frac{Z_\infty^2}{m^2} \right). \end{aligned}$$

Thus,

$$\begin{aligned} & - \frac{\eta}{m^2} \sum_{i=1}^m \langle \nabla_{g^i} J^i(\bar{K}_t, g_t^i), \nabla_{g^i} J_r^i(K_t^i, g_t^i) \rangle \\ & \leq -\frac{1}{2} \left\| \nabla_G \bar{J}(\bar{K}_t, G_t) \right\|^2 + 2L^2 \eta \left(\frac{r^2}{m} + \eta^2 H^2 \frac{Z_\infty^2}{m^2} \right) \end{aligned}$$

We turn our attention now to $\mathbb{E}^t \left\| \frac{1}{m} \sum_{i=1}^m z_t^i(K_t^i, g_t^i) \right\|^2$. For notational convenience we use $\text{Var}(\cdot)$ to denote $\mathbb{E}[\|\cdot\|^2]$. Then,

$$\begin{aligned} & \mathbb{E}^t \left\| \frac{1}{m} \sum_{i=1}^m (z_t^i)_K(K_t^i, g_t^i) \right\|^2 \\ & = \frac{1}{m^2} \sum_{i=1}^m \text{Var}((z_t^i)_K) + \left\| \frac{1}{m} \sum_{i=1}^m \nabla_K J_r^i(K_t^i, g_t^i) \right\|^2 \\ & = \frac{1}{m^2} \sum_{i=1}^m \text{Var}((z_t^i)_K) + 2 \left\| \frac{1}{m} \sum_{i=1}^m \nabla_K J^i(K_t^i, g_t^i) \right\|^2 + 2 \left\| \frac{1}{m} \sum_{i=1}^m \nabla_K J_r^i(K_t^i, g_t^i) - \nabla_K J^i(K_t^i, g_t^i) \right\|^2 \\ & = \frac{1}{m^2} \sum_{i=1}^m \text{Var}((z_t^i)_K) + 4 \left\| \frac{1}{m} \sum_{i=1}^m \nabla_K J^i(\bar{K}_t, g_t^i) \right\|^2 \\ & \quad + 4 \left\| \frac{1}{m} \sum_{i=1}^m \nabla_K J^i(\bar{K}_t, g_t^i) - J^i(K_t^i, g_t^i) \right\|^2 + 2 \left\| \frac{1}{m} \sum_{i=1}^m \nabla_K J_r^i(K_t^i, g_t^i) - \nabla J^i(K_t^i, g_t^i) \right\|^2 \\ & \leq \frac{1}{m} Z_{2,K} + 4 \left\| \nabla_K \bar{J}(\bar{K}_t, G_t) \right\|^2 + 8H^2 L^2 \eta^2 Z_\infty^2 + 2L^2 r^2. \end{aligned}$$

For the first line, we utilized independence of the random perturbation $\{u_t^i\}$ amongst the agents $i \in [m]$. For the last line, we utilized the definition $Z_{2,K} := \max_t \frac{1}{m} \sum_{i=1}^m \text{Var}((z_t^i)_K)$.

Next, note that

$$\begin{aligned}
& \frac{1}{m^2} \sum_{i=1}^m \mathbb{E}^t \left\| (z_t^i)_{g^i}(K_t^i, g_t^i) \right\|^2 \\
&= \frac{1}{m^2} \sum_{i=1}^m \text{Var}((z_t^i)_{g^i}) + \left\| \nabla_{g^i} J_r^i(K_t^i, g_t^i) \right\|^2 \\
&= \frac{1}{m^2} \sum_{i=1}^m \text{Var}((z_t^i)_{g^i}) + 2 \left\| \nabla_{g^i} J^i(K_t^i, g_t^i) \right\|^2 + 2 \left\| \nabla_{g^i} J^i(K_t^i, g_t^i) - \nabla_{g^i} J_r^i(K_t^i, g_t^i) \right\|^2 \\
&\leq \frac{1}{m^2} \left(\sum_{i=1}^m \text{Var}((z_t^i)_{g^i}) + 4 \left\| \nabla_{g^i} J^i(\bar{K}_t, g_t) \right\|^2 + 4 \left\| \nabla_{g^i} J^i(\bar{K}_t, g_t^i) - \nabla_{g^i} J^i(K_t^i, g_t^i) \right\|^2 \right) \\
&\quad + 2 \frac{1}{m^2} \left\| \nabla_{g^i} J^i(K_t^i, g_t^i) - \nabla_{g^i} J_r^i(K_t^i, g_t^i) \right\|^2 \\
&\leq \frac{Z_{2,G}}{m} + 4 \left\| \nabla_G \bar{J}(\bar{K}_t, G_t) \right\|^2 + 8H^2 L^2 \eta^2 \frac{Z_\infty^2}{m^2} + 2L^2 \frac{r^2}{m}.
\end{aligned}$$

Putting everything together, we find that

$$\begin{aligned}
& \mathbb{E}^t [\bar{J}(\bar{K}_{t+1}, G_{t+1})] \\
&\leq \bar{J}(\bar{K}_t, G_t) - \frac{\eta}{2} \left\| \nabla \bar{J}(\bar{K}_t, G_t) \right\|^2 + 2L^2 \eta \left(2r^2 + \eta^2 H^2 (Z_\infty^2 \frac{Z_\infty^2}{m^2}) \right) \\
&\quad + \frac{L\eta^2}{2} \mathbb{E}^t \left(\left\| \frac{1}{m} \sum_{i=1}^m (z_t^i)_K(K_t^i, g_t^i) \right\|^2 + \sum_{i=1}^m \left\| (z_t^i)_{g^i}(K_t^i, g_t^i) \right\|^2 \right) \\
&\leq \bar{J}(\bar{K}_t, G_t) - \frac{\eta}{2} \left\| \nabla \bar{J}(\bar{K}_t, G_t) \right\|^2 + 2L^2 \eta (2r^2 + 2\eta^2 H^2 Z_\infty^2) \\
&\quad + \frac{L\eta^2}{2} \left(\frac{Z_{2,K} + Z_{2,G}}{m} + 4 \left\| \nabla \bar{J}(\bar{K}_t, G_t) \right\|^2 + 16H^2 L^2 \eta^2 Z_{\infty,K}^2 + 4L^2 r^2 \right) \\
&\leq \bar{J}(\bar{K}_t, G_t) - \left(\frac{\eta}{2} - \frac{4L\eta^2}{2} \right) \left\| \nabla \bar{J}(\bar{K}_t, G_t) \right\|^2 \\
&\quad + (2L^2 \eta + L^3 \eta^2) 2r^2 + (2L^2 \eta^3 H^2 + 8H^2 L^3 \eta^4) Z_\infty^2 + \frac{L\eta^2}{2m} (Z_{2,K} + Z_{2,G}) \\
&\leq \left(1 - \frac{\eta}{4\mu} \right) \bar{J}(\bar{K}_t, G_t) + 6\eta L^2 r^2 + 3L^2 \eta^3 H^2 Z_\infty^2 + \frac{L\eta^2}{2m} (Z_{2,K} + Z_{2,G}).
\end{aligned}$$

The penultimate inequality holds by picking $\eta L \leq \frac{1}{8}$, so that $\frac{\eta}{2} - \frac{4\eta^2 L}{2} \geq \frac{\eta}{4}$. The last inequality is a result of the PL inequality and some algebraic simplifications using $\eta L \leq \frac{1}{8}$.

Continuing, using the choice

$$r \leq \sqrt{\frac{\epsilon}{720L\mu}},$$

we get that

$$\mathbb{E}^t[\Delta_{t+1}] \leq \left(1 - \frac{\eta}{4\mu}\right)\Delta_t + \frac{\eta^2 L}{2} \left(\frac{Z_{2,K} + Z_{2,G}}{m} + 6\eta L H^2 Z_\infty^2\right) + \frac{\eta}{120\mu}\epsilon.$$

□

We next show that with a relatively large (but constant) probability, the joint sequence (\bar{K}_t, G_t) produced by the algorithm remains in the stable region \mathcal{G}^0 , facilitating the local smoothness conditions necessary for Lemma 21 to hold.

Lemma 22. *With probability larger than $4/5$, (K_t, G_t) remains in the region \mathcal{G}_0 for the duration of the algorithm.*

Proof. We define a supermartingale sequence Y_t . For convenience, we suppose we actually run the algorithm for $2T$ iterations. Also, to lighten the notation define

$$M_2 := \frac{Z_{2,K} + Z_{2,G}}{m} + 6\eta L H^2 Z_\infty^2$$

Define $\tau = \min_{t \in [2T]} \{\Delta_t > 10\bar{J}(K_0, G_0)\}$. The sequence Y_t is defined as follows:

$$Y_t = \Delta_{t \wedge \tau} 1_{\tau > e} + (2T - \tau) \left[\frac{\eta^2 L}{2} M_2 + \frac{\eta}{120\mu} \epsilon \right], \quad t = 0, 1, \dots, 2T.$$

By a similar calculation to the proof of Theorem 2 in [20], we can show that this is a supermartingale sequence. Then via Doob's supermartingale inequality, we have that

$$\begin{aligned} \mathbb{P} \left(\max_{0 \leq t \leq 2T} Y_t \geq \nu \right) &\leq \frac{\mathbb{E}[Y_0]}{\nu} \\ &\leq \frac{1}{\nu} \left(\mathbb{E}[\Delta_0 1_{\tau > 0}] + (2T) \left[\frac{\eta^2 L}{2} M_2 + \frac{\eta}{120\mu} \epsilon \right] \right) \\ &\leq \frac{1}{\nu} \left(\Delta_0 + 2\eta\mu L M_2 + \frac{\epsilon}{30} \right) \\ &\quad + \frac{1}{\nu} \left((2T) \left[\frac{\eta^2 L}{2} M_2 + \frac{\eta}{120\mu} \epsilon \right] \right). \end{aligned}$$

Then, continuing with the calculations, choosing $\nu = \frac{1}{10\bar{J}_0}$, $\eta \leq \frac{\epsilon}{240L\mu M_2}$ and $T = \frac{4\mu}{\eta} \log(120\Delta_0/\epsilon)$, it follows that

$$\mathbb{P} \left(\max_{0 \leq t \leq 2T} Y_t \geq \nu \right) \leq \frac{1}{\nu} \left(m\Delta_0 + \frac{m\epsilon}{5} \log(120\Delta_0/\epsilon) \right) \leq \frac{1}{5},$$

if we choose ϵ small enough such that $\epsilon \log(120\Delta_0/\epsilon) \leq 5\Delta_0$. Suppose $\tau = t$ for some $0 \leq t \leq 2T$, i.e. an escape from \mathcal{G}^0 happened during the first $2T$ iterations. This implies that $Y_t \geq \Delta_t \geq \nu$. Therefore,

$$\mathbb{P}(\tau \leq 2T) \leq \mathbb{P} \left(\max_{0 \leq t \leq 2T} Y_t \geq \nu \right) \leq \frac{1}{5}.$$

This completes the proof. □

Proof of Theorem 1. Observe that for any $0 \leq t \leq T$,

$$\begin{aligned}\mathbb{E}^t[\Delta_{t+1}1_{\tau>t+1}] &\leq \mathbb{E}^t[\Delta_{t+1}1_{\tau>t}] \\ &= \mathbb{E}^t[\Delta_{t+1}]1_{\tau>t}\end{aligned}$$

We will then bound $\mathbb{E}^t[\Delta_{t+1}1_{\tau>t+1}]$ by bounding $\mathbb{E}^t[\Delta_{t+1}]1_{\tau>t}$. There are two cases to consider.

1. The first is when $\tau > t$. In this case, by Lemma 21, then

$$\mathbb{E}^t[\Delta_{t+1}] \leq \left(1 - \frac{\eta}{4\mu}\right)\Delta_t + \frac{\eta^2 L}{2}(M_2) + \frac{\eta}{120\mu}\epsilon.$$

2. The second case is when $\tau \leq t$. In this case,

$$\mathbb{E}^t[\Delta_{t+1}]1_{\tau>t} = 0.$$

Thus, combining the two cases, we obtain that

$$\mathbb{E}^t[\Delta_{t+1}]1_{\tau>t} \leq \left(1 - \frac{\eta}{4\mu}\right)\Delta_t 1_{\tau>t} + \frac{\eta^2 L}{2}(M_2) + \frac{\eta}{120\mu}\epsilon.$$

Taking expectations over \mathcal{F}_t and then using induction, we find that

$$\begin{aligned}\mathbb{E}^t[\Delta_{t+1}]1_{\tau>t+1} &\leq \left(1 - \frac{\eta}{4\mu}\right)^{t+1}\Delta_0 + \left(\frac{\eta L^2}{2}M_2 + \eta\frac{\epsilon}{120\mu}\right)\sum_{i=0}^t\left(1 - \frac{\eta}{4\mu}\right)^i \\ &\leq \left(1 - \frac{\eta}{4\mu}\right)^{t+1}\Delta_0 + \frac{2\eta}{L}(M_2) + \frac{4\epsilon}{120}.\end{aligned}$$

After substituting $T = \frac{4\mu}{\eta} \log\left(\frac{120\Delta_0}{\epsilon}\right)$, we find that (assuming for simplicity that T is a multiple of H),

$$\mathbb{E}[\Delta_T 1_{\tau>T}] \leq \frac{\epsilon}{20}.$$

By Lemma 22, we find that

$$\begin{aligned}\mathbb{P}(\Delta_T \geq \epsilon) &\leq \mathbb{P}(\Delta_T 1_{\tau>T} \geq \epsilon) + \mathbb{P}(1_{\tau \leq T}) \\ &\leq \frac{\mathbb{E}[\Delta_T 1_{\tau>T}]}{\epsilon} + \frac{1}{5} \\ &\leq \frac{1}{20} + \frac{1}{5} = \frac{1}{4}.\end{aligned}$$

□

D.2 Theorem 2: Benchmark algorithm convergence

We next analyze the convergence rate of Algorithm 3, shown in Theorem 2. In this subsection, we define $\mathbf{K} = [(K^1)^\top \quad (K^2)^\top \quad \dots \quad (K^m)^\top]$.

Lemma 23. *Suppose $\{(K_t^i, G_t^i)\} \in \mathcal{G}'_0$, where*

$$\mathcal{G}'_0 := \left\{ \{K^i, g^i\}_{i \in [m]} : \frac{1}{m} \sum_{i=1}^m J^i(K^i, g^i) \leq 10 \frac{1}{m} \sum_{i=1}^m J^i(K_0^i, g_0^i) \right\}.$$

Then, if we choose step-size $\eta > 0$ and smoothing radius $r > 0$ such that

$$\eta \leq \min \left\{ \frac{\rho}{2Z_\infty}, \frac{L}{8} \right\}, \quad r \leq \min \left\{ \sqrt{\frac{\epsilon}{720\mu L}}, \rho \right\},$$

the optimality gap satisfies the following bound

$$\mathbb{E}[\Delta_{t+1}] \leq (1 - \frac{\eta}{4\mu})\Delta_t + \frac{\eta^2 L}{2} (Z_{2,K} + Z_{2,G}) + \frac{\eta}{120\mu} \epsilon,$$

where $\Delta_t = \frac{1}{m} \sum_{i=1}^m J^i(K_t^i, g_t^i) - J_{avg}^*$.

Proof. For each agent i , by choosing $\eta \leq \frac{\rho}{2Z_\infty}$, we find that

$$\begin{aligned} \|(K_{t+1}^i, g_{t+1}^i) - (K_t^i, g_t^i)\|^2 &= \|K_t^i - \eta(z_t^i)_K(K_t^i, g_t^i) - K_t\|^2 + \|g_t^i - \eta(z_t^i)_{g^i}(K_t^i, g_t^i) - g_t\|^2 \\ &\leq \eta^2 (Z_\infty^2 + Z_\infty^2) \leq \rho^2. \end{aligned}$$

Since $(\mathbf{K}_t, G_t) \in \mathcal{G}'_0$, by local smoothness, and taking conditional expectation on \mathcal{F}_t (which we denote by \mathbb{E}^t), it follows that

$$\begin{aligned} \mathbb{E}^t J^i(K_{t+1}^i, g_{t+1}^i) &\leq J^i(K_t^i, g_t^i) - \eta \mathbb{E}^t \langle \nabla_K J^i(K_t^i, g_t^i), (z_t^i)_K(K_t^i, g_t^i) \rangle + \frac{L}{2} \mathbb{E}^t \|\eta z_t^i(K_t^i, g_t^i)\|^2 \\ &\quad - \eta \sum_{i=1}^m \mathbb{E}^t \langle \nabla_{g^i} J^i(K_t^i, g_t^i), (z_t^i)_{g^i}(K_t^i, g_t^i) \rangle + \frac{L}{2} \mathbb{E}^t \|\eta (z_t^i)_{g^i}(K_t^i, g_t^i)\|^2 \\ &= J^i(K_t^i, g_t^i) - \eta \langle \nabla J^i(K_t^i, g_t^i), \nabla J_r^i(K_t^i, g_t^i) \rangle + \frac{L\eta^2}{2} \mathbb{E}^t \|(z_t^i)_K(K_t^i, g_t^i)\|^2 \\ &\quad + \frac{L\eta^2}{2} \mathbb{E}^t \|(z_t^i)_{g^i}(K_t^i, g_t^i)\|^2 \end{aligned}$$

We first handle the term $-\eta \langle \nabla J^i(K_t^i, g_t^i), \nabla J_r^i(K_t^i, g_t^i) \rangle$. Observe that

$$\begin{aligned} &-\eta \langle \nabla J^i(K_t^i, g_t^i), \nabla J_r^i(K_t^i, g_t^i) \rangle \\ &= -\eta \langle \nabla J^i(K_t^i, g_t^i), \nabla J^i(K_t^i, g_t^i) \rangle - \eta \langle \nabla_K J^i(K_t^i, g_t^i), \nabla_K J_r^i(K_t^i, g_t^i) - \nabla_K J^i(K_t^i, g_t^i) \rangle \\ &= -\eta \|\nabla J^i(K_t^i, g_t^i)\|^2 + \eta \left(\frac{\|\nabla J^i(K_t^i, g_t^i)\|^2}{2} + \frac{\|\nabla J_r^i(K_t^i, g_t^i) - \nabla J^i(K_t^i, g_t^i)\|^2}{2} \right) \\ &\leq -\frac{\eta}{2} \|\nabla J^i(K_t^i, g_t^i)\|^2 + \frac{\eta}{2} L^2 r^2. \end{aligned}$$

In the last inequality, we used local smoothness and the fact that $r \leq \rho$.

We turn our attention now to $\mathbb{E}^t \|(z_t^i)_K(K_t^i, g_t^i)\|^2$. For notational convenience we use $\text{Var}(\cdot)$ to denote $\mathbb{E}^t[\|\cdot\|^2]$. Then,

$$\begin{aligned} \mathbb{E}^t \|(z_t^i)_K(K_t^i, g_t^i)\|^2 &= \text{Var}((z_t^i)_K) + \|\nabla_K J_r^i(K_t^i, g_t^i)\|^2 \\ &= \text{Var}((z_t^i)_K) + 2\|\nabla_K J^i(K_t^i, g_t^i)\|^2 + 2\|\nabla_K J_r^i(K_t^i, g_t^i) - \nabla_K J^i(K_t^i, g_t^i)\|^2 \\ &\leq Z_{2,K^i} + 2\|\nabla_K J^i(K_t^i, g_t^i)\|^2 + 2L^2r^2. \end{aligned}$$

Next, note that

$$\begin{aligned} \mathbb{E}^t \|(z_t^i)_{g^i}(K_t^i, g_t^i)\|^2 &= \text{Var}((z_t^i)_{g^i}) + \|\nabla_{g^i} J_r^i(K_t^i, g_t^i)\|^2 \\ &= \text{Var}((z_t^i)_{g^i}) + 2\|\nabla_{g^i} J^i(K_t^i, g_t^i)\|^2 + 2\|\nabla_{g^i} J^i(K_t^i, g_t^i) - \nabla_{g^i} J_r^i(K_t^i, g_t^i)\|^2 \\ &\leq Z_{2,g^i} + 2\|\nabla_{g^i} J^i(K_t^i, g_t^i)\|^2 + 2L^2r^2. \end{aligned}$$

Putting everything together, and summing over the agents i , we find that

$$\begin{aligned} &\mathbb{E}^t \left[m\tilde{J}(\mathbf{K}_t, G_t) \right] \\ &\leq \sum_{i=1}^m \left(J^i(K_t^i, g_t^i) - \frac{\eta}{2} \|\nabla J^i(K_t^i, g_t^i)\|^2 + 2L^2\eta r^2 + \frac{L\eta^2}{2} \mathbb{E}^t \left(\|(z_t^i)_K(K_t^i, g_t^i)\|^2 + \|(z_t^i)_{g^i}(K_t^i, g_t^i)\|^2 \right) \right) \\ &\leq m\tilde{J}(\mathbf{K}_t, G_t) - \frac{\eta}{2} \left\| m\nabla \tilde{J}(\mathbf{K}_t, G_t) \right\|^2 + 2mL^2\eta r^2 \\ &\quad + \frac{L\eta^2}{2} \left(mZ_{2,K} + mZ_{2,G} + 2\left\| m\nabla \tilde{J}(\mathbf{K}_t, G_t) \right\|^2 + 4mL^2r^2 \right) \\ &\leq m\tilde{J}(\mathbf{K}_t, G_t) - \left(\frac{\eta}{2} - L\eta^2 \right) \left\| \nabla m\tilde{J}(\mathbf{K}_t, G_t) \right\|^2 \\ &\quad + 2m(L^2\eta + L^3\eta^2)r^2 + mL\eta^2(Z_{2,K} + Z_{2,G}) \\ &\leq \left(1 - \frac{\eta}{4\mu} \right) m\tilde{J}(\mathbf{K}_t, G_t) + 3\eta L^2mr^2 + mL\eta^2(Z_{2,K} + Z_{2,G}). \end{aligned}$$

The penultimate inequality holds by picking $\eta L \leq \frac{1}{8}$, so that $\frac{\eta}{2} - \eta^2 L \geq \frac{\eta}{4}$. The last inequality is a result of the PL inequality and some algebraic simplifications using $\eta L \leq \frac{1}{8}$.

Continuing, using the choice

$$r \leq \sqrt{\frac{\epsilon}{360L\mu}},$$

we get that

$$\mathbb{E}^t[m\Delta_{t+1}] \leq \left(1 - \frac{\eta}{4\mu} \right) m\Delta_t + \frac{\eta^2 L}{2} (m(Z_{2,K} + Z_{2,G})) + \frac{\eta}{120\mu} m\epsilon.$$

□

We next show that with a relatively large (but constant) probability, the joint sequence (\mathbf{K}_t, G_t) produced by the algorithm remains in the stable region \mathcal{G}'_0 , facilitating the local smoothness conditions necessary for Lemma 23 to hold.

Lemma 24. *With probability larger than $4/5$, (\mathbf{K}_t, G_t) remains in the region \mathcal{G}'_0 for the duration of the algorithm.*

Proof. We define a supermartingale sequence Y_t . For convenience, we suppose we actually run the algorithm for $2T$ iterations. Also, to lighten the notation define

$$M_2 := Z_{2,K} + Z_{2,G}$$

Define $\tau := \min_{t \in [T]} \{\Delta_t > 10\tilde{J}_0\}$, where $\tilde{J}_0 := \frac{1}{m} \sum_{i=1}^m J^i(K_0^i, g_0^i)$. The sequence Y_t is defined as follows:

$$Y_t = m\Delta_{t \wedge \tau} 1_{\tau > e} + (T + 1 - \tau) \left[\frac{\eta^2 L}{2} m M_2 + \frac{\eta}{120\mu} m \epsilon \right], \quad t = 0, 1, \dots, 2T.$$

By a similar calculation to the proof of Theorem 2 in [20], we can show that this is a supermartingale sequence. Then via Doob's supermartingale inequality, we have that

$$\begin{aligned} \mathbb{P} \left(\max_{0 \leq t \leq 2T} Y_t \geq \nu \right) &\leq \frac{\mathbb{E}[Y_0]}{\nu} \\ &\leq \frac{1}{\nu} \left(\mathbb{E}[m\Delta_0 1_{\tau > 0}] + m(2T) \left[\frac{\eta^2 L}{2} M_2 + \frac{\eta}{120\mu} \epsilon \right] \right) \\ &\leq \frac{1}{\nu} \left(m\Delta_0 + (2T) \left[\frac{\eta^2 L}{2} m M_2 + \frac{\eta}{120\mu} m \epsilon \right] \right). \end{aligned}$$

Then, continuing with the calculations, choosing $\nu = \frac{1}{10m\tilde{J}_0}$, it follows by $\eta \leq \frac{\epsilon}{240L\mu M_2}$ and $T = \frac{4\mu}{\eta} \log(120\Delta_0/\epsilon)$ that

$$\mathbb{P} \left(\max_{0 \leq t \leq 2T} Y_t \geq \nu \right) \leq \frac{1}{\nu} (m\Delta_0 + \frac{m\epsilon}{5} \log(120\Delta_0/\epsilon)) \leq \frac{1}{5},$$

if we choose ϵ small enough such that $\epsilon \log(120\Delta_0/\epsilon) \leq 5\Delta_0$. Suppose $\tau = t$ for some $0 \leq t \leq 2T$, i.e. an escape from \mathcal{G}'_0 happened during the first $2T$ iterations. This implies that $Y_t \geq \Delta_t \geq \nu$. Therefore,

$$\mathbb{P}(\tau \leq 2T) \leq \mathbb{P} \left(\max_{0 \leq t \leq 2T} Y_t \geq \nu \right) \leq \frac{1}{5}.$$

This completes the proof. \square

Proof of Theorem 2. The proof is similar to that of Theorem 1. Observe that for any $0 \leq t \leq T$,

$$\begin{aligned} \mathbb{E}^t[\Delta_{t+1} 1_{\tau > t+1}] &\leq \mathbb{E}^t[\Delta_{t+1} 1_{\tau > t}] \\ &= \mathbb{E}^t[\Delta_{t+1}] 1_{\tau > t} \end{aligned}$$

We will then bound $\mathbb{E}^t[\Delta_{t+1}1_{\tau>t+1}]$ by bounding $\mathbb{E}^t[\Delta_{t+1}]1_{\tau>t}$. There are two cases to consider.

1. The first is when $\tau > t$. In this case, by Lemma 21, then

$$\mathbb{E}^t[\Delta_{t+1}] \leq \left(1 - \frac{\eta}{4\mu}\right)\Delta_t + \frac{\eta^2 L}{2}(Z_{2,K} + Z_{2,G}) + \frac{\eta}{120\mu}\epsilon.$$

2. The second case is when $\tau \leq t$. In this case,

$$\mathbb{E}^t[\Delta_{t+1}]1_{\tau>t} = 0.$$

Thus, combining the two cases, we obtain that

$$\mathbb{E}^t[\Delta_{t+1}]1_{\tau>t} \leq \left(1 - \frac{\eta}{4\mu}\right)\Delta_t 1_{\tau>t} + \frac{\eta^2 L}{2}(Z_{2,K} + Z_{2,G}) + \frac{\eta}{120\mu}\epsilon.$$

Taking expectations over \mathcal{F}_t and then using induction, we find that

$$\begin{aligned} \mathbb{E}^t[\Delta_{t+1}]1_{\tau>t+1} &\leq \left(1 - \frac{\eta}{4\mu}\right)^{t+1}\Delta_0 + \left(\frac{\eta L^2}{2}(Z_{2,K} + Z_{2,G}) + \eta\frac{\epsilon}{120\mu}\right) \sum_{i=0}^t \left(1 - \frac{\eta}{4\mu}\right)^i \\ &\leq \left(1 - \frac{\eta}{4\mu}\right)^{t+1}\Delta_0 + \frac{2\eta}{L}(Z_{2,K} + Z_{2,G}) + \frac{4\epsilon}{120}. \end{aligned}$$

After substituting $T = \frac{4\mu}{\eta} \log\left(\frac{120\Delta_0}{\epsilon}\right)$, we find that (assuming for simplicity that T is a multiple of H),

$$\mathbb{E}[\Delta_T 1_{\tau>T}] \leq \frac{\epsilon}{20}.$$

By Lemma 24, we find that

$$\begin{aligned} \mathbb{P}(\Delta_T \geq \epsilon) &\leq \mathbb{P}(\Delta_T 1_{\tau>T} \geq \epsilon) + \mathbb{P}(1_{\tau \leq T}) \\ &\leq \frac{\mathbb{E}[\Delta_T 1_{\tau>T}]}{\epsilon} + \frac{1}{5} \\ &\leq \frac{1}{20} + \frac{1}{5} = \frac{1}{4}. \end{aligned}$$

□

D.3 Theorem 3: Comparison of federated and benchmark algorithm

D.3.1 Properties of the one-point and two-point estimators

We first analyze properties of the one-point and two-point estimators. Recall that the one-point and two-point zeroth order estimators are defined as follows for the LQR tracking problem with random initialization:

$$\begin{aligned} z_r^1((K_t^i, g_t^i), u_t^i, (x_0)_t^i) &:= \frac{d}{r} \hat{J}^i((K_t^i, g_t^i) + r u_t^i; (x_0)_t^i) u_t^i, \\ z_r^2((K_t^i, g_t^i), u_t^i, (x_0)_t^i) &:= \frac{d}{2r} (J^i((K_t^i, g_t^i) + r u_t^i; (x_0)_t^i) - J^i((K_t^i, g_t^i) - r u_t^i; (x_0)_t^i)) u_t^i, \end{aligned}$$

where $d = n \cdot k + k$ (i.e. dimension of K plus dimension of g). We also define a smoothed version of J^i with smoothing radius r , as follows

$$J_r^i(K, g) = \mathbb{E}_U[J^i((K, g) + rU)], \quad U \sim \text{Unif}(\mathbb{B}^d).$$

We first study the one-point estimator

Proposition 9 (Properties of one-point estimator). *For each $i \in [m]$, the one-point estimator of J^i , $z_r^1(\cdot)$, satisfies (ignoring effects of the truncation \hat{J}^i)*

$$\mathbb{E}[z_r^1((K_t^i, g_t^i), u_t^i, (x_0)_t^i)] = \nabla J_r^i(K_t^i, g_t^i),$$

and if $r \leq \rho$, the local radius of smoothness defined earlier,

$$\|\nabla J_r^i(\cdot) - \nabla J^i(\cdot)\| \leq rL\|\cdot\|.$$

In addition,

$$\max\|z_r^1((K_t^i, g_t^i), u_t^i, (x_0)_t^i)\| \leq \frac{10d}{r} J_0^i.$$

Finally,

$$\max \mathbb{E}[\|z_r^1((K_t^i, g_t^i), u_t^i, (x_0)_t^i)\|^2] \leq \frac{100d^2}{r^2} (J_0^i)^2$$

Therefore, for the one-point estimator, letting $J_0 := \max_{i \in [m]} J_0^i$, we can set

$$Z_\infty^i = 10 \frac{d}{r} J_0^i, \quad Z_\infty = 10 \frac{d}{r} J_0, \quad Z_{2,G} + Z_{2,K} = \frac{1}{m} \sum_{i=1}^m \frac{100d^2}{r^2} (J_0^i)^2$$

Proof. The first result regarding the expectation of z_r^1 and bias of ∇J_r^i follows from analysis in [20]. The bound on the maximum size of z_r^1 follows from the bound which we imposed on the size of the estimator. Finally the last result on the maximum variance of the estimator is a consequence of the bound on the maximum size of the estimator. \square

We next study the two-point estimator.

Proposition 10 (Properties of two-point estimator). *For each $i \in [m]$, the two-point estimator of J^i , $z_r^2(\cdot)$, satisfies*

$$\mathbb{E}[z_r^2((K_t^i, g_t^i), u_t^i, (x_0)_t^i)] = \nabla J_r^i(K_t^i, g_t^i),$$

and if $r \leq \rho$, the local radius of smoothness defined earlier,

$$\|\nabla J_r^i(\cdot) - \nabla J^i(\cdot)\| \leq rL\|\cdot\|.$$

In addition,

$$\max\|z_r^2((K_t^i, g_t^i), u_t^i, (x_0)_t^i)\| \leq d\lambda,$$

where λ is the local Lipschitz parameter defined earlier. Finally,

$$\max \mathbb{E}[\|z_r^2((K_t^i, g_t^i), u_t^i, (x_0)_t^i)\|^2] \leq d\lambda^2$$

Therefore, for the two-point estimator, we can set

$$Z_\infty^i = d\lambda, \quad Z_{2,G} + Z_{2,K} = d\lambda^2$$

Proof. The first result regarding the expectation of z_r^2 and bias of ∇J_r^i follows from analysis in [11]. The bound on the maximum size and variance of z_r^2 follows from the proof of Corollary 2 in [20]. \square

D.3.2 Convergence for one-point estimator

We next specialize Theorem 1 to the case when one-point estimator is used.

Corollary 1 (Convergence of Theorem 1 for one-point estimator). *Consider the parameter settings in Theorem 1. Suppose $1 \leq H \leq \frac{1}{\sqrt{\eta L m}}$ (if $\eta L m \geq 1$, then we just set $H = 1$). Then, to ensure the ϵ -convergence in Theorem 1, we need at least*

$$T = \tilde{O}\left(\frac{d^2 J_0^2 L^2 \mu^2}{m\epsilon^2}\right)$$

steps, where \tilde{O} hides log terms.

Proof. Suppose $H \leq \frac{1}{\sqrt{\eta L m}}$. Then, recall in Theorem 1 we need $\eta \leq \frac{\rho}{2HZ_\infty}$ as well as $\eta \leq \frac{m\epsilon}{240\mu L(Z_{2,G} + Z_{2,K} + 6m\eta L H^2 Z_\infty^2)}$. From the first requirement, we need that

$$\eta \leq \frac{\rho}{2Z_\infty/(\sqrt{\eta m L})} \implies \sqrt{\eta} \leq \frac{\rho m L}{2Z_\infty}.$$

From the second requirement, we need that

$$\eta \leq \frac{m\epsilon}{240\mu L(Z_{2,G} + Z_{2,K} + 6m\eta L H^2 Z_\infty^2)} \implies \eta \leq \frac{m\epsilon}{240\mu L(Z_{2,G} + Z_{2,K} + 6Z_\infty^2)}.$$

Asymptotically as $\epsilon \rightarrow 0$, the second requirement dominates, so we focus on this. Since $Z_\infty = \frac{10d}{r} J_0$ and $Z_{2,G} + Z_{2,K} = \frac{100d^2}{r^2} J_0^2$, we have that to ensure convergence we need

$$\eta \leq \frac{m\epsilon r^2}{240\mu L(8d^2 J_0^2)} \implies T \geq \tilde{O}\left(\frac{d^2 J_0^2 L^2 \mu^2}{m\epsilon^2}\right),$$

where \tilde{O} hides log terms. \square

In a similar vein, we specialize Theorem 2 to the case when the one-point estimator is used.

Corollary 2 (Convergence of Theorem 2 for one-point estimator). *Consider the parameter settings in Theorem 2. Then, to ensure the ϵ -convergence in Theorem 2, we need at least*

$$T = \tilde{O}\left(\frac{d^2 J_0^2 L^2 \mu^2}{\epsilon^2}\right)$$

steps, where \tilde{O} hides log terms.

Proof. Since $Z_\infty = \frac{10d}{r} J_0$ and $Z_{2,G} + Z_{2,K} = \frac{100d^2}{r^2} J_0^2$, we have that to ensure convergence we need

$$\eta \leq \frac{\epsilon r^2}{240\mu L(2d^2 J_0^2)} \implies T \geq \tilde{O}\left(\frac{d^2 J_0^2 L^2 \mu^2}{\epsilon^2}\right),$$

where \tilde{O} hides log terms. □

D.3.3 Convergence for two-point estimator

We next specialize Theorem 1 to the case when the two-point estimator is used.

Corollary 3 (Convergence of Theorem 1 for two-point estimator). *Consider the parameter settings in Theorem 1. Suppose $1 \leq H \leq \frac{1}{\sqrt{\eta L m d}}$. Then, to ensure the ϵ -convergence in Theorem 1, we need at least*

$$T = \tilde{O}\left(\frac{d\lambda^2 L^2 \mu^2}{m\epsilon^2}\right)$$

steps, where \tilde{O} hides log terms.

Proof. Suppose $H \leq \frac{1}{\sqrt{d\eta L m}}$ (if $\eta L m d \geq 1$, we just set $H = 1$). Then, recall in Theorem 1 we need $\eta \leq \frac{\rho}{2HZ_\infty}$ as well as $\eta \leq \frac{m\epsilon}{240\mu L(Z_{2,G} + Z_{2,K} + 6m\eta L H^2 Z_\infty^2)}$. From the first requirement, we need that

$$\eta \leq \frac{\rho}{2Z_\infty/(\sqrt{d\eta m L})} \implies \sqrt{\eta} \leq \frac{d\rho m L}{2Z_\infty}.$$

From the second requirement, we need that

$$\eta \leq \frac{m\epsilon}{240\mu L(Z_{2,G} + Z_{2,K} + 6m\eta L H^2 Z_\infty^2)} \implies \eta \leq \frac{m\epsilon}{240\mu L(Z_{2,G} + Z_{2,K} + 6Z_\infty^2/d)}.$$

Asymptotically as $\epsilon \rightarrow 0$, the second requirement dominates, so we focus on this. Since for the two-point estimator, $Z_\infty = d\lambda$ and $Z_{2,G} + Z_{2,K} = d\lambda^2$, we have that to ensure convergence we need

$$\eta \leq \frac{m\epsilon r^2}{240\mu L(8d\lambda^2)} \implies T \geq \tilde{O}\left(\frac{d\lambda^2 L^2 \mu^2}{m\epsilon^2}\right),$$

where \tilde{O} hides log terms. □

We next specialize Theorem 2 to the case when the two-point estimator is used.

Corollary 4 (Convergence of Theorem 2 for two-point estimator). *Consider the parameter settings in Theorem 2. Then, to ensure the ϵ -convergence in Theorem 2, we need at least*

$$T = \tilde{O}\left(\frac{d\lambda^2 L^2 \mu^2}{\epsilon^2}\right)$$

steps, where \tilde{O} hides log terms.

Proof. Since for the two-point estimator, $Z_\infty = d\lambda$ and $Z_{2,G} + Z_{2,K} = d\lambda^2$, we have that to ensure convergence we need

$$\eta \leq \frac{\epsilon r^2}{240\mu L(2d\lambda^2)} \implies T \geq \tilde{O}\left(\frac{d\lambda^2 L^2 \mu^2}{\epsilon^2}\right),$$

where \tilde{O} hides log terms. □

Proof of Theorem 3. The result in Theorem 3 is a consequence of Corollary 3 and Corollary 4. □

D.4 Discussion

We provide a discussion here of the main results in the paper.

Probabilistic convergence Inherently, maintaining stability of the matrix K , *i.e.* $\rho(A+BK) < 1$, is a critical issue for LQR learning, since when K is unstable, the infinite-horizon LQR cost can diverge [7]. This stability requirement is in tension with the constant accumulation of noise in the learning procedure due to the stochastic zeroth-order updates. For this reason, the convergence result in Theorems 1 and 2 holds only with a certain level of probability, and we note that convergence need not hold if the total number of steps T were to be much larger than the specified number in the theorem [14].

Dependence on d and ϵ From Corollaries 1 through 4, we note that in the simple case $H = 1$, the convergence result for both Algorithm 2 and Algorithm 3 display a (1) linear dependence on d for one-point estimator and quadratic dependence on d for two-point estimator, and (2) ϵ^{-2} scaling for one-point estimator and ϵ^{-1} for two-point estimator. These observations are consistent with the sample complexity results for the centralized LQR (where x^* is 0) in [20], and reflect the significantly lower variance that the two-point zeroth-order estimator inherently enjoys over the one-point estimator [12].

Dependence on m and effect of H Under appropriate settings of H (scaling as $\max\left(\tilde{O}\left(\frac{\sqrt{T}}{\sqrt{d\lambda m}}\right), 1\right)$), Theorem 3 says that convergence for the federated approach is $\tilde{O}(m)$ times faster. This follows primarily from the periodic averaging of the K matrices, which reduces the variance in the optimization problem up to a factor of m in a similar vein to the results in the federated averaging/local SGD literature [24], [27]. However, while ideally we can allow the step-size η to increase as m increases, thus accelerating convergence, in practice, due to stability concerns, there exists some maximum step-size which permits convergence. This echoes the observation about mini-batching in [20]. In addition, we note that for the communication interval H , its effect asymptotically as $\epsilon \rightarrow 0$ is clear — in that case, $\eta \leq \frac{1}{Lm}$ (one-point estimator requirement so that H exceeds 1) or $\eta \leq \frac{1}{Lmd}$ (two-point estimator requirement so that H exceeds 1) will both hold, since we also need $\eta \leq \tilde{O}(\epsilon^2)$ and $\eta \leq \tilde{O}(\epsilon)$ to hold. In that case, H can indeed exceed 1, and as long as it does not exceed $\frac{1}{\sqrt{\eta Lm}}$ (one-point estimator) or $\frac{1}{\sqrt{\eta Lmd}}$ (two-point estimator), the convergence speedup in terms of m is preserved. However, in the case when ϵ is merely a small but fixed constant such that $\frac{1}{mL} \leq \eta \leq \frac{mer^2}{240\mu L(2d^2J_0^2)}$ (one-point estimator) or $\frac{1}{dmL} \leq \eta \leq \frac{mer^2}{240\mu L(8d\lambda^2)}$, to theoretically preserve the convergence result we need to set $H = 1$. As seen from the analysis of Theorem 1, larger H causes the local controllers K_t^i to deviate from their average, causing the algorithm to accumulate variance. Empirically, we found that $H = 1$ consistently worked best.

Dependence on local radius ρ , local Lipschitz parameter λ and local smoothness parameter L We note that we adopted conservation parameter choices for the local radius/Lipschitz/smoothness constants — by Proposition 2, we know there exists ρ' , L' such that (a) if $\{(K^i, g^i)\}_{i \in [m]} \in \mathcal{G}'_0$, for any single $i \in [m]$, when $\|(K', (g^i)') - (K, g^i)\| \leq \rho'$, $\|\nabla J^i(K', (g^i)') - \nabla J^i(K, g^i)\| \leq L'\|(K', (g^i)') - (K, g^i)\|$, and that (b) there exists ρ'' , L'' such that if $(K, G) \in \mathcal{G}_0$, if $\|(K', G') - (K, G)\| \leq \rho''$, then $\|\nabla J(K', G') - \nabla J(K, G)\| \leq L''\|(K', G') - (K, G)\|$. In our assumptions for the main results, we used $\rho = \min(\rho', \rho'')$ and $L = \max(L', L'')$. However, it could well be that the true local radii and smoothness constants in cases (a) and (b) differ. In particular, as m increases, the Clarifying the local radii, Lipschitz and smoothness parameters in different cases requires working through the local Lipschitz and smoothness analysis we performed earlier, and we leave it to future work.

Appendix E More details about numerical results

E.1 Linear system

We compare our proposed algorithm with the benchmark algorithm for a federated LQR tracking problem. We use a LQR problem with A,B, Q, R matrices

each in $\mathbb{R}^{3 \times 3}$, which are as follows:

$$A = \begin{bmatrix} 1 & 0 & -10 \\ -1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, B = \begin{bmatrix} 1 & -10 & 0 \\ 0 & 1 & 0 \\ -1 & 0 & 1 \end{bmatrix}, Q = \begin{bmatrix} 2 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 2 \end{bmatrix}, R = \begin{bmatrix} 5 & -3 & 0 \\ -3 & 5 & -2 \\ 0 & -2 & 5 \end{bmatrix}.$$

Each initial state x_0^i is sampled uniformly at random from the canonical basis vectors and the discount factor γ is set to 0.9. The tracking targets x^{i*} for each agent are randomly sampled from a zero mean Gaussian with covariance matrix (i) $\frac{1}{10} \times I$ and (ii) $\frac{1}{2} \times I$. A two-point estimator is used. In the main paper, case (i) where the targets are drawn from $N(0, \frac{1}{10} \times I)$ was shown. For case (ii), where the targets are drawn from $N(0, \frac{1}{10} \times I)$, we first see in Figure 2a that the average optimality gap decreases significantly faster for the federated algorithm (with 8 agents) compared to that for a single agent. In addition, for case (ii), we see that for $H = 1$, the as m increases, the maximum step-size permitting convergence does seem to be m times larger for the federated algorithm than for the independent algorithm. However, the improvement weakens as H increases.

In general, empirically, we found that increasing H led to a decrease in the probability of convergence, suggesting that where possible, agents should communicate more frequently, as seen from Figure 3. In Figure 4a and Figure 4b, we see the number of iterations required to converge to an average optimality gap of $\epsilon = 0.05$ as the number of agents m increases. We see that the decrease in number of iterations is linear as m increases from 10^0 to 10^1 in both cases, and plateaus beyond that.

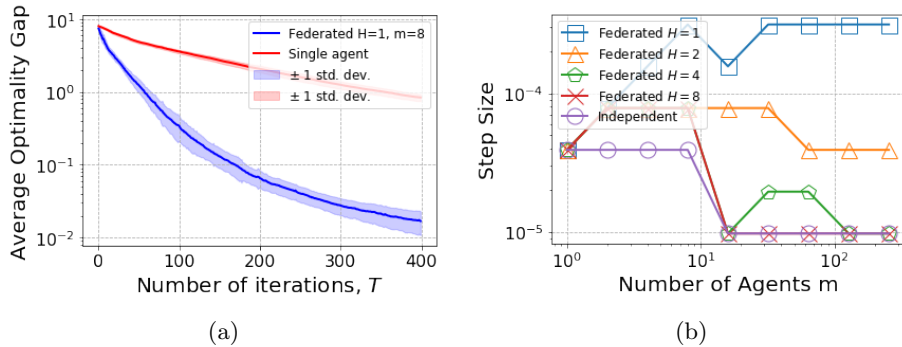


Figure 2: (a): Errors at each iteration of proposed algorithm and single agent baseline on a linear system.(mean and standard deviation for 20 runs) (b): Maximum step-size that allows for convergence to an error tolerance of $\epsilon = 0.1$.

E.2 Nonlinear system

We provide more details about cartpole problem. The problem is also known as the cart-inverted pendulum [35], and the underlying physics can be described as follows. Consider an inverted pendulum on the top of a cart with mass M ,

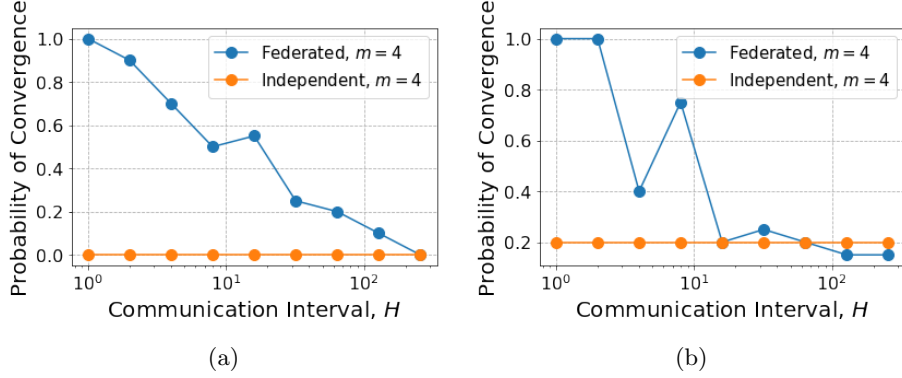


Figure 3: (a): The probability of convergence with different communication intervals H for case(i). (b): The probability of convergence with different communication intervals H for case(ii).

where a force (input) F is applied to the cart. We assume the rod has length L , the mass at the end of the rod is m , and the moment of inertia is I about its center of gravity. The distance of the hinge from the origin is given by p . Other parameters include gravitational acceleration g , viscous friction of the cart c , and viscous friction at the hinge v . Every parameter has SI units. The state is given by

$$x = \begin{bmatrix} p \\ \theta \\ \dot{p} \\ \dot{\theta} \end{bmatrix}$$

The dynamic equations of motions for θ and p are given as follows:

$$\begin{aligned} (M + m)\ddot{p} - mL \cos(\theta)\ddot{\theta} + c\dot{p} + mL \sin(\theta)\dot{\theta}^2 &= F \\ -mL \cos(\theta)\ddot{p} + (I + mL^2)\ddot{\theta} + v\dot{\theta} - mgL \sin(\theta) &= 0. \end{aligned}$$

When there is no force F , it can be shown that an equilibrium point of the system satisfies $\theta = 0, \dot{\theta} = 0, \dot{p} = 0$. Linearizing around this equilibrium point yields a continuous-time linear system,

$$\frac{d}{dt}x = Ax + BF$$

for some $A \in \mathbb{R}^{4 \times 4}, B \in \mathbb{R}^{4 \times 1}$. It can be shown that this linearized system is controllable [35]. Empirically, when initial conditions are not very far away from the equilibrium point, a linear controller can stabilize the system [35]. We set $M = 1, m = 0.1, g = 9.8, L = 1, R = 0.1$, and

$$Q = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 100 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \times 0.1$$

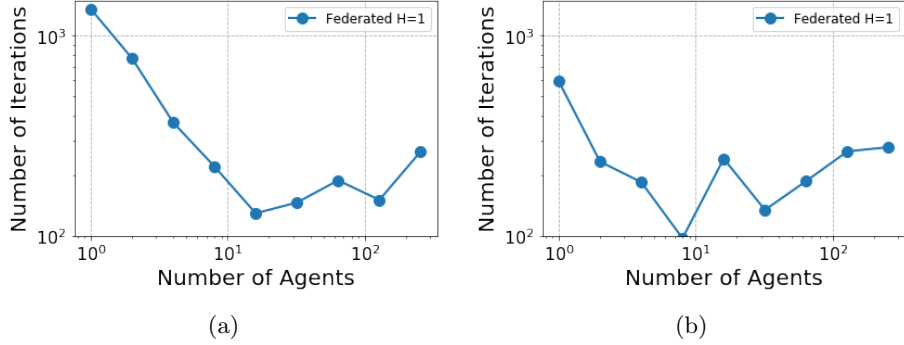


Figure 4: (a): Numbers of iterations used to reach an error tolerance of $\epsilon = 0.05$ with different numbers of agents m for case(i). (b): Numbers of iterations used to reach an error tolerance of $\epsilon = 0.1$ with different numbers of agents m for case(ii). Average numbers of iterations of successful runs within 20 runs are shown in both (a) and(b) using the Maximum step-size that allows convergence.

The discount factor γ is set to 0.95, and every episode ends when time step reaches 300. We use $m = 5$ agents and set the tracking targets for agent i as

$$x^{i*} = \begin{bmatrix} (-2 + i) \times 0.5 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

We show the convergence process for our proposed algorithm and independently learning baseline in Figure 5. We can observe that using federated learning indeed can accelerate the speed of convergence and make the convergence more stable.

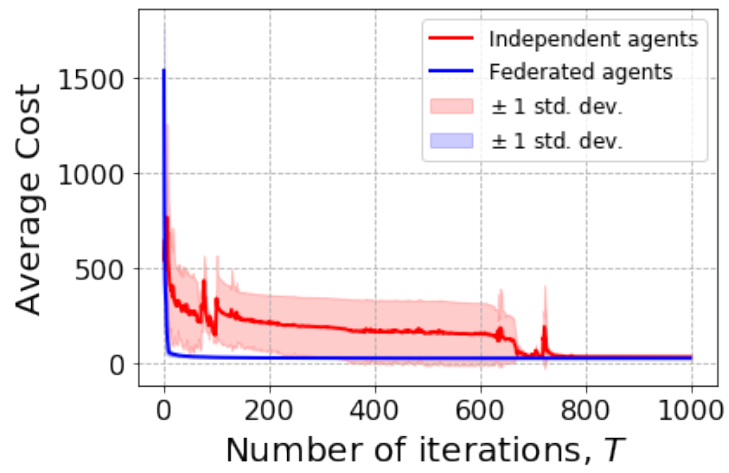


Figure 5: Average cost at each iterations for the non-linear cartpole system in 10 runs.