



HAL
open science

Thompson Sampling for Linear-Quadratic Control Problems

Marc Abeille, Alessandro Lazaric

► **To cite this version:**

Marc Abeille, Alessandro Lazaric. Thompson Sampling for Linear-Quadratic Control Problems. AIS-TATS 2017 - 20th International Conference on Artificial Intelligence and Statistics, Apr 2017, Fort Lauderdale, United States. hal-01493564

HAL Id: hal-01493564

<https://inria.hal.science/hal-01493564>

Submitted on 21 Mar 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Thompson Sampling for Linear-Quadratic Control Problems

Marc Abeille

Inria Lille - Nord Europe, Team SequeL

Alessandro Lazaric

Abstract

We consider the exploration-exploitation tradeoff in linear quadratic (LQ) control problems, where the state dynamics is linear and the cost function is quadratic in states and controls. We analyze the regret of Thompson sampling (TS) (a.k.a. posterior-sampling for reinforcement learning) in the frequentist setting, i.e., when the parameters characterizing the LQ dynamics are fixed. Despite the empirical and theoretical success in a wide range of problems from multi-armed bandit to linear bandit, we show that when studying the frequentist regret TS in control problems, we need to trade-off the frequency of sampling optimistic parameters and the frequency of switches in the control policy. This results in an overall regret of $O(T^{2/3})$, which is significantly worse than the regret $O(\sqrt{T})$ achieved by the optimism-in-face-of-uncertainty algorithm in LQ control problems.

1 Introduction

One of the most challenging problems in reinforcement learning (RL) is how to effectively trade off exploration and exploitation in an unknown environment. A number of learning methods has been proposed in finite Markov decision processes (MDPs) and they have been analyzed in the PAC-MDP (see e.g., [13]) and the regret framework (see e.g., [8]). The two most popular approaches to address the exploration-exploitation trade-off are the optimism-in-face-of-uncertainty (OFU) principle, where optimistic policies are selected according to upper-confidence bounds on the true MDP parameters, and the Thompson sampling (TS) strategy¹, where

¹In RL literature, TS has been introduced by Strens [14] and it is often referred to as posterior-sampling for reinforcement learning (PSRL).

random MDP parameters are selected from a posterior distribution and the corresponding optimal policy is executed. Despite their success in finite MDPs, extensions of these methods and their analyses to continuous state-action spaces are still rather limited. Osband et al. [11] study how to randomize the parameters of a linear function approximator to induce exploration and prove regret guarantees in the finite MDP case. Osband and Van Roy [10] develops a specific TS method applied to the more complex case of neural architectures with significant empirical improvements over alternative exploration strategies, although with no theoretical guarantees. In this paper, we focus on a specific family of continuous state-action MDPs, the linear quadratic (LQ) control problems, where the state transition is linear and the cost function is quadratic in the state and the control. Despite their specific structure, LQ models are very flexible and widely used in practice (e.g., to track a reference trajectory). If the parameter θ defining dynamics and cost is known, the optimal control can be computed explicitly as a linear function of the state with an appropriate gain. On the other hand, when θ is unknown, an exploration-exploitation trade-off needs to be solved. Bittanti and Campi [6] and Campi and Kumar [7], first proposed an optimistic approach to this problem, showing that the performance of an adaptive control strategy asymptotically converges to the optimal control. Building on this approach and the OFU principle, Abbasi-Yadkori and Szepesvári [1] proposed a learning algorithm (OFU-LQ) with $O(\sqrt{T})$ cumulative regret. Abbasi-Yadkori and Szepesvári [2] further studied how the TS strategy, could be adapted to work in the LQ control problem. Under the assumption that the true parameters of the model are drawn from a known prior, they show that the so-called Bayesian regret matches the $O(\sqrt{T})$ bound of OFU-LQ.

In this paper, we analyze the regret of TS in LQ problems in the more challenging frequentist case, where θ is a fixed parameter, with no prior assumption of its value. The analysis of OFU-LQ relies on three main ingredients: **1**) optimistic parameters, **2**) lazy updates (the control policy is updated only a logarithmic number of times) and **3**) concentration inequalities for

regularized least-squares used to estimate the unknown parameter θ . While we build on previous results for the least-squares estimates of the parameters, points **1**) and **2**) should be adapted for TS. Unfortunately, the Bayesian regret analysis of TS in [2] does not apply in this case, since no prior is available on θ . Furthermore, we show that existing frequentist regret analysis for TS in linear bandit [5] cannot be generalized to the LQ case. This requires deriving a novel line of proof in which we first prove that TS has a constant probability to sample an optimistic parameter (i.e., an LQ system whose optimal expected average cost is smaller than the true one) and then we exploit the LQ structure to show how being optimistic allows to directly link the regret to the controls operated by TS over time and eventually bound them. Nonetheless, this analysis reveals a critical trade-off between the frequency with which new parameters are sampled (and thus the chance of being optimistic) and the regret cumulated every time the control policy changes. In OFU-LQ this trade-off is easily solved by construction: the lazy update guarantees that the control policy changes very rarely and whenever a new policy is computed, it is guaranteed to be optimistic. On the other hand, TS relies on the *random* sampling process to obtain optimistic models and if this is not done *frequently enough*, the regret can grow unbounded. This forces TS to favor short episodes and we prove that this leads to an overall regret of order $O(T^{2/3})$ in the one-dimensional case (i.e., both states and controls are scalars), which is significantly worse than the $O(\sqrt{T})$ regret of OFU-LQ.

2 Preliminaries

The control problem. We consider the discrete-time infinite-horizon linear quadratic (LQ) control problem. Let $x_t \in \mathbb{R}^n$ be the state of the system and $u_t \in \mathbb{R}^d$ be the control at time t ; an LQ problem is characterized by linear dynamics and a quadratic cost function

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

where A_* and B_* are *unknown* matrices and Q and R are *known* positive definite matrices of appropriate dimension. We summarize the unknown parameters in $\theta_*^\top = (A_*, B_*)$. The noise process ϵ_{t+1} is zero-mean and it satisfies the following assumption.

Assumption 1. $\{\epsilon_t\}_t$ is a \mathcal{F}_t -martingale difference sequence, where \mathcal{F}_t is the filtration which represents the information knowledge up to time t .

In LQ, the objective is to design a closed-loop control policy $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^d$ mapping states to controls that minimizes the average expected cost

$$J_\pi(\theta_*) = \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=0}^T c(x_t, u_t) \right], \quad (2)$$

with $x_0 = 0$ and $u_t = \pi(x_t)$. Standard theory for LQ control guarantees that the optimal policy is linear in the state and that the corresponding average expected cost is the solution of a Riccati equation.

Proposition 1 (Thm.16.6.4 in [9]). *Under Asm. 1 and for any LQ system with parameters $\theta^\top = (A, B)$ such that (A, B) is stabilizable², and p.d. cost matrices Q and R , the optimal solution of Eq. 2 is given by*

$$\begin{aligned} \pi(\theta) &= K(\theta)x_t, \quad J(\theta) = \text{Tr}(P(\theta)), \\ K(\theta) &= -(R + B^\top P(\theta)B)^{-1} B^\top P(\theta)A, \\ P(\theta) &= Q + A^\top P(\theta)A + A^\top P(\theta)BK(\theta) \end{aligned} \quad (3)$$

where $\pi(\theta)$ is the optimal policy, $J(\theta)$ is the corresponding average expected cost, $K(\theta)$ is the optimal gain, and $P(\theta)$ is the unique solution to the Riccati equation associated with the control problem. Finally, we also have that $A + BK(\theta)$ is asymptotically stable.

For notational convenience, we use $H(\theta) = (I \ K(\theta)^\top)^\top$, so that the closed loop dynamics $A + BK(\theta)$ can be equivalently written as $\theta^\top H(\theta)$. We introduce further assumptions about the LQ systems we consider.

Assumption 2. *We assume that the LQ problem is characterized by parameters (A_*, B_*, Q, R) such that the cost matrices Q and R are symmetric p.d., and $\theta_* \in \mathcal{S}$ where³ $\mathcal{S} = \{\theta \in \mathbb{R}^{(n+d) \times n} \text{ s.t. } \text{Tr}(P(\theta)) \leq D \text{ and } \text{Tr}(\theta\theta^\top) \leq S^2\}$.*

While Asm. 1 basically guarantees that the linear model in Eq. 1 is correct, Asm. 2 restricts the control parameters to the admissible set \mathcal{S} . This is used later in the learning process and it replaces Asm. A2-4 in [1] in a synthetic way, as shown in the following proposition.

Proposition 2. *Given an admissible set \mathcal{S} as defined in Asm. 2, we have **1**) $\mathcal{S} \subset \{\theta^\top = (A, B) \text{ s.t. } (A, B) \text{ is stabilizable}\}$, **2**) \mathcal{S} is compact, and **3**) there exists $\rho < 1$ and $C < \infty$ positive constants such that $\rho = \sup_{\theta \in \mathcal{S}} \|A + BK(A, B)\|_2$ and $C = \sup_{\theta \in \mathcal{S}} \|K(\theta)\|_2$ ⁴.*

As an immediate result, any system with $\theta \in \mathcal{S}$ is stabilizable, and therefore, Asm. 2 implies that Prop. 1 holds. Finally, we derive a result about the regularity of the Riccati solution, which we later use to relate the regret to the controls performed by TS.

Lemma 1. *Under Asm. 1 and for any LQ with parameters $\theta^\top = (A, B)$ and cost matrices Q and R satisfying Asm. 2, let $J(\theta) = \text{Tr}(P(\theta))$ be the optimal solution of*

² (A, B) is stabilizable if there exists a control gain matrix K s.t. $A + BK$ is stable (i.e., all eigenvalues are in $(-1, 1)$).

³Even if $P(\theta)$ is not defined for every θ , we extend its domain of definition by setting $P(\theta) = +\infty$.

⁴We use $\|\cdot\|$ and $\|\cdot\|_2$ to denote the Frobenius and the 2-norm respectively.

Eq. 2. Then, the mapping $\theta \in \mathcal{S} \rightarrow \text{Tr}(P(\theta))$ is continuously differentiable. Furthermore, let $A_c(\theta) = \theta^\top H(\theta)$ be the closed-loop matrix, then the directional derivative of $J(\theta)$ in a direction $\delta\theta$, denoted as $\nabla J(\theta)^\top \delta\theta$, where $\nabla J(\theta) \in \mathbb{R}^{(n+d) \times n}$ is the gradient of J , is the solution of the Lyapunov equation

$$\nabla J(\theta)^\top \delta\theta = A_c(\theta)^\top \nabla J(\theta)^\top \delta\theta A_c(\theta) + C(\theta, \delta\theta) + C(\theta, \delta\theta)^\top,$$

where $C(\theta, \delta\theta) = A_c(\theta)^\top P(\theta) \delta\theta^\top H(\theta)$.

The learning problem. At each time t , the learner chooses a policy π_t , it executes the induced control $u_t = \pi_t(x_t)$ and suffers a cost $c_t = c(x_t, u_t)$. The performance is measured by the cumulative regret up to time T as $R_T = \sum_{t=0}^T (c_t^{\pi_t} - J_{\pi_*}(\theta_*))$, where at each step the difference between the cost of the controller c_t^π and the expected average cost $J_{\pi_*}(\theta_*)$ of the optimal controller π_* is measured. Let (u_0, \dots, u_t) be a sequence of controls and $(x_0, x_1, \dots, x_{t+1})$ be the corresponding states, then θ^* can be estimated by regularized least-squares (RLS). Let $z_t = (x_t, u_t)^\top$, for any regularization parameter $\lambda \in \mathbb{R}_+^*$, the design matrix and the RLS estimate are defined as

$$V_t = \lambda I + \sum_{s=0}^{t-1} z_s z_s^\top; \quad \hat{\theta}_t = V_t^{-1} \sum_{s=0}^{t-1} z_s x_{s+1}.$$

For notational convenience, we use $W_t = V_t^{-1/2}$. We recall a concentration inequality for RLS estimates.

Proposition 3 (Thm. 2 in [3]). *We assume that ϵ_t are conditionally and component-wise sub-Gaussian of parameter L and that $\mathbb{E}(\epsilon_{t+1} \epsilon_{t+1}^\top | \mathcal{F}_t) = I$. Then for any $\delta \in (0, 1)$ and any \mathcal{F}_t -adapted sequence (z_0, \dots, z_t) , the RLS estimator $\hat{\theta}_t$ is such that*

$$\text{Tr}\left((\hat{\theta}_t - \theta_*)^\top V_t (\hat{\theta}_t - \theta_*)\right) \leq \beta_t(\delta)^2, \quad (4)$$

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

$$\beta_t(\delta) = nL \sqrt{2 \log \left(\frac{\det(V_t)^{1/2}}{\det(\lambda I)^{1/2}} \right)} + \lambda^{1/2} S. \quad (5)$$

Further, when $\|z_t\| \leq Z$,

$$\frac{\det(V_t)}{\det(\lambda I)} \leq (n+d) \log \left(1 + TZ^2 / \lambda(n+d) \right).$$

At any step t , we define the ellipsoid $\mathcal{E}_t^{\text{RLS}} = \{\theta \in \mathbb{R}^d \mid \|\theta - \hat{\theta}_t\|_{V_t} \leq \beta_t(\delta')\}$ centered in $\hat{\theta}_t$ with orientation V_t and radius $\beta_t(\delta')$, with $\delta' = \delta/(4T)$. Finally, we report a standard result of RLS that, together with Prop. 3, shows that the prediction error on the points z_t used to construct the estimator $\hat{\theta}_t$ is cumulatively small.

```

Input:  $\hat{\theta}_0, V_0 = \lambda I, \delta, T, \tau, t_0 = 0$ 
1: Set  $\delta' = \delta/(8T)$ 
2: for  $t = \{0, \dots, T\}$  do
3:   if  $\det(V_t) > 2 \det(V_0)$  or  $t \geq t_0 + \tau$  then
4:     while  $\tilde{\theta}_t \notin \mathcal{S}$  do
5:       Sample  $\eta_t \sim \mathcal{D}^{\text{TS}}$ 
6:       Compute  $\tilde{\theta}_t = \hat{\theta}_t + \beta_t(\delta') V_t^{-1/2} \eta_t$ 
7:     end while
8:     Let  $V_0 = V_t, t_0 = t,$ 
9:   else
10:     $\tilde{\theta}_t = \hat{\theta}_{t-1}$ 
11:   end if
12:   Execute control  $u_t = K(\tilde{\theta}_t) x_t$ 
13:   Move to state  $x_{t+1}$ , receive cost  $c(x_t, u_t)$ 
14:   Compute  $V_{t+1}$  and  $\hat{\theta}_{t+1}$ 
15: end for
    
```

Figure 1: Thompson sampling algorithm.

Proposition 4 (Lem. 10 in [1]). *Let $\lambda \geq 1$, for any arbitrary \mathcal{F}_t -adapted sequence (z_0, z_1, \dots, z_t) , let V_{t+1} be the corresponding design matrix, then*

$$\sum_{s=0}^t \min(\|z_s\|_{V_s^{-1}}^2, 1) \leq 2 \log \frac{\det(V_{t+1})}{\det(\lambda I)}. \quad (6)$$

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

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

3 Thompson Sampling for LQR

We introduce a specific instance of TS for learning in LQ problems obtained as a modification of the algorithm proposed in [2], where we replace the Bayesian structure and the Gaussian prior assumption with a generic randomized process and we modify the update rule. The algorithm is summarized in Alg. 1. At any step t , given the RLS-estimate $\hat{\theta}_t$ and the design matrix V_t , TS samples a *perturbed* parameter $\tilde{\theta}_t$. In order to ensure that the sampling parameter is indeed admissible, we re-sample it until a valid $\tilde{\theta}_t \in \mathcal{S}$ is obtained. Denoting as $\mathcal{R}_{\mathcal{S}}$ the rejection sampling operator associated with the admissible set \mathcal{S} , we define $\tilde{\theta}_t$ as

$$\tilde{\theta}_t = \mathcal{R}_{\mathcal{S}}(\hat{\theta}_t + \beta_t(\delta') W_t \eta_t), \quad (7)$$

where $W_t = V_t^{-1/2}$ and every coordinate of the matrix $\eta_t \in \mathbb{R}^{(n+d) \times (n+d)}$ is a random sample drawn i.i.d. from $\mathcal{N}(0, 1)$. We refer to this distribution as \mathcal{D}^{TS} . Notice that such sampling does not need to be associated with an actual posterior over θ^* but it just need to randomize parameters coherently with the RLS estimate and the uncertainty captured in V_t . Let $\gamma_t(\delta) = \beta_t(\delta') n \sqrt{2(n+d) \log(2n(n+d)/\delta)}$,

then the high-probability TS ellipsoid $\mathcal{E}_t^{\text{TS}} = \{\theta \in \mathbb{R}^d \mid \|\theta - \hat{\theta}_t\|_{V_t} \leq \gamma_t(\delta')\}$ is defined so that any parameter $\tilde{\theta}_t$ belongs to it with $1 - \delta/8$ probability.

Given the parameter $\tilde{\theta}_t$, the gain matrix $K(\tilde{\theta}_t)$ is computed and the corresponding optimal control $u_t = K(\tilde{\theta}_t)x_t$ is applied. As a result, the learner observes the cost $c(x_t, u_t)$ and the next state x_{t+1} , and V_t and $\hat{\theta}_t$ are updated accordingly. Similar to most of RL strategies, the updates are not performed at each step and the same estimated optimal policy $K(\tilde{\theta}_t)$ is kept constant throughout an *episode*. Let V_0 be the design matrix at the beginning of an episode, then the episode is terminated upon two possible conditions: **1**) the determinant condition of the design matrix is doubled (i.e., $\det(V_t) \geq 2\det(V_0)$) or **2**) a maximum length condition is reached. While the first condition is common to all RL strategies, here we need to force the algorithm to interrupt episodes as soon as their length exceeds τ steps. The need for this additional termination condition is intrinsically related to the TS nature and it is discussed in detail in the next section.

4 Theoretical analysis

We prove the first frequentist regret bound for TS in LQ systems of dimension 2 ($n = 1, d = 1$). In order to isolate the steps which explicitly rely on this restriction, whenever possible we derive the proof in the general $n + d$ -dimensional case.

Theorem 1. *Consider the LQ system in Eq. 1 of dimension $n = 1$ and $d = 1$. Under Asm. 1 and 2 for any $0 < \delta < 1$, the cumulative regret of TS over T steps is bounded w.p. at least $1 - \delta$ as⁵*

$$R(T) = \tilde{O}\left(T^{2/3}\sqrt{\log(1/\delta)}\right).$$

This result is in striking contrast with previous results in multi-armed and linear bandit where the frequentist regret of TS is $O(\sqrt{T})$ and the Bayesian analysis of TS in control problems where the regret is also $O(\sqrt{T})$. As discussed in the introduction, the frequentist regret analysis in control problems introduces a critical trade-off between the frequency of selecting optimistic models, which guarantees small regret in bandit problems, and the reduction of the number of policy switches, which leads to small regret in control problems. Unfortunately, this trade-off cannot be easily balanced and this leads to a final regret of $O(T^{2/3})$. Sect. 4.2 provides a more detailed discussion on the challenges of bounding the frequentist regret of TS in LQ problems.

4.1 Setting the Stage

Concentration events. We introduce the following high probability events.

⁵Further details can be recovered from the proof.

Definition 1. *Let $\delta \in (0, 1)$ and $\delta' = \delta/(8T)$ and $t \in [0, T]$. We define the event (RLS estimate concentration) $\hat{E}_t = \{\forall s \leq t, \|\hat{\theta}_s - \theta^*\|_{V_s} \leq \beta_s(\delta')\}$ and the event (parameter $\tilde{\theta}_s$ concentrates around $\hat{\theta}_s$) $\tilde{E}_t = \{\forall s \leq t, \|\tilde{\theta}_s - \hat{\theta}_s\|_{V_s} \leq \gamma_s(\delta')\}$.*

We also introduce a high probability event on which the states x_t are bounded almost surely.

Definition 2. *Let $\delta \in (0, 1)$, X, X' be two problem dependent positive constants and $t \in [0, T]$. We define the event (bounded states) $\bar{E}_t = \{\forall s \leq t, \|x_s\| \leq X \log \frac{X'}{\delta}\}$.*

Then we have that $\hat{E} := \hat{E}_T \subset \dots \subset \hat{E}_1$, $\tilde{E} := \tilde{E}_T \subset \dots \subset \tilde{E}_1$ and $\bar{E} := \bar{E}_T \subset \dots \subset \bar{E}_1$. We show that these events do hold with high probability.

Lemma 2. $\mathbb{P}(\hat{E} \cap \tilde{E}) \geq 1 - \delta/4$.

Corollary 1. *On $\hat{E} \cap \tilde{E}$, $\mathbb{P}(\bar{E}) \geq 1 - \delta/4$. Thus, $\mathbb{P}(\hat{E} \cap \tilde{E} \cap \bar{E}) \geq 1 - \delta/2$.*

Lem. 2 leverages Prop. 3 and the sampling distribution \mathcal{D}^{TS} to ensure that $\hat{E} \cap \tilde{E}$ holds w.h.p. Furthermore, Corollary 1 ensures that the states remains bounded w.h.p. on the events $\hat{E} \cap \tilde{E}$.⁶ As a result, the proof can be derived considering that both parameters concentrate and that states are bounded, which we summarize in the sequence of events $E_t = \hat{E}_t \cap \tilde{E}_t \cap \bar{E}_t$, which holds with probability at least $1 - \delta/2$ for all $t \in [0, T]$.

Regret decomposition. Conditioned on the filtration \mathcal{F}_t and event E_t , we have $\theta^* \in \mathcal{E}_t^{\text{RLS}}$, $\tilde{\theta}_t \in \mathcal{E}_t^{\text{TS}}$ and $\|x_t\| \leq X$. We directly decompose the regret and bound it on this event as [1, Sect. 4.2]

$$R(T) = \underbrace{\sum_{t=0}^T \{J(\tilde{\theta}_t) - J(\theta^*)\} \mathbb{1}\{E_t\}}_{R^{\text{TS}}} + \underbrace{(R_1^{\text{RLS}} + R_2^{\text{RLS}} + R_3^{\text{RLS}}) \mathbb{1}\{E_t\}}_{R^{\text{RLS}}} \quad (8)$$

where R^{RLS} is decomposed into the three components

$$\begin{aligned} R_1^{\text{RLS}} &= \sum_{t=0}^T \{\mathbb{E}(x_{t+1}^\top P(\tilde{\theta}_{t+1})x_{t+1} | \mathcal{F}_t) - x_t^\top P(\tilde{\theta}_t)x_t\}, \\ R_2^{\text{RLS}} &= \sum_{t=0}^T \mathbb{E}[x_{t+1}^\top (P(\tilde{\theta}_t) - P(\tilde{\theta}_{t+1}))x_{t+1} | \mathcal{F}_t], \\ R_3^{\text{RLS}} &= \sum_{t=0}^T \{z_t^\top \tilde{\theta}_t P(\tilde{\theta}_t) \tilde{\theta}_t^\top z_t - z_t^\top \theta^* P(\tilde{\theta}_t) \theta^{*\top} z_t\}. \end{aligned}$$

Before entering into the details of how to bound each of these components, in the next section we discuss what are the main challenges in bounding the regret.

⁶This non-trivial result is directly collected from the bounding-the-state section of [1].

4.2 Related Work and Challenges

Since the RLS estimator is the same in both TS and OFU, the regret terms R_1^{RLS} and R_3^{RLS} can be bounded as in [1]. In fact, R_1^{RLS} is a martingale by construction and it can be bounded by Azuma inequality. The term R_3^{RLS} is related to the difference between the *true* next expected state $\theta_*^\top z_t$ and the *predicted* next expected state $\tilde{\theta}_t^\top z_t$. A direct application of RLS properties makes this difference small by construction, thus bounding R_3^{RLS} . Finally, the R_2^{RLS} term is directly affected by the changes in model from any two time instants (i.e., $\tilde{\theta}_t$ and $\tilde{\theta}_{t+1}$), while R^{TS} measures the difference in optimal average expected cost between the true model θ_* and the sampled model $\tilde{\theta}_t$. In the following, we denote by $R_{2,t}^{\text{RLS}}$ and R_t^{TS} the elements at time t of these two regret terms and we refer to them as *consistency regret* and *optimality regret* respectively.

Optimistic approach. OFU-LQ explicitly bounds both regret terms directly by construction. In fact, the lazy update of the control policy allows to set to zero the consistency regret $R_{2,t}^{\text{RLS}}$ in all steps but when the policy changes between two episodes. Since in OFU-LQ an episode terminates only when the determinant of the design matrix is doubled, it is easy to see that the number of episodes is bounded by $O(\log(T))$, which bounds R_2^{RLS} as well (with a constant depending on the bounding of the state X and other parameters specific of the LQ system).⁷ At the same time, at the beginning of each episode an optimistic parameter $\tilde{\theta}_t$ is chosen, i.e., $J(\tilde{\theta}_t) \leq J(\theta_*)$, which directly ensures that R_t^{TS} is upper bounded by 0 at each time step.

Bayesian regret. The lazy PSRL algorithm in [2] has the same lazy update as OFUL and thus it directly controls R_2^{RLS} by a small number of episodes. On the other hand, the random choice of $\tilde{\theta}_t$ does not guarantee optimism at each step anymore. Nonetheless, the regret is analyzed in the Bayesian setting, where θ_* is drawn from a known prior and the regret is evaluated *in expectation* w.r.t. the prior. Since $\tilde{\theta}_t$ is drawn from a posterior constructed from the same prior as θ_* , in expectation its associated $J(\tilde{\theta}_t)$ is the same as $J(\theta_*)$, thus ensuring that $\mathbb{E}[R_t^{\text{TS}}] = 0$.

Frequentist regret. When moving from Bayesian to frequentist regret, this argument does not hold anymore and the (positive) deviations of $J(\tilde{\theta}_t)$ w.r.t. $J(\theta_*)$ has to be bounded in high probability. Abbasi-Yadkori and Szepesvári [1] exploits the linear structure of LQ problems to reuse arguments originally developed in

the linear bandit setting. Similarly, we could leverage on the analysis of TS for linear bandit by Agrawal and Goyal [5] to derive a frequentist regret bound. Agrawal and Goyal [5] partition the (potentially infinite) arms into *saturated* and *unsaturated* arms depending on their estimated value and their associated uncertainty (i.e., an arm is saturated when the uncertainty of its estimate is smaller than its performance gap w.r.t. the optimal arm). In particular, the uncertainty is measured using confidence intervals derived from a concentration inequality similar to Prop. 3. This suggests to use a similar argument and classify policies as saturated and unsaturated depending on their value. Unfortunately, this proof direction cannot be applied in the case of LQR. In fact, in an LQ system θ the performance of a policy π is evaluated by the function $J_\pi(\theta)$ and the policy uncertainty should be measured by a confidence interval constructed as $|J_\pi(\theta_*) - J_\pi(\tilde{\theta}_t)|$. Despite the concentration inequality in Prop. 3, we notice that neither $J_\pi(\theta_*)$ nor $J_\pi(\tilde{\theta}_t)$ may be finite, since π may not stabilize the system θ_* (or $\tilde{\theta}_t$) and thus incur an infinite cost. As a result, it is not possible to introduce the notion of saturated and unsaturated policies in this setting and another line of proof is required. Another key element in the proof of [5] for TS in linear bandit is to show that TS has a constant probability p to select optimistic actions and that this contributes to reduce the regret of any non-optimistic step. In our case, this translates to requiring that TS selects a system $\tilde{\theta}_t$ whose corresponding optimal policy is such that $J(\tilde{\theta}_t) \leq J(\theta_*)$. Lem. 3 shows that this happens with a constant probability p . Furthermore, we can show that optimistic steps reduce the regret of non-optimistic steps, thus effectively bounding the optimality regret R^{TS} . Nonetheless, this is not compatible with a small consistency regret. In fact, we need optimistic parameters $\tilde{\theta}_t$ to be sampled *often enough*. On the other hand, bounding the consistency regret R_2^{RLS} requires to reduce the switches between policies as much as possible (i.e., number of episodes). If we keep the same number of episodes as with the lazy update of OFUL (i.e., about $\log(T)$ episodes), then the number of sampled points is as small as $T/(T - \log(T))$. While OFU-LQ guarantees that any policy update is optimistic by construction, with TS, only a fraction $T/(p(T - \log(T)))$ of steps would be optimistic *on average*. Unfortunately, such small number of optimistic steps is no longer enough to derive a bound on the optimality regret R^{TS} . Summarizing, in order to derive a frequentist regret bound for TS in LQ systems, we need the following ingredient **1)** constant probability of optimism, **2)** connection between optimism and R^{TS} without using the saturated and unsaturated argument, **3)** a suitable trade-off between lazy updates to bound the consistency regret and frequent updates to guarantee small optimality regret.

⁷Notice that the consistency regret is not specific to LQ systems but it is common to all regret analyses in RL (see e.g., UCRL [8]) except for episodic MDPs and it is always bounded by keeping under control the number of switches of the policy (i.e., number of episodes).

4.3 Bounding the Optimality Regret R^{TS}

R^{TS} decomposition. We define the ‘‘extended’’ filtration $\mathcal{F}_t^x = (\mathcal{F}_{t-1}, x_t)$. Let K be the (random) number of episodes up to time T , $\{t_k\}_{k=1}^K$ be the steps when the policy is updated, i.e., when a new parameter $\tilde{\theta}$ is sampled, and let T_k be the associated length of each episode, then we can further decompose R^{TS} as

$$R^{\text{TS}} = \sum_{k=0}^K T_k \underbrace{\left(J(\tilde{\theta}_{t_k}) - \mathbb{E}[J(\tilde{\theta}_{t_k}) | \mathcal{F}_{t_k}^x, E_{t_k}] \right)}_{R_{t_k}^{\text{TS},1}} \mathbf{1}_{E_{t_k}} + \sum_{k=0}^K T_k \underbrace{\left\{ \mathbb{E}[J(\tilde{\theta}_{t_k}) | \mathcal{F}_{t_k}^x, E_{t_k}] - J(\theta_*) \right\}}_{R_{t_k}^{\text{TS},2}} \mathbf{1}_{E_{t_k}}. \quad (9)$$

We focus on the second regret term that we redefine $R_{t_k}^{\text{TS},2} = \Delta_t$ for any $t = t_k$ for notational convenience.

Optimism and expectation. Let $\Theta^{\text{opt}} = \{\theta : J(\theta) \leq J(\theta_*)\}$ be the set of optimistic parameters (i.e., LQ systems whose optimal average expected cost is lower than the true one). Then, for any $\theta \in \Theta^{\text{opt}}$, the per-step regret Δ_t is bounded by:

$$\begin{aligned} \Delta_t &\leq (\mathbb{E}[J(\tilde{\theta}_t) | \mathcal{F}_t^x, E_t] - J(\theta)) \mathbf{1}_{E_t}, \\ &\leq \left| J(\theta) - \mathbb{E}[J(\tilde{\theta}_t) | \mathcal{F}_t^x, E_t] \right| \mathbf{1}_{E_t}, \text{ which implies that} \\ \Delta_t &\leq \mathbb{E} \left[\left| J(\tilde{\theta}) - \mathbb{E}[J(\tilde{\theta}) | \mathcal{F}_t^x, E_t] \right| \mathbf{1}_{\tilde{E}_t} \mid \mathcal{F}_t^x, \hat{E}_t, \bar{E}_t, \tilde{\theta} \in \Theta^{\text{opt}} \right], \end{aligned}$$

where we use first the definition of the optimistic parameter set, then bounding the resulting quantity by its absolute value, and finally switch to the expectation over the optimistic set, since the inequality is true for any $\tilde{\theta} \in \Theta^{\text{opt}}$. While this inequality is true for any sampling distribution, it is convenient to select it equivalent to the sampling distribution of TS. Thus, we set $\tilde{\theta} = \mathcal{R}_S(\hat{\theta}_t + \beta_t(\delta')W_t\eta)$ with η is component wise Gaussian $\mathcal{N}(0, 1)$ and obtain

$$\begin{aligned} \Delta_t &\leq \mathbb{E} \left[\left| J(\tilde{\theta}_t) - \mathbb{E}[J(\tilde{\theta}_t) | \mathcal{F}_t^x, E_t] \right| \mathbf{1}_{\tilde{E}_t} \mid \mathcal{F}_t^x, \hat{E}_t, \bar{E}_t, \tilde{\theta}_t \in \Theta^{\text{opt}} \right], \\ &\leq \frac{\mathbb{E} \left[\left| J(\tilde{\theta}_t) - \mathbb{E}[J(\tilde{\theta}_t) | \mathcal{F}_t^x, E_t] \right| \mathbf{1}_{\tilde{E}_t} \mid \mathcal{F}_t^x, \hat{E}_t, \bar{E}_t \right]}{\mathbb{P}(\tilde{\theta}_t \in \Theta^{\text{opt}} \mid \mathcal{F}_t^x, \hat{E}_t)}. \end{aligned}$$

At this point we need to show that the probability of sampling an optimistic parameter $\tilde{\theta}_t$ is constant at any step t . This result is proved in the following lemma.

Lemma 3. *Let $\Theta^{\text{opt}} := \{\theta \in \mathbb{R}^d \mid J(\theta) \leq J(\theta^*)\}$ be the set of optimistic parameters and $\hat{\theta}_t = \mathcal{R}_S(\hat{\theta}_t + \beta_t(\delta')W_t\eta)$ with η be component-wise normal $\mathcal{N}(0, 1)$, then in the one-dimensional case ($n=1$ and $d=1$)*

$$\forall t \geq 0, \mathbb{P}(\tilde{\theta}_t \in \Theta^{\text{opt}} \mid \mathcal{F}_t^x, \hat{E}_t) \geq p,$$

where p is a strictly positive constant.

Integrating this result into the previous expression gives

$$\Delta_t \leq \frac{1}{p} \mathbb{E} \left[\left| J(\tilde{\theta}_t) - \mathbb{E}[J(\tilde{\theta}_t) | \mathcal{F}_t^x, E_t] \right| \mid \mathcal{F}_t^x, E_t \right]. \quad (10)$$

The most interesting aspect of this result is that the constant probability of being optimistic allows us to bound the worst-case non-stochastic quantity $\mathbb{E}[J(\tilde{\theta}_t) | \mathcal{F}_t^x] - J(\theta_*)$ depending on $J(\theta_*)$ by an expectation $\mathbb{E}[\left| J(\tilde{\theta}_t) - \mathbb{E}[J(\tilde{\theta}_t) | \mathcal{F}_t^x] \right| \mid \mathcal{F}_t^x]$ up to a multiplicative constant (we drop the events E for notational convenience). The last term is the conditional *absolute deviation* of the performance J w.r.t. the TS distribution. This connection provides a major insight about the functioning of TS, since it shows that TS does not need to have an accurate estimate of θ_* but it should rather reduce the estimation errors of θ_* only on the directions that may translate in larger errors in estimating the objective function J . In fact, we show later that at each step TS chooses a sampling distribution that tends to minimize the expected absolute deviations of J , thus contributing to reduce the deviations in R_t^{TS} .

Variance and gradient. Let $d' = \sqrt{n(n+d)}$, we introduce the mapping f_t from the ball $\mathcal{B}(0, d')$ to \mathbb{R}_+ defined as

$$f_t(\eta) = J(\hat{\theta}_t + \beta_t(\delta')W_t\eta) - \mathbb{E}[J(\tilde{\theta}_t) | \mathcal{F}_t^x, E_t]$$

where the restriction on the ball is here to meet the $\mathcal{E}_t^{\text{TS}}$ confidence ellipsoid of the sampling. Since the perturbation $\eta \sim \mathcal{D}^{\text{TS}}$ is independent of the past, we can rewrite Eq. 10 as

$$\Delta_t \leq \mathbb{E}_{\eta \sim \mathcal{D}^{\text{TS}}} \left[|f_t(\eta)| \mid \eta \in \mathcal{B}(0, d'), \hat{\theta}_t + \beta_t(\delta')W_t\eta \in \mathcal{S} \right].$$

We now need to show that this formulation of the regret is strictly related to the policy executed by TS. We prove the following result (proof in the supplement).

Lemma 4. *Let $\Omega \subset \mathbb{R}^d$ be a convex domain with finite diameter diam . Let p be a non-negative log-concave function on Ω with continuous derivative up to the second order. Then, for all $u \in W^{1,1}(\Omega)$ ⁸ such that $\int_{\Omega} u(z)p(z)dz = 0$ one has*

$$\int_{\Omega} |f(z)|p(z)dz \leq 2\text{diam} \int_{\Omega} \|\nabla f(z)\|p(z)dz$$

Before using the previous result, we relate the gradient of f_t to the gradient of J . Since for any η and any $\theta = \hat{\theta}_t + \beta_t(\delta')W_t\eta$, we have

$$\nabla f_t(\eta) = \beta_t(\delta')W_t \nabla J(\theta)$$

To obtain a bound on the norm of ∇f_t , we apply Prop. 5 (derived from Lem. 1) to get a bound on $\|\nabla J(\theta)\|_{W_t^2}$:

$$\begin{aligned} \|\nabla J(\theta)\|_{W_t^2} &\leq \|A_c(\theta)\|_2^2 \|\nabla J(\theta)\|_{W_t^2} \\ &\quad + 2\|P(\theta)\| \|A_c(\theta)\|_2 \|H(\theta)\|_{W_t^2}. \end{aligned}$$

⁸ $W^{1,1}(\Omega)$ is the Sobolev space of order 1 in $L^1(\Omega)$.

Making use of $\|M\| \leq \text{Tr}(M)$ for any positive definite matrix together with $\text{Tr}(P(\theta)) \leq D$ (Asm. 2) and $\|A_c(\theta)\|_2 \leq \rho$ (Prop. 2),

$$\|\nabla J(\theta)\|_{W_t^2} \leq \rho^2 \|\nabla J(\theta)\|_{W_t^2} + 2D\rho \|H(\theta)\|_{W_t^2},$$

which leads to

$$\|\nabla J(\theta)\|_{W_t^2} \leq 2D\rho/(1 - \rho^2) \|H(\theta)\|_{W_t^2}.$$

We are now ready to use the weighted Poincaré inequality of Lem. 4 to link the expectation of $|f_t|$ to the expectation of its gradient. From Lem. 1, we have $f_t \in W^{1,1}(\Omega)$ and its expectation is zero by construction. On the other hand, the rejection sampling procedure impose that we conditioned the expectation with $\tilde{\theta}_t + \beta_t(\delta')W_t\eta \in \mathcal{S}$ which is unfortunately not convex. However, we can still apply Lem. 4 considering the function $\tilde{f}_t(\eta) = f_t(\eta)\mathbb{1}(\tilde{\theta}_t + \beta_t(\delta')W_t\eta \in \mathcal{S})$ and diameter $\text{diam} = d'$. As a result, we finally obtain

$$\Delta_t \leq \gamma \mathbb{E} \left[\|H(\tilde{\theta}_t)\|_{W_t^2} | \mathcal{F}_t^x \right],$$

where $\gamma = 8\sqrt{n(n+d)}\beta_T(\delta')D\rho/(p(1-\rho^2))$.

From gradient to actions. Recalling the definition of $H(\theta) = (I - K(\theta)^\top)^\top$ we notice that the previous expression bound the regret Δ_t with a term involving the gain $K(\theta)$ of the optimal policy for the sampled parameter θ . This shows that the R^{TS} regret is directly related to the policies chosen by TS. To make such relationship more apparent, we now elaborate the previous expression to reveal the sequence of state-control pairs z_t induced by the policy with gain $K(\tilde{\theta}_t)$. We first plug the bound on Δ_t back into Eq. 9 as

$$R^{\text{TS}} \leq \sum_{k=1}^K T_k \left(R_{t_k}^{\text{TS},1} + \gamma \mathbb{E} \left[\|H(\tilde{\theta}_{t_k})\|_{V_{t_k}^{-1}} | \mathcal{F}_{t_k}^x \right] \right) \mathbb{1}_{E_{t_k}}.$$

We remove the expectation by adding and subtracting the actual realizations of $\tilde{\theta}_{t_k}$ as

$$R_{t_k}^{\text{TS},3} = \mathbb{E} \left[\|H(\tilde{\theta}_{t_k})\|_{V_{t_k}^{-1}} | \mathcal{F}_{t_k}^x \right] - \|H(\tilde{\theta}_{t_k})\|_{V_{t_k}^{-1}}.$$

Thus, one obtains

$$R^{\text{TS}} \leq \sum_{k=1}^K T_k \left(R_{t_k}^{\text{TS},1} + R_{t_k}^{\text{TS},3} + \gamma \|H(\tilde{\theta}_{t_k})\|_{V_{t_k}^{-1}} \right) \mathbb{1}_{E_{t_k}}.$$

Now we want to relate the cumulative sum of the last regret term to $\sum_{t=1}^T \|z_t\|_{V_t^{-1}}$. This quantity represents the prediction error of the RLS, and we know from Prop. 6 that it is bounded w.h.p. We now focus on the one-dimensional case, where x_t is just a scalar value. Noticing that $\|z_t\|_{V_t^{-1}} = |x_t| \|H(\tilde{\theta}_t)\|_{V_t^{-1}}$, one has:

$$\sum_{t=0}^T \|z_t\|_{V_t^{-1}} = \sum_{k=1}^K \left(\sum_{t=t_k}^{t_{k+1}-1} |x_t| \right) \|H(\tilde{\theta}_{t_k})\|_{V_{t_k}^{-1}}.$$

Intuitively, it means that over each episode, the more states are excited (e.g., the larger $\sum_{t=t_k}^{t_{k+1}-1} |x_t|$), the more V_t^{-1} reduces in the direction $H(\tilde{\theta}_{t_k})$. As a result, to ensure that the term $\sum_{k=1}^K T_k \|H(\tilde{\theta}_{t_k})\|_{V_{t_k}^{-1}}$ in R^{TS} is small, it would be sufficient to show that $\sum_{t=t_k}^{t_{k+1}-1} |x_t| \sim T_k$, i.e., that the states provides enough information to learn the system in each chosen direction $H(\tilde{\theta}_{t_k})$. More formally, let assume that there exists a constant α such that $T_k \leq \alpha \sum_{t=t_k}^{t_{k+1}-1} |x_t|$ for all $k \leq K$. Then,

$$\sum_{k=1}^K T_k \|H(\tilde{\theta}_{t_k})\|_{V_{t_k}^{-1}} \leq \alpha \sum_{t=0}^T \|z_t\|_{V_{t_k}^{-1}} \leq 2\alpha \sum_{t=0}^T \|z_t\|_{V_t^{-1}},$$

where we use that $\det(V_t) \leq 2 \det(V_{t_k})$ as guaranteed by the termination condition. Unfortunately, the intrinsic randomness of x_t (triggered by the noise ξ_t) is such that the assumption above is violated w.p. 1. However, in the one-dimensional case, the regret over the episode k can be conveniently written as

$$R_k(T) = \left(\sum_{t=t_k}^{t_{k+1}-1} |x_t|^2 \right) (Q + K(\theta_{t_k})^2 R) - T_k J(\theta_*).$$

As a result, if we set

$$\alpha := X \frac{Q + RC^2}{J(\theta_*)} \geq X \frac{Q + RK(\theta_{t_k})^2}{J(\theta_*)}, \quad (11)$$

whenever $\sum_{t=t_k}^{t_{k+1}-1} \|x_t\| \leq \frac{1}{\alpha} T_k$ then we can directly conclude that $R_k(T)$ is zero. On the other hand, in the opposite case, we have $T_k \leq \alpha \sum_{t=t_k}^{t_{k+1}-1} |x_t|$ and thus we can upper bound the last term in R^{TS} as

$$R^{\text{TS}} \leq \sum_{k=1}^K T_k \left(R_{t_k}^{\text{TS},1} + R_{t_k}^{\text{TS},3} \right) \mathbb{1}_{E_{t_k}} + 2\gamma\alpha \sum_{t=0}^T \|z_t\|_{V_t^{-1}}.$$

4.4 Final bound

Bounding R_1^{RLS} and R_3^{RLS} . These two terms can be bounded following similar steps as in [1]. We report the detailed derivation in the supplement while here we simply report the final bounds

$$R_1^{\text{RLS}} \leq \underbrace{2DX^2 \sqrt{2 \log(4/\delta)}}_{:=\gamma_1} \sqrt{T},$$

and

$$R_3^{\text{RLS}} \leq \underbrace{4SD \sqrt{(1+C^2)X^2 \mu_T(\delta')}}_{:=\gamma_3} \sum_{t=0}^T \|z_t\|_{V_t^{-1}} \mathbb{1}_{E_t},$$

where $\mu_T(\delta') = \beta_T(\delta') + \gamma_T(\delta')$.

Bounding R_2^{RLS} . Since the policy is updated from time to time, the difference of the optimal values

$P(\tilde{\theta}_t) - P(\tilde{\theta}_{t+1})$ is zero unless when the parameters are updated. When it is the case, thanks to the rejection sampling procedure which ensures that every parameters belong to the set \mathcal{S} of Asm. 2, it is trivially bounded by $2D$. Therefore, on event E , one has:

$$R_2^{\text{RLS}} \leq 2X^2DK,$$

where K is the (random) number of episodes. By definition of TS, the updates are triggered either when the $\det(V_t)$ increases by a factor 2 or when the length of the episode is greater than τ . Hence, the number of update can be split into $K = K^{\text{det}} + K^{\text{len}}$, where K^{det} and K^{len} are the number of updates triggered by the two conditions respectively. From Cor. 2, one gets:

$$K \leq (T/\tau + (n+d) \log_2(1 + TX^2(1 + C^2)/\lambda)),$$

and thus

$$R_2^{\text{RLS}} \leq \underbrace{2X^2D(n+d) \log_2(1 + TX^2(1 + C^2)/\lambda)}_{:=\gamma_2} T/\tau.$$

Plugging everything together. We are now ready to bring all the regret terms together and obtain

$$\begin{aligned} R(T) &\leq (2\gamma\alpha + \gamma_3) \sum_{t=0}^T \|z_t\|_{V_t^{-1}} \mathbb{1}_{E_t} + \gamma_2 T/\tau \\ &\quad + \gamma_1 \sqrt{T} + \sum_{k=1}^K T_k (R_{t_k}^{\text{TS},1} + R_{t_k}^{\text{TS},3}) \mathbb{1}_{E_{t_k}} \end{aligned}$$

At this point, the regret bound is decomposed into several parts: 1) the first term can be bounded as $\sum_{t=0}^T \|z_t\|_{V_t^{-1}} = \tilde{O}(\sqrt{T})$ on E using Prop. 4 (see App. E for details) 2) two terms which are already conveniently bounded as T/τ and \sqrt{T} , and 3) two remaining terms from R^{TS} that are almost exact martingales. In fact, T_k is random w.r.t. \mathcal{F}_{t_k} and thus the terms $T_k R_{t_k}^{\text{TS},1}$ and $T_k R_{t_k}^{\text{TS},3}$ are not proper martingale difference sequences. However, we can leverage on the fact that on most of the episodes, the length T_k is not random since the termination of the episode is triggered by the (deterministic) condition $T_k \leq \tau$. Let $\alpha_k = (R_{t_k}^{\text{TS},1} + R_{t_k}^{\text{TS},3}) \mathbb{1}_{E_{t_k}}$, \mathcal{K}^{det} and \mathcal{K}^{len} two set of indexes of cardinality K^{det} and K^{len} respectively, which correspond to the episodes terminated following the determinant or the limit condition respectively. Then, we can write

$$\begin{aligned} \sum_{k=1}^K T_k \alpha_k &= \sum_{k \in \mathcal{K}^{\text{det}}} T_k \alpha_k + \tau \sum_{k \in \mathcal{K}^{\text{len}}} \alpha_k \\ &\leq \sum_{k \in \mathcal{K}^{\text{det}}} T_k \alpha_k + \sum_{k \in \mathcal{K}^{\text{len}}} \tau \alpha_k + \sum_{k \in \mathcal{K}^{\text{det}}} \tau \alpha_k + \sum_{k \in \mathcal{K}^{\text{det}}} \tau \|\alpha_k\| \\ &\leq 2\tau \sum_{k \in \mathcal{K}^{\text{det}}} \|\alpha_k\| + \tau \sum_{k=1}^K \alpha_k. \end{aligned}$$

The first term can be bounded using Lem. 6, which implies that the number of episodes triggered by the determinant condition is only logarithmic. On the other hand the remaining term $\sum_{k=1}^K \alpha_k$ is now a proper martingale and, together with the boundedness of α_k on event E , Azuma inequality directly holds. We obtain

$$\sum_{k=1}^K T_k (R_{t_k}^{\text{TS},1} + R_{t_k}^{\text{TS},3}) \mathbb{1}_{E_{t_k}} = \tilde{O}(\tau \sqrt{K}).$$

w.p. $1 - \delta/2$. Grouping all higher-order terms w.r.t. to T and applying Cor. 2 to bound K , we finally have

$$R(T) \leq C_1 \frac{T}{\tau} + C_2 \tau \sqrt{T/\tau},$$

where C_1 and C_2 are suitable problem-dependent constants. This final bound is optimized for $\tau = O(T^{1/3})$ and it induces the final regret bound $R(T) = O(T^{2/3})$. More details are reported in App. E.

5 Discussion

We derived the first frequentist regret for TS in LQ control systems. Despite the existing results in LQ for optimistic approaches (OFU-LQ), the Bayesian analysis of TS in LQ, and its frequentist analysis in linear bandit, we showed that controlling the frequentist regret induced by the randomness of the sampling process in LQ systems is considerably more difficult and it requires developing a new line of proof that directly relates the regret of TS and the controls executed over time. Furthermore, we show that TS has to solve a trade-off between frequently updating the policy to guarantee enough optimistic samples and reducing the number of policy switches to limit the regret incurred at each change. This gives rise to a final bound of $O(T^{2/3})$. This opens a number of questions. **1)** The current analysis is derived in the general n/d -dimensional case except for Lem. 3 and the steps leading to the introduction of the state in Sect. 4.4, where we set $n = d = 1$. We believe that these steps can be extended to the general case without affecting the final result. **2)** The final regret bound is in striking contrast with previous results for TS. While we provide a rather intuitive reason on the source of this extra regret, it is an open question whether a different TS or analysis could allow to improve the regret to $O(\sqrt{T})$ or whether this result reveals an intrinsic limitation of the randomized approach of TS.

Acknowledgement This research is supported in part by a grant from CPER Nord-Pas de Calais/FEDER DATA Advanced data science and technologies 2015-2020, CRISAL (Centre de Recherche en Informatique et Automatique de Lille), and the French National Research Agency (ANR) under project ExTra-Learn n.ANR-14-CE24-0010-01.

References

- [1] Yasin Abbasi-Yadkori and Csaba Szepesvári. Regret bounds for the adaptive control of linear quadratic systems. In *COLT*, pages 1–26, 2011.
- [2] Yasin Abbasi-Yadkori and Csaba Szepesvári. Bayesian optimal control of smoothly parameterized systems. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 2015.
- [3] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Proceedings of the 25th Annual Conference on Neural Information Processing Systems (NIPS)*, 2011.
- [4] Gabriel Acosta and Ricardo G Durán. An optimal poincaré inequality in l_1 for convex domains. *Proceedings of the american mathematical society*, pages 195–202, 2004.
- [5] Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. *arXiv preprint arXiv:1209.3352*, 2012.
- [6] S Bittanti and MC Campi. Adaptive control of linear time invariant systems: the “bet on the best” principle. *Communications in Information & Systems*, 6(4):299–320, 2006.
- [7] Marco C Campi and PR Kumar. Adaptive linear quadratic gaussian control: the cost-biased approach revisited. *SIAM Journal on Control and Optimization*, 36(6):1890–1907, 1998.
- [8] Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *J. Mach. Learn. Res.*, 11:1563–1600, August 2010.
- [9] P. Lancaster and L. Rodman. *Algebraic riccati equations*. Oxford University Press, 1995.
- [10] Ian Osband and Benjamin Van Roy. Deep exploration via bootstrapped dqn. *Proceedings of the 30th Annual Conference on Neural Information Processing Systems (NIPS)*, 2016.
- [11] Ian Osband, Benjamin Van Roy, and Zheng Wen. Generalization and exploration via randomized value functions. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pages 2377–2386, 2016.
- [12] Lawrence E Payne and Hans F Weinberger. An optimal poincaré inequality for convex domains. *Archive for Rational Mechanics and Analysis*, 5(1):286–292, 1960.
- [13] Alexander L. Strehl, Lihong Li, and Michael L. Littman. Reinforcement learning in finite mdps: Pac analysis. *J. Mach. Learn. Res.*, 10:2413–2444, December 2009.
- [14] Malcolm J. A. Strens. A bayesian framework for reinforcement learning. In *Proceedings of the Seventeenth International Conference on Machine Learning*, pages 943–950, 2000.

A Control theory

A.1 Proof of Prop. 2

1. When $\theta^\top = (A, B)$ is not stabilizable, there exists no linear control K such that the controlled process $x_{t+1} = Ax_t + BKx_t + \epsilon_{t+1}$ is stationary. Thus, the positiveness of Q and R implies $J(\theta) = \text{Tr}(P(\theta)) = +\infty$. As a consequence, $\theta^\top \notin \mathcal{S}$.
2. The mapping $\theta \rightarrow \text{Tr}(P(\theta))$ is continuous (see Lem. 1). Thus, \mathcal{S} is compact as the intersection between a closed and a compact set.
3. The continuity of the mapping $\theta \rightarrow K(\theta)$ together with the compactness of \mathcal{S} justifies the finite positive constants ρ and C . Moreover, since every $\theta \in \mathcal{S}$ are stabilizable pairs, $\rho < 1$.

A.2 Proof of Lem. 1

Let $\theta^\top = (A, B)$ where A and B are matrices of size $n \times n$ and $n \times d$ respectively. Let $\mathcal{R} : \mathbb{R}^{n+d,n} \times \mathbb{R}^{n,n} \rightarrow \mathbb{R}^{n,n}$ be the Riccati operator defined by:

$$\mathcal{R}(\theta, P) := Q - P + A^\top P A - A^\top P B (R + B^\top P B)^{-1} B^\top P A, \quad (12)$$

where Q, R are positive definite matrices. Then, the solution $P(\theta)$ of the Riccati equation of Thm. 1 is the solution of $\mathcal{R}(\theta, P) = 0$. While Prop. 2 guarantees that there exists a unique admissible solution as soon as $\theta \in \mathcal{S}$, addressing the regularity of the function $\theta \rightarrow P(\theta)$ requires the use of the implicit function theorem.

Theorem 2 (Implicit function theorem). *Let E and F be two banach spaces, let $\Omega \subset E \times F$ be an open subset. Let $f : \Omega \rightarrow F$ be a C^1 -map and let (x_0, y_0) be a point of Ω such that $f(x_0, y_0) = 0$. We denote as $d_y f(x_0, y_0) : F \rightarrow F$ the differential of the function f with respect to the second argument at point (x_0, y_0) . Assume that this linear transformation is bounded and invertible. Then, there exists*

1. two open subsets U and V such that $(x_0, y_0) \in U \times V \subset \Omega$,
2. a function $g : U \rightarrow V$ such that $g(x) = y$ for all $(x, y) \in U \times V$.

Moreover, g is C^1 and $dg(x) = -d_y f(x, g(x))^{-1} d_x f(x, g(x))$ for all $(x, y) \in U \times V$.

Since R is positive definite, the Riccati operator is clearly a C^1 -map. Moreover, thanks to Thm. 1, to any $\theta \in \mathcal{S}$, there exists an admissible P such that $\mathcal{R}(\theta, P) = 0$. Thanks to Thm. 2, a sufficient condition for $\theta \rightarrow P(\theta)$ to be C^1 on \mathcal{S} is that the linear map $d_P \mathcal{R}(\theta, P(\theta)) : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}$ is a bounded invertible transformation i.e.

- **Bounded.** There exists M such that, for any $P \in \mathbb{R}^{n \times n}$, $\|d_P \mathcal{R}(\theta, P(\theta))(P)\| \leq M \|P\|$.
- **Invertible.** There exists a bounded linear operator $S : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}$ such that $SP = I_{n,n}$ and $PS = I_{n,n}$.

Lemma 5. *Let $\theta^\top = (A, B)$ and \mathcal{R} be the Riccati operator defined in equation (12). Then, the differential of \mathcal{R} w.r.t P taken in $(\theta, P(\theta))$ denoted as $d_P \mathcal{R}(\theta, P(\theta))$ is defined by:*

$$d_P \mathcal{R}(\theta, P(\theta))(\delta P) := A_c^\top \delta P A_c - \delta P, \quad \text{for any } \delta P \in \mathbb{R}^{n \times n},$$

where $A_c = A - B(R + B^\top P B)^{-1} B^\top P A$.

Proof. The proof is straightforward using the standard composition/multiplication/inverse operations for the differential operator together with an appropriate rearranging. \square

Clearly, $d_P \mathcal{R}(\theta, P(\theta))$ is a bounded linear map. Moreover, thanks to the Lyapunov theory, for any stable matrix $\|A_c\|_2 < 1$ and for any matrix Q , the Lyapunov equation $A_c^\top X A_c - X = Q$ admits a unique solution. From Thm. 1, the optimal matrix $P(\theta)$ is such that the corresponding A_c is stable. This implies that $d_P \mathcal{R}(\theta, P(\theta))$ is an invertible operator, and $\theta \rightarrow P(\theta)$ is C^1 on \mathcal{S} .

Therefore, the differential of $\theta \rightarrow P(\theta)$ can be deduced from the implicit function theorem. After tedious yet standard operations, one gets that for any $\theta \in \mathcal{S}$ and direction $\delta\theta \in \mathbb{R}^{(n+d) \times n}$:

$$dJ(\theta)(\delta\theta) = \text{Tr}(dP(\theta)(\delta\theta)) = \text{Tr}(\nabla J(\theta)^\top \delta\theta),$$

where $\nabla J(\theta) \in \mathbb{R}^{(n+d) \times n}$ is the jacobian matrix of J in θ . For any $\delta\theta \in \mathbb{R}^{(n+d) \times n}$, one has:

$$\nabla J(\theta)^\top \delta\theta = A_c(\theta)^\top \nabla J(\theta)^\top \delta\theta A_c(\theta) + C(\theta, \delta\theta) + C(\theta, \delta\theta)^\top, \quad \text{where } C(\theta, \delta\theta) = A_c(\theta)^\top P(\theta) \delta\theta^\top H(\theta). \quad (13)$$

Proposition 5. *For any $\theta \in \mathcal{S}$ and any positive definite matrix V , one has the following inequality for the weighted norm of the gradient of J :*

$$\|\nabla J(\theta)\|_V \leq \|A_c(\theta)\|_2^2 \|\nabla J(\theta)\|_V + 2\|P(\theta)\| \|A_c(\theta)\|_2 \|H(\theta)\|_V.$$

Proof. For any $\theta \in \mathcal{S}$ and any positive definite matrix $V \in \mathbb{R}^{(n+d) \times (n+d)}$. Applying (13) to $\delta\theta = V\nabla J(\theta)$ leads to:

$$\nabla J(\theta)^\top V \nabla J(\theta) = A_c(\theta)^\top \nabla J(\theta)^\top V \nabla J(\theta) A_c(\theta) + C(\theta, V\nabla J(\theta)) + C(\theta, V\nabla J(\theta))^\top,$$

where $C(\theta, V\nabla J(\theta))^\top = (V^{1/2}H(\theta))^\top V^{1/2} \nabla J(\theta) P(\theta) A_c(\theta)$. Let $\langle A, B \rangle = \text{Tr}A^\top B$ be the Frobenius inner product, then taking the trace of the above equality, one gets:

$$\|\nabla J(\theta)\|_V^2 = \|\nabla J(\theta) A_c(\theta)\|_V^2 + 2\langle V^{1/2}H(\theta), V^{1/2} \nabla J(\theta) P(\theta) A_c(\theta) \rangle.$$

Using the Cauchy-Schwarz inequality and that the Frobenius norm is sub-multiplicative together with $\text{Tr}(M_1 M_2) \leq \|M_1\|_2 \text{Tr}(M_2)$ for any M_1, M_2 symmetric positive definite matrices, one obtains:

$$\|\nabla J(\theta)\|_V^2 \leq \|A_c(\theta)\|_2^2 \|\nabla J(\theta)\|_V^2 + 2\|H(\theta)\|_V \|P(\theta)\| \|A_c(\theta)\|_2 \|\nabla J(\theta)\|_V.$$

Finally, dividing by $\|\nabla J(\theta)\|_V$ provides the desired result. \square

B Material

Theorem 3 (Azuma's inequality). *Let $\{M_s\}_{s \geq 0}$ be a super-martingale such that $|M_s - M_{s-1}| \leq c_s$ almost surely. Then, for all $t > 0$ and all $\epsilon > 0$,*

$$\mathbb{P}(|M_t - M_0| \geq \epsilon) \leq 2 \exp\left(\frac{-\epsilon^2}{2 \sum_{s=1}^t c_s^2}\right).$$

Lemma 6 (Lemma. 8 from Abbasi-Yadkori and Szepesvári [1]). *Let K^{det} be the number of changes in the policy of Algorithm 1 due to the determinant trigger $\det(V_t) \geq 2 \det(V_0)$. Then, on E , K^{det} is at most*

$$K^{\text{den}} \leq (n+d) \log_2(1 + TX^2(1 + C^2)/\lambda).$$

Corollary 2. *Let K be the number of policy changes of Algorithm 1, K^{det} be defined as in Lem. 6 and $K^{\text{len}} = K - K^{\text{det}}$ be the number of policy changes due to the length trigger $t \geq t_0 + \tau$. Then, on E , K is at most*

$$K \leq K^{\text{det}} + K^{\text{len}} \leq (n+d) \log_2(1 + TX^2(1 + C^2)/\lambda) + T/\tau.$$

Moreover, assuming that $T \geq \frac{\lambda}{X^2(1+C^2)}$, one gets $K \leq (n+d) \log_2(1 + TX^2(1 + C^2)/\lambda)T/\tau$.

Lemma 7 (Chernoff bound for Gaussian r.v.). *Let $X \sim \mathcal{N}(0, 1)$. For any $0 < \delta < 1$, for any $t \geq 0$, then,*

$$\mathbb{P}(|X| \geq t) \leq 2 \exp\left(-\frac{t^2}{2}\right).$$

Proof of Lem.2. Let $\delta' = \delta/8T$.

1. From Prop. 3, $\mathbb{P}(\|\widehat{\theta}_t - \theta_*\|_{V_t} \leq \beta_t(\delta')) \geq 1 - \delta'$. Hence,

$$\begin{aligned} \mathbb{P}(\widehat{E}) &= \mathbb{P}\left(\bigcap_{t=0}^T (\|\widehat{\theta}_t - \theta_*\|_{V_t} \leq \beta_t(\delta'))\right) \\ &= 1 - \mathbb{P}\left(\bigcup_{t=0}^T (\|\widehat{\theta}_t - \theta_*\|_{V_t} \geq \beta_t(\delta'))\right) \\ &\geq 1 - \sum_{t=0}^T \mathbb{P}(\|\widehat{\theta}_t - \theta_*\|_{V_t} \geq \beta_t(\delta')) \\ &\geq 1 - T\delta' \geq 1 - \delta/8 \end{aligned}$$

2. From Lem. 7, let $\eta \sim \mathcal{D}^{\text{TS}}$ then, for any $\epsilon > 0$, making use of the fact that $\|\eta\| \leq n\sqrt{n+d} \max_{i \leq n+d, j \leq n} |\eta_{i,j}|$,

$$\mathbb{P}(\|\eta\| \leq \epsilon) \geq \mathbb{P}(n\sqrt{n+d} \max_{i,j} |\eta_{i,j}| \leq \epsilon) \geq 1 - \prod_{i,j} \mathbb{P}(|\eta_{i,j}| \geq \frac{\epsilon}{n\sqrt{n+d}}) \geq 1 - n(n+d) \mathbb{P}_{X \sim \mathcal{N}(0,1)}(|X| \geq \frac{\epsilon}{n\sqrt{n+d}}).$$

Hence,

$$\begin{aligned} \mathbb{P}(\widetilde{E}) &= \mathbb{P}\left(\bigcap_{t=0}^T (\|\widetilde{\theta}_t - \widehat{\theta}_t\|_{V_t} \leq \gamma_t(\delta'))\right) = 1 - \mathbb{P}\left(\bigcup_{t=0}^T (\|\widetilde{\theta}_t - \widehat{\theta}_t\|_{V_t} \geq \gamma_t(\delta'))\right) \\ &\geq 1 - \sum_{t=0}^T \mathbb{P}(\|\widetilde{\theta}_t - \widehat{\theta}_t\|_{V_t} \geq \gamma_t(\delta')) \geq 1 - \sum_{t=0}^T \mathbb{P}(\|\eta\| \geq \gamma_t(\delta')/\beta_t(\delta')) \\ &\geq 1 - \sum_{t=0}^T \mathbb{P}\left(\|\eta\| \geq n\sqrt{2(n+d) \log(2n(n+d)/\delta')}\right) \\ &\geq 1 - T\delta' \geq 1 - \delta/8. \end{aligned}$$

3. Finally, a union bound argument ensures that $\mathbb{P}(\widehat{E} \cap \widetilde{E}) \geq 1 - \delta/4$.

Proof of Cor. 1. This result comes directly from Sec. 4.1. and App. D of Abbasi-Yadkori and Szepesvári [1]. The proof relies on the fact that, on \widehat{E} , because $\widehat{\theta}_t$ is chosen within the confidence ellipsoid $\mathcal{E}_t^{\text{RLS}}$, the number of time steps the true closed loop matrix $A_* + B_*K(\widehat{\theta}_t)$ is unstable is small. Intuitively, the reason is that as soon as the true closed loop matrix is unstable, the state process explodes and the confidence ellipsoid is drastically changed. As the ellipsoid can only shrink over time, the state is well controlled expect for a small number of time steps.

Since the only difference is that, on $\widehat{E} \cap \widetilde{E}$, $\widetilde{\theta}_t \in \mathcal{E}_t^{\text{TS}}$, the same argument applies and the same bound holds replacing β_t with γ_t . Therefore, there exists appropriate problem dependent constants X, X' such that $\mathbb{P}(\widehat{E} | \widehat{E} \cap \widetilde{E}) \geq 1 - \delta/4$. Finally, a union bound argument ensures that $\mathbb{P}(\widehat{E} \cap \widetilde{E} \cap \bar{E}) \geq 1 - \delta/2$.

C Proof of Lem. 3

We prove here that, on E , the sampling $\widetilde{\theta} \sim \mathcal{R}_S(\widehat{\theta}_t + \beta_t(\delta')V_t^{1/2})$ guarantees a fixed probability of sampling an optimistic parameter, i.e. which belongs to $\Theta_t^{\text{opt}} := \{\theta \in \mathbb{R}^d \mid J(\theta) \leq J(\theta^*)\}$. However, our result only holds for the 1-dimensional case as we deeply leverage on the geometry of the problem. Figure 2 synthesizes the properties of the optimal value function and the geometry of the problem w.r.t the probability of being optimistic.

1. First, we introduce a simpler subset of optimistic parameters which involves hyperplanes rather than complicated J level sets. Without loss of generality we assume that $A_* + B_*K_* = \rho_* \geq 0$ and introduce $H_* = \begin{pmatrix} 1 \\ K_* \end{pmatrix} \in \mathbb{R}^2$ so that $A_* + B_*K_* = \theta^\top H_*$. Let $\Theta^{\text{lin,opt}} = \{\theta \in \mathbb{R}^d \mid |\theta^\top H_*| \leq \rho_*\}$. Intuitively, $\Theta^{\text{lin,opt}}$ consists in the set of systems θ which are more stable under control K_* . The following proposition ensures those systems to be optimistic.




Figure 2: **Optimism and worst case configuration.** **1)** In 1-D, the Riccati solution is well-defined expect for $\{(A, B) \in]-\infty, -1] \cup [1, \infty[\times\{0\}\}$. The rejection sampling procedure into \mathcal{S} ensures $P(\hat{\theta}_t)$ to be well-defined. Moreover, \mathcal{S}^c does not overlap with Θ^{opt} . **2)** The introduction of the subset $\Theta^{\text{lin,opt}}$ prevents using the actual - yet complicated - optimistic set Θ^{opt} to lower bound the probability of being optimistic. **3)** Even if the event \mathcal{E}^{RLS} holds, there exists an ellipsoid configuration which does not contain any optimistic point. This justifies the over-sampling to guarantee a fixed probability of being optimistic.

Proposition 6. $\Theta^{\text{lin,opt}} \subset \Theta_t^{\text{opt}}$.

Proof. Leveraging on the expression of J , one has when $n = d = 1$,

$$J(\theta) = \text{Tr}(P(\theta)) = P(\theta) = \lim_{T \rightarrow \infty} \sum_{t=0}^T x_t^2 (Q + K(\theta)^2 R) = (Q + K(\theta)^2 R) \mathbb{V}(x_t),$$

where $\mathbb{V}(x_t) = (1 - |\theta^\top H(\theta)|^2)^{-1}$ is the steady-state variance of the stationary first order autoregressive process $x_{t+1} = \theta^\top H(\theta) x_t + \epsilon_{t+1}$ where ϵ_t is zero mean noise of variance 1 and $H(\theta) = \begin{pmatrix} 1 \\ K(\theta) \end{pmatrix}$. Thus,

$$J(\theta) = (Q + K(\theta)^2 R) (1 - |\theta^\top H(\theta)|^2)^{-1}.$$

Hence, for any $\theta \in \Theta^{\text{lin,opt}}$, $(1 - |\theta^\top H_*|^2)^{-1} \leq (1 - |\theta_*^\top H_*|^2)^{-1}$ which implies that

$$(Q + K_*^2 R) (1 - |\theta^\top H_*|^2)^{-1} \leq (Q + K_*^2 R) (1 - |\theta_*^\top H_*|^2)^{-1} = J(\theta_*).$$

However, since $K(\theta)$ is the optimal control associated with θ ,

$$\begin{aligned} J(\theta) &= (Q + K(\theta)^2 R) (1 - |\theta^\top H(\theta)|^2)^{-1} \\ &= \min_K (Q + K^2 R) (1 - |(1 \quad K) \theta|^2)^{-1} \\ &\leq (Q + K_*^2 R) (1 - |\theta^\top H_*|^2)^{-1} \\ &\leq J(\theta_*) \end{aligned}$$

□

As a result, $\mathbb{P}(\tilde{\theta}_t \in \Theta^{\text{opt}} \mid \mathcal{F}_t^x, \hat{E}_t) \geq \mathbb{P}(\tilde{\theta}_t \in \Theta^{\text{lin,opt}} \mid \mathcal{F}_t^x, \hat{E}_t)$ and we can focus on $\Theta^{\text{lin,opt}}$.

2. To ensure the sampling parameter to be admissible, we perform a rejection sampling until $\tilde{\theta}_t \in \mathcal{S}$. Noticing that $\Theta^{lin,opt} \subset \Theta^{opt} \subset \mathcal{S}$ by construction, the rejection sampling is always favorable in terms of probability of being optimistic. Since we seek for a lower bound, we can get rid of it and consider $\tilde{\theta}_t = \hat{\theta}_t + \beta_t(\delta')V_t^{-1/2}\eta$ where $\eta \sim \mathcal{N}(0, I_2)$.⁹
3. On \hat{E}_t , $\theta_* \in \mathcal{E}_t^{RLS}$, where \mathcal{E}_t^{RLS} is the confidence RLS ellipsoid centered in $\hat{\theta}_t$. Since θ_* is fixed (by definition), we lower bound the probability by considering the worst possible $\hat{\theta}_t$ such that \hat{E}_t holds. Intuitively, we consider the worst possible center for the RLS ellipsoid such that θ_* still belong in \mathcal{E}_t^{RLS} and that the probability of being optimistic is minimal. Formally,

$$\begin{aligned} \mathbb{P}(\tilde{\theta}_t \in \Theta^{lin,opt} \mid \mathcal{F}_t^x, \hat{E}_t) &= \mathbb{P}_{\tilde{\theta}_t \sim \mathcal{N}(\hat{\theta}_t, \beta_t^2(\delta')V_t^{-1})}(\tilde{\theta}_t \in \Theta^{lin,opt} \mid \mathcal{F}_t^x, \hat{E}_t) \\ &\geq \min_{\hat{\theta}: \|\hat{\theta} - \theta_*\|_{V_t} \leq \beta_t(\delta')} \mathbb{P}_{\tilde{\theta}_t \sim \mathcal{N}(\hat{\theta}, \beta_t^2(\delta')V_t^{-1})}(\tilde{\theta}_t \in \Theta^{lin,opt} \mid \mathcal{F}_t^x) \end{aligned}$$

Moreover, by Cauchy-Schwarz inequality, for any $\hat{\theta}$,

$$|(\hat{\theta} - \theta_*)^\top H_*| \leq \|\hat{\theta} - \theta_*\|_{V_t} \|H_*\|_{V_t^{-1}} \leq \beta_t(\delta') \|H_*\|_{V_t^{-1}},$$

thus,

$$\begin{aligned} \mathbb{P}(\tilde{\theta}_t \in \Theta^{lin,opt} \mid \mathcal{F}_t^x, \hat{E}_t) &\geq \min_{\hat{\theta}: \|\hat{\theta} - \theta_*\|_{V_t} \leq \beta_t(\delta')} \mathbb{P}_{\tilde{\theta}_t \sim \mathcal{N}(\hat{\theta}, \beta_t^2(\delta')V_t^{-1})}(\tilde{\theta}_t \in \Theta^{lin,opt} \mid \mathcal{F}_t^x) \\ &\geq \min_{\hat{\theta}: |(\hat{\theta} - \theta_*)^\top H_*| \leq \beta_t(\delta') \|H_*\|_{V_t^{-1}}} \mathbb{P}_{\tilde{\theta}_t \sim \mathcal{N}(\hat{\theta}, \beta_t^2(\delta')V_t^{-1})}(\tilde{\theta}_t \in \Theta^{lin,opt} \mid \mathcal{F}_t^x) \quad (14) \\ &= \min_{\hat{\theta}: |\hat{\theta}^\top H_* - \rho_*| \leq \beta_t(\delta') \|H_*\|_{V_t^{-1}}} \mathbb{P}_{\tilde{\theta}_t \sim \mathcal{N}(\hat{\theta}, \beta_t^2(\delta')V_t^{-1})}(|\tilde{\theta}_t^\top H_*| \leq \rho_* \mid \mathcal{F}_t^x) \end{aligned}$$

Cor. 3 provides us with an explicit expression of the worst case ellipsoid. Introducing $x = \hat{\theta}_t^\top H_*$, one has $x \sim \mathcal{N}(\bar{x}, \sigma_x^2)$ with $\bar{x} = \hat{\theta}_t^\top H_*$ and $\sigma_x = \beta_t(\delta') \|H_*\|_{V_t^{-1}}$. Applying Cor. 3 with $\alpha = \rho_*$, $\rho = \rho_*$ and $\beta = \beta_t(\delta') \|H_*\|_{V_t^{-1}}$, inequality (14) becomes

$$\begin{aligned} \mathbb{P}(\tilde{\theta}_t \in \Theta^{lin,opt} \mid \mathcal{F}_t^x, \hat{E}_t) &\geq \min_{\hat{\theta}: |\hat{\theta}^\top H_* - \rho_*| \leq \beta_t(\delta') \|H_*\|_{V_t^{-1}}} \mathbb{P}_{\eta \sim \mathcal{N}(0, I_2)}(|\hat{\theta}^\top H_* + \beta_t(\delta') \eta^\top V_t^{-1/2} H_*| \leq \rho_* \mid \mathcal{F}_t^x) \\ &\geq \mathbb{P}_{\eta \sim \mathcal{N}(0, I_2)}(|\rho_* + \beta_t(\delta') \|H_*\|_{V_t^{-1}} + \beta_t(\delta') \eta^\top V_t^{-1/2} H_*| \leq \rho_* \mid \mathcal{F}_t^x) \end{aligned}$$

Introducing the vector $u_t = \beta_t(\delta') V_t^{-1/2} H_*$, one can simplify

$$\begin{aligned} &|\rho_* + \beta_t(\delta') \|H_*\|_{V_t^{-1}} + \beta_t(\delta') \eta^\top V_t^{-1/2} H_*| \leq \rho_*, \\ &\Leftrightarrow -\rho_* \leq \rho_* + \|u_t\| + \eta^\top u_t \leq \rho_*, \\ &\Leftrightarrow -\frac{\rho_*}{\|u_t\|} - 1 \leq \eta^\top \frac{u_t}{\|u_t\|} \leq -1. \end{aligned}$$

Since $\eta \sim \mathcal{N}(0, I_2)$ is rotationally invariant, $\mathbb{P}(\tilde{\theta}_t \in \Theta^{lin,opt} \mid \mathcal{F}_t^x, \hat{E}_t) \geq \mathbb{P}_{\epsilon \sim \mathcal{N}(0,1)}(\epsilon \in [1, 1 + \frac{2\rho_*}{\|u_t\|}] \mid \mathcal{F}_t^x, \hat{E}_t)$. Finally, for all $t \leq T$, u_t is almost surely bounded: $\|u_t\| \leq \beta_T(\delta') \sqrt{(1+C^2)/\lambda}$. Therefore,

$$\mathbb{P}(\tilde{\theta}_t \in \Theta^{lin,opt} \mid \mathcal{F}_t^x, \hat{E}_t) \geq \mathbb{P}_{\epsilon \sim \mathcal{N}(0,1)}(\epsilon \in [1, 1 + 2\rho_*/\beta_T(\delta') \sqrt{(1+C^2)/\lambda}]) := p$$

Corollary 3. For any $\rho, \sigma_x > 0$, for any $\alpha, \beta \geq 0$, $\arg \min_{\bar{x}: |\bar{x} - \alpha| \leq \beta} \mathbb{P}_{x \sim \mathcal{N}(\bar{x}, \sigma_x^2)}(|x| \leq \rho) = \alpha + \beta$.

This corollary is a direct consequence of the properties of standard gaussian r.v.

Lemma 8. Let x be a real random variable. For any $\rho, \sigma_x > 0$ Let $f: \mathbb{R} \rightarrow [0, 1]$ be the continuous mapping defined by $f(\bar{x}) = \mathbb{P}_{x \sim \mathcal{N}(\bar{x}, \sigma_x^2)}(|x| \leq \rho)$. Then, f is increasing on \mathbb{R}_- and decreasing on \mathbb{R}_+ .

⁹In the 1-dimensional case, η is just a 2d standard gaussian r.v.

Proof. Without loss of generality, one can assume that $\sigma_x = 1/\sqrt{2}$ (otherwise, modify ρ), and that $\bar{x} \geq 0$ (by symmetry). Denoting as Φ and erf the standard gaussian cdf and the error function, one has:

$$\begin{aligned} f(\bar{x}) &= \mathbb{P}_{x \sim \mathcal{N}(\bar{x}, \sigma_x^2)}(-\rho \leq x \leq \rho), = \mathbb{P}_{x \sim \mathcal{N}(\bar{x}, \sigma_x^2)}(x \leq \rho) - \mathbb{P}_{x \sim \mathcal{N}(\bar{x}, \sigma_x^2)}(x \leq -\rho), \\ &= \mathbb{P}_{x \sim \mathcal{N}(\bar{x}, \sigma_x^2)}((x - \bar{x})/\sigma_x \leq (\rho - \bar{x})/\sigma_x) - \mathbb{P}_{x \sim \mathcal{N}(\bar{x}, \sigma_x^2)}((x - \bar{x})/\sigma_x \leq (-\rho - \bar{x})/\sigma_x), \\ &= \Phi((\rho - \bar{x})/\sigma_x) - \Phi(-(\rho + \bar{x})/\sigma_x), \\ &= \frac{1}{2} + \frac{1}{2}\text{erf}((\rho - \bar{x})/\sqrt{2}\sigma_x) - \frac{1}{2} - \frac{1}{2}\text{erf}(-(\rho + \bar{x})/\sqrt{2}\sigma_x), \\ &= \frac{1}{2}(\text{erf}(\rho - \bar{x}) - \text{erf}(-(\rho + \bar{x}))). \end{aligned}$$

Since erf is odd, one obtains $f(\bar{x}) = \frac{1}{2}(\text{erf}(\rho - \bar{x}) + \text{erf}(\rho + \bar{x}))$. The error function is differentiable with $\text{erf}'(z) = \frac{2}{\pi}e^{-z^2}$, thus

$$\begin{aligned} f'(\bar{x}) &= \frac{1}{\pi} \left(\exp(-(\rho + \bar{x})^2) - \exp(-(\rho - \bar{x})^2) \right) \\ &= -\frac{2}{\pi} \sinh((\rho - \bar{x})^2) \leq 0 \end{aligned}$$

Hence, f is decreasing on \mathbb{R}_+ and by symmetry, is increasing on \mathbb{R}_- . □

D Weighted L1 Poincaré inequality (proof of Lem. 4)

This result is build upon the following theorem which links the function to its gradient in L^1 norm:

Theorem 4 (see Acosta and Durán [4]). *Let $W^{1,1}(\Omega)$ be the Sobolev space on $\Omega \subset \mathbb{R}^d$. Let Ω be a convex domain bounded with diameter D and $f \in W^{1,1}(\Omega)$ of zero average on Ω then*

$$\int_{\Omega} |f(x)| dx \leq \frac{D}{2} \int_{\Omega} \|\nabla f(x)\| dx \tag{15}$$

Lem. 4 is an extension of Thm. 4. In practice, we show that their proof still holds for log-concave weight.

Theorem 5. *Let $L > 0$ and ρ any non negative and log-concave function on $[0, L]$. Then for any $f \in W^{1,1}(0, L)$ such that*

$$\int_0^L f(x)\rho(x)dx = 0$$

one has:

$$\int_0^L |f(x)|\rho(x)dx \leq 2L \int_0^L |f'(x)|\rho(x)dx \tag{16}$$

The proof is based on the following inequality for log-concave function.

Lemma 9. *Let ρ be any non negative log-concave function on $[0, 1]$ such that $\int_0^1 \rho(x) = 1$ then*

$$\forall x \in (0, 1), \quad H(\rho, x) := \frac{1}{\rho(x)} \int_0^x \rho(t)dt \int_x^1 \rho(t)dt \leq 1 \tag{17}$$

Proof. Since any non-negative log-concave function on $[0, 1]$ can be rewritten as $\rho(x) = e^{\nu(x)}$ where ν is a concave function on $[0, 1]$ and since $x \rightarrow e^x$ is increasing, the monotonicity of ν is preserved and as for concave function, ρ can be either increasing, decreasing or increasing then decreasing on $[0, 1]$.

Hence, $\forall x \in (0, 1)$, either

1. $\rho(t) \leq \rho(x)$ for all $t \in [0, x]$,
2. $\rho(t) \leq \rho(x)$ for all $t \in [x, 1]$.

Assume that $\rho(t) \leq \rho(x)$ for all $t \in [0, x]$ without loss of generality. Then,

$$\begin{aligned} \forall x \in (0, 1), \quad H(\rho, x) &:= \frac{1}{\rho(x)} \int_0^x \rho(t) dt \int_x^1 \rho(t) dt \\ &= \int_0^x \frac{\rho(t)}{\rho(x)} \int_x^1 \rho(t) dt \\ &\leq \int_0^x dt \int_x^1 \rho(t) dt \\ &\leq x \int_0^1 \rho(t) dt \leq x \leq 1 \end{aligned}$$

□

Proof of theorem 5. This proof is exactly the same as [4] where we use lemma 9 instead of a concave inequality. We provide it for sake of completeness.

A scaling argument ensures that it is enough to prove it for $L = 1$. Moreover, dividing both side of (16) by $\int_0^1 \rho(x) dx$, we can assume without loss of generality that $\int_0^1 \rho(x) dx = 1$.

Since $\int_0^1 f(x) \rho(x) dx = 0$ by integration part by part one has:

$$\begin{aligned} f(y) &= \int_0^y f'(x) \int_0^x \rho(t) dt - \int_y^1 f'(x) \int_x^1 \rho(t) dt \\ |f(y)| &\leq \int_0^y |f'(x)| \int_0^x \rho(t) dt + \int_y^1 |f'(x)| \int_x^1 \rho(t) dt \end{aligned}$$

Multiplying by $\rho(y)$, integrating on y and applying Fubini's theorem leads to

$$\int_0^1 |f(y)| \rho(y) dy \leq 2 \int_0^1 |f'(x)| \int_0^x \rho(t) dt \int_x^1 \rho(t) dt$$

and applying (17) of lemma 9 ends the proof. □

While theorem 5 provides a 1 dimensional weighed Poincaré inequality, we actually seek for one in \mathbb{R}^d . The idea of [4] is to use arguments of [12] to reduce the d -dimensional problem to a $1-d$ problem by splitting any convex set Ω into subspaces Ω_i thin in all but one direction and such that an average property is preserved. We just provide their result.

Lemma 10. *Let $\Omega \subset \mathbb{R}^d$ be a convex domain with finite diameter D and $u \in L^1(\Omega)$ such that $\int_{\Omega} u = 0$. Then, for any $\delta > 0$, there exists a decomposition of Ω into a finite number of convex domains Ω_i satisfying*

$$\Omega_i \cap \Omega_j = \emptyset \text{ for } i \neq j, \quad \bar{\Omega} = \bigcup \bar{\Omega}_i, \quad \int_{\Omega_i} u = 0$$

and each Ω_i is thin in all but one direction i.e. in an appropriate rectangular coordinate system $(x, y) = (x, y_1, \dots, y_{d-1})$ the set Ω_i is contained in

$$\{(x, y) : 0 \leq x \leq D, \quad 0 \leq y_i \leq \delta \text{ for } i = 1, \dots, d-1\}$$

This decomposition together with theorem 5 allow us to prove the d -dimensional weighted Poincaré inequality.

Proof of Lem. 4. By density, we can assume that $u \in C^\infty(\bar{\Omega})$. Hence, $up \in C^2(\bar{\Omega})$. Let M be a bound for up and all its derivative up to the second order.

Given $\delta > 0$ decompose the set Ω into Ω_i as in lemma 10 and express $z \in \Omega_i$ into the appropriate rectangular basis $z = (x, y)$, where $x \in [0, d_i]$, $y \in [0, \delta]$. Define as $\rho(x_0)$ the $d-1$ volume of the intersection between Ω_i and the hyperplan $\{x = x_0\}$. Since Ω_i is convex, ρ is concave and from the smoothness of up one has:

$$\left| \int_{\Omega_i} |u(x, y)|p(x, y)dxdy - \int_0^{d_i} |u(x, 0)|p(x, 0)\rho(x)dx \right| \leq (d-1)M|\Omega_i|\delta \quad (18)$$

$$\left| \int_{\Omega_i} \left| \frac{\partial u}{\partial x}(x, y) \right| p(x, y) dxdy - \int_0^{d_i} \left| \frac{\partial u}{\partial x}(x, 0) \right| p(x, 0) \rho(x) dx \right| \leq (d-1)M|\Omega_i|\delta \quad (19)$$

$$\left| \int_{\Omega_i} u(x, y)p(x, y)dxdy - \int_0^{d_i} u(x, 0)p(x, 0)\rho(x)dx \right| \leq (d-1)M|\Omega_i|\delta \quad (20)$$

Those equation allows us to switch from d -dimensional integral to 1-dimensional integral for which we can apply theorem 5 at the condition that $\int_0^{d_i} u(x, 0)p(x, 0)\rho(x)dx = 0$ (which is not satisfied here). On the other hand, we can apply theorem 5 to

$$g(x) = u(x, 0) - \int_0^{d_i} u(x, 0)p(x, 0)\rho(x)dx / \int_0^{d_i} p(x, 0)\rho(x)dx$$

with weighed function $x \rightarrow p(x, 0)\rho(x)$. Indeed, $x \rightarrow p(x, 0)$ is log-concave - as restriction along one direction of log-concave function, $x \rightarrow \rho(x)$ is log-concave - as a concave function, and so is $x \rightarrow p(x, 0)\rho(x)$ - as product of log-concave function. Moreover, $g \in W^{1,1}(0, d_i)$ and $\int_0^{d_i} g(x)p(x, 0)\rho(x)dx = 0$ by construction. Therefore, applying theorem 5 one gets:

$$\begin{aligned} \int_0^{d_i} |g(x)|p(x, 0)\rho(x)dx &\leq 2d_i \int_0^{d_i} |g'(x)|p(x, 0)\rho(x)dx \\ \int_0^{d_i} |u(x, 0)|p(x, 0)\rho(x)dx &\leq 2d_i \int_0^{d_i} \left| \frac{\partial u}{\partial x}(x, 0) \right| p(x, 0) \rho(x) dx - \left| \int_0^{d_i} u(x, 0)p(x, 0)\rho(x)dx \right| \\ \int_0^{d_i} |u(x, 0)|p(x, 0)\rho(x)dx &\leq 2d_i \int_0^{d_i} \left| \frac{\partial u}{\partial x}(x, 0) \right| p(x, 0) \rho(x) dx + (d-1)M|\Omega_i|\delta \end{aligned} \quad (21)$$

where we use equation (20) together with $\int_{\Omega_i} u(z)p(z)dz = 0$ to obtain the last inequality.

Finally, from (18)

$$\int_{\Omega_i} |u(x, y)|p(x, y)dxdy \leq \int_0^{d_i} |u(x, 0)|p(x, 0)\rho(x)dx + (d-1)M|\Omega_i|\delta$$

from (21)

$$\int_{\Omega_i} |u(x, y)|p(x, y)dxdy \leq 2d_i \int_0^{d_i} \left| \frac{\partial u}{\partial x}(x, 0) \right| p(x, 0) \rho(x) dx + (d-1)M|\Omega_i|\delta(1 + 2d_i)$$

from (19)

$$\begin{aligned} \int_{\Omega_i} |u(x, y)|p(x, y)dxdy &\leq 2d_i \int_{\Omega_i} \left| \frac{\partial u}{\partial x}(x, y) \right| p(x, y) dxdy + (d-1)M|\Omega_i|\delta(1 + 4d_i) \\ \int_{\Omega_i} |u(x, y)|p(x, y)dxdy &\leq 2d_i \int_{\Omega_i} \|\nabla u(x, y)\| p(x, y) dxdy + (d-1)M|\Omega_i|\delta(1 + 4d_i) \end{aligned}$$

Summing up on Ω_i leads to

$$\int_{\Omega} |u(z)|p(z)dz \leq 2D \int_{\Omega} \|\nabla u(z)\| p(z) dz + (d-1)M|\Omega|\delta(1 + 4D)$$

and since δ is arbitrary one gets the desired result. \square

E Regret proofs

Bounding R_1^{RLS} . On E , $\|x_t\| \leq X$ for all $t \in [0, T]$. Moreover, since $\tilde{\theta}_t \in \mathcal{S}$ for all $t \in [0, T]$ due to the rejection sampling, $\text{Tr}(P(\tilde{\theta}_t)) \leq D$. From the definition of the matrix 2-norm, $\sup_{\|x\| \leq X} x^\top P(\tilde{\theta}_t) x \leq X^2 \|P(\tilde{\theta}_t)^{1/2}\|_2^2$. Since for any $A \in \mathbb{R}^{m,n}$, $\|A\|_2 \leq \|A\|$, one has $\|P(\tilde{\theta}_t)^{1/2}\|_2^2 \leq \|P(\tilde{\theta}_t)^{1/2}\|^2 = \text{Tr}P(\tilde{\theta}_t)$. As a consequence, for any $t \in [0, T]$, $\sup_{\|x\| \leq X} x^\top P(\tilde{\theta}_t) x \leq X^2 D$ and the martingale increments are bounded almost surely on E by $2DX^2$. Applying Thm. 3 to R_1^{RLS} with $\epsilon = 2DX^2 \sqrt{2T \log(4/\delta)}$ one obtains that

$$R_1^{\text{RLS}} = \sum_{t=0}^T \{ \mathbb{E}(x_{t+1}^\top P(\tilde{\theta}_{t+1}) x_{t+1} | \mathcal{F}_t) - x_t^\top P(\tilde{\theta}_t) x_t \} \mathbf{1}\{E_t\} \leq 2DX^2 \sqrt{2T \log(4/\delta)}$$

with probability at least $1 - \delta/2$.

Bounding R_3^{RLS} . The derivation of this bound is directly collected from Abbasi-Yadkori and Szepesvári [1]. Since our framework slightly differs, we provide it for the sake of completeness. The whole derivation is performed conditioned on the event E .

$$\begin{aligned} R_3^{\text{RLS}} &= \sum_{t=0}^T \{ z_t^\top \tilde{\theta}_t P(\tilde{\theta}_t) \tilde{\theta}_t^\top z_t - z_t^\top \theta_* P(\tilde{\theta}_t) \theta_*^\top z_t \} = \sum_{t=0}^T \{ \|\tilde{\theta}_t^\top z_t\|_{P(\tilde{\theta}_t)}^2 - \|\theta_*^\top z_t\|_{P(\tilde{\theta}_t)}^2 \}, \\ &= \sum_{t=0}^T (\|\tilde{\theta}_t^\top z_t\|_{P(\tilde{\theta}_t)} - \|\theta_*^\top z_t\|_{P(\tilde{\theta}_t)}) (\|\tilde{\theta}_t^\top z_t\|_{P(\tilde{\theta}_t)} + \|\theta_*^\top z_t\|_{P(\tilde{\theta}_t)}) \end{aligned}$$

By the triangular inequality, $\|\tilde{\theta}_t^\top z_t\|_{P(\tilde{\theta}_t)} - \|\theta_*^\top z_t\|_{P(\tilde{\theta}_t)} \leq \|P(\tilde{\theta}_t)^{1/2}(\tilde{\theta}_t^\top z_t - \theta_*^\top z_t)\| \leq \|P(\tilde{\theta}_t)\| \|\tilde{\theta}_t^\top - \theta_*^\top\| z_t$. Making use of the fact that $\tilde{\theta}_t \in \mathcal{S}$ by construction of the rejection sampling, $\theta_* \in \mathcal{S}$ by Asm. 2 and that $\sup_{t \in [0, T]} \|z_t\| \leq \sqrt{(1+C^2)X^2}$ thanks to the conditioning on E and Prop. 2, one gets:

$$R_3^{\text{RLS}} \leq \sum_{t=0}^T (\sqrt{D} \|\tilde{\theta}_t^\top - \theta_*^\top\| z_t) (2S\sqrt{D} \sqrt{(1+C^2)X^2}) \leq 2SD \sqrt{(1+C^2)X^2} \sum_{t=0}^T \|\tilde{\theta}_t^\top - \theta_*^\top\| z_t$$

and one just has to bound $\sum_{t=0}^T \|\tilde{\theta}_t^\top - \theta_*^\top\| z_t$. Let $\tau(t) \leq t$ be the last time step before t when the parameter was updated. Using Cauchy-Schwarz inequality, one has:

$$\sum_{t=0}^T \|\tilde{\theta}_t^\top - \theta_*^\top\| z_t = \sum_{t=0}^T \|(V_\tau^{1/2}(t)(\tilde{\theta}_{\tau(t)} - \theta_*)^\top V_{\tau(t)}^{-1/2} z_t)\| \leq \sum_{t=0}^T \|\tilde{\theta}_{\tau(t)} - \theta_*\|_{V_{\tau(t)}} \|z_t\|_{V_{\tau(t)}^{-1}}$$

However, on E , $\|\tilde{\theta}_{\tau(t)} - \theta_*\|_{V_{\tau(t)}} \leq \|\tilde{\theta}_{\tau(t)} - \hat{\theta}_{\tau(t)}\|_{V_{\tau(t)}} + \|\theta_* - \hat{\theta}_{\tau(t)}\|_{V_{\tau(t)}} \leq \beta_{\tau(t)}(\delta') + \gamma_{\tau(t)}(\delta') \leq \beta_T(\delta') + \gamma_T(\delta')$ and, thanks to the lazy update rule $\|z_t\|_{V_{\tau(t)}^{-1}} \leq \|z_t\|_{V_t^{-1}} \frac{\det(V_t)}{\det(V_{\tau(t)})} \leq 2\|z_t\|_{V_t}$. Therefore,

$$R_3^{\text{RLS}} \leq 4SD \sqrt{(1+C^2)X^2} (\beta_T(\delta') + \gamma_T(\delta')) \sum_{t=0}^T \|z_t\|_{V_t^{-1}}.$$

Bounding $\sum_{k=1}^K T_k \alpha_k$. From section 4.4,

$$\sum_{k=1}^K T_k \alpha_k \leq 2\tau \sum_{k \in \mathcal{K}^{den}} \|\alpha_k\| + \tau \sum_{k=1}^K \alpha_k.$$

First, it is clear from

$$\begin{aligned} \alpha_k &= (R_{t_k}^{\text{TS},1} + R_{t_k}^{\text{TS},3}) \{E_{t_k}\} \\ &= (J(\tilde{\theta}_{t_k}) - \mathbb{E}[J(\tilde{\theta}_{t_k}) | \mathcal{F}_{t_k}^x, E_{t_k}]) \mathbf{1}\{E_{t_k}\} + (\mathbb{E} \left[\left\| \left(\frac{I}{K(\tilde{\theta}_{t_k})^\top} \right) \right\|_{V_{t_k}^{-1}} | \mathcal{F}_{t_k}^x \right] - \left\| \left(\frac{I}{K(\tilde{\theta}_{t_k})^\top} \right) \right\|_{V_{t_k}^{-1}} \right), \end{aligned}$$

that the sequence $\{\alpha_k\}_{k=1}^K$ is a martingale difference sequence with respect to $\mathcal{F}_{t_k}^x$. Moreover, since $\tilde{\theta}_{t_k} \in \mathcal{S}$ for all $k \in [1, K]$, $\|\alpha_k\| \leq 2D + 2\sqrt{(1+C^2)/\lambda}$. Therefore,

1. $\sum_{k \in \mathcal{K}^{den}} \|\alpha_k\| \leq (2D + 2\sqrt{(1+C^2)})|\mathcal{K}^{den}|$,
2. with probability at least $1 - \delta/2$, Azuma's inequality ensures that $\sum_{k=1}^K \alpha_k \leq (2D + 2\sqrt{(1+C^2)})\sqrt{2|K|\log(4/\delta)}$.

From Lem. 6 and Cor. 2, $|\mathcal{K}^{det}| \leq (n+d)\log_2(1+TX^2(1+C^2)/\lambda)$ and $|K| \leq (n+d)\log_2(1+TX^2(1+C^2)/\lambda)T/\tau$. Finally, one obtains:

$$\sum_{k=1}^K T_k \alpha_k \leq 4(2D + 2\sqrt{(1+C^2)})(n+d)\log_2(1+TX^2(1+C^2)/\lambda)\sqrt{\log(4/\delta)}T/\tau$$

Bounding $\sum_{t=0}^T \|z_t\|_{V_t^{-1}}$. On E , for all $t \in [0, T]$, $\|z_t\|^2 \leq (1+C^2)X^2$. Thus, from Cauchy-Schwarz inequality and Prop. 4,

$$\sum_{t=0}^T \|z_t\|_{V_t^{-1}} \leq \sqrt{T} \left(\sum_{t=0}^T \|z_t\|_{V_t^{-1}}^2 \right)^{1/2} \leq \sqrt{T} \sqrt{2(n+d)(1+C^2)X^2/\lambda} \log^{1/2} \left(1 + \frac{T(1+C^2)X^2}{\lambda(n+d)} \right).$$