

Input Perturbations for Adaptive Control and Learning

Mohamad Kazem Shirani Faradonbeh^a, Ambuj Tewari^b, George Michailidis^a

^a*Department of Statistics, University of Florida, Gainesville, FL, USA, 32611*

^b*Department of Statistics, University of Michigan, Ann Arbor, MI, USA, 48109*

Abstract

This paper studies adaptive algorithms for simultaneous regulation (i.e., control) and estimation (i.e., learning) of Multiple Input Multiple Output (MIMO) linear dynamical systems. It proposes *practical*, easy to implement control policies based on *perturbations* of input signals. Such policies are shown to achieve a worst-case regret that scales as the square-root of the time horizon, and holds uniformly over time. Further, it discusses specific settings where such greedy policies attain the information theoretic lower bound of logarithmic regret. To establish the results, recent advances on self-normalized martingales together with a novel method of policy decomposition are leveraged.

Key words: Finite-time Optimality; Greedy Policies; Adaptive LQRs; System Identification; Decision-making under Uncertainty; Linear-Quadratic; Exploration-Exploitation.

1 Introduction

MIMO systems with linear dynamics represent canonical models in system engineering [1]. Their linear structure renders them amenable to rigorous mathematical analysis of their performance, as well as computationally feasible implementation of control policies [2,3]. Further, they provide insights on how to deal with nonlinear dynamical models [4,5], and accurately represent the behavior of nonlinear systems around their operating equilibrium [6]. In addition, they are also providing benchmarks for prediction tasks [7,8].

The standard formulation is that the system is characterized by autoregressive dynamics (in discrete time), while its operating cost is described by a quadratic function [9]. An extensive and mature literature exists on optimal Linear-Quadratic (LQ) policies when the dynamical model is exactly known [10]. However, for systems whose dynamics are unknown, an adaptive policy is needed in order to simultaneously learn (identify) the model parameters, and control (regulate) its behavior. This gives rise to designing adaptive Linear-Quadratic-Regulators (LQR) in order to balance accurate learning (exploration) with efficient regulation (exploitation) [11]. A common approach in adaptive

control uses the principle of optimism in the face of uncertainty by designing policies assuming that optimistic approximations of the unknown parameters correspond to the truth [12]. However, existing asymptotic [13,14] and non-asymptotic [15,16,17] analyses of optimism-based adaptive regulators indicate that the proposed algorithms are not computationally tractable. That is because non-convex matrix optimization problems, requiring high precision accuracy of the optimal solution, need to be solved repeatedly. In addition, optimism-based policies need to have access to information regarding spectral properties of the system's closed-loop matrix, as well as the statistics of the noise process.

On the other hand, a greedy policy (also referred to as Certainty Equivalent (CE)) uses the optimal feedback gain of the current estimates of the dynamics parameters. Then, upon collecting new input-output observations, this greedy policy learns the unknown parameters through a least squares procedure, and uses the above learning-planning steps in an alternating manner. Hence, CE based policies do not require the aforementioned additional knowledge that optimism-based adaptive policies demand. A recent result establishes that if one randomizes the greedy adaptive regulator, it provides a regret of asymptotically square-root magnitude (with respect to time). In fact, it suffices to obtain the least squares estimate (LSE) of the unknown dynamics matrix, and randomize the estimate by adding a statistically independent random matrix to it [18].

Email addresses: mfaradonbeh@ufl.edu (Mohamad Kazem Shirani Faradonbeh), tewaria@umich.edu (Ambuj Tewari), gmichail@ufl.edu (George Michailidis).

However, the *non-asymptotic* performance of the above randomized greedy algorithm is not satisfactory, due to the presence of (frequent and large) fluctuations in the trajectory of the system. This renders the finite-time performance of this regulator sub-optimal, since the regret can become large over a short time horizon. Technically, the probability of “the regret exceeding a factor of the optimal value” decays polynomially (thus, slowly) as a function of the magnitude of that factor. Hence, in order to ensure that with high probability the regret remains within a margin of the optimal rate, the size of the posited margin needs to be large. In conclusion, the magnitude of the regret can exceed the square root of time¹. Note that the above fluctuations do not compromise the asymptotic performance of randomized greedy regulators, as the time horizon grows to infinity.

This study presents the *first finite-time* theoretical guarantee for such practical algorithms for both system identification and adaptive control. It establishes that a greedy algorithm, subject to a suitable input *perturbation*, achieves finite-time efficiency of cost minimization, as well as learning accuracy. Namely, it is shown that at *all* time steps, the *worst-case* regret scales at a square-root of time rate; in addition, the *uniform* non-asymptotic learning rates of the unknown dynamics are provided. Further, in Section 3, a comprehensive analysis of input perturbation, applicable to various problems in stochastic control [19,20] is presented. To obtain the results, tools from the theory of martingale concentration [21,22] are employed, together with an extension of the policy decomposition method [18] to general control laws with perturbed input.

The remainder of the paper is organized as follows. In Section 2, we provide a precise formulation of the problem under consideration, and discuss the necessary preliminaries. We then address the behavior of the non-asymptotic regret for general control policies with perturbed inputs in Subsection 3.1. Subsequently, the effect of the perturbation signal on the high probability estimation error of the true model is given in Subsection 3.2. In Section 4, we study the growth rate of the regret, as well as the accuracy of learning procedure for the perturbed greedy regulator. We also discuss the case of restricted uncertainty in Subsection 4.1, where side information (such as the support or the rank of the dynamics matrices) is available to the system’s operator.

Remark 1 *Unless otherwise explicitly stated, all stochastic statements in this work hold almost surely.*

Remark 2 *Due to space constraints, all proofs are delegated to the supplementary material [23].*

¹ For more details, see Theorem 3 and the ensuing discussion.

The following notation will be used throughout this paper. For a matrix $A \in \mathbb{C}^{p \times q}$, A' denotes its transpose. When $p = q$, the smallest (respectively largest) eigenvalue of A (in magnitude) is denoted by $\lambda_{\min}(A)$ (respectively $\lambda_{\max}(A)$). For $v \in \mathbb{C}^d$, define the norm $\|v\| = \left(\sum_{i=1}^d |v_i|^2 \right)^{1/2}$. We also use the following notation for the operator norm of matrices. For $A \in \mathbb{C}^{p \times q}$ let, $\|A\| = \sup_{\|v\|=1} \|Av\|$. The symbol \vee (\wedge) will be used to denote the maximum (minimum) of two or more quantities.

2 Problem Formulation

The input-output evolution of the system is governed by the following stochastic dynamical model. Starting from an arbitrary initial $x(0) \in \mathbb{R}^p$, at time $t = 0, 1, \dots$, the transition to the next state (or output) is according to

$$x(t+1) = A_0 x(t) + B_0 u(t) + w(t+1). \quad (1)$$

Namely, $x(t+1)$ is determined by the current state vector $x(t) \in \mathbb{R}^p$, the input signal $u(t) \in \mathbb{R}^r$, and the noise component $w(t+1) \in \mathbb{R}^p$, the latter being a mean-zero random vector: $\mathbb{E}[w(t+1)] = 0$. This study aims to address the problems of system identification and adaptive regulation, when the dynamics parameters A_0, B_0 are *unknown*. Specifically, it aims to design the input signal to minimize the cost, and to learn the matrices A_0, B_0 accurately. We start by presenting some necessary background material, followed by a rigorous formulation of the problems under consideration.

The transition matrix $A_0 \in \mathbb{R}^{p \times p}$ models the effect of the current state signal, while $B_0 \in \mathbb{R}^{p \times r}$ is the input matrix indicating the influence of the control signal. The noise vectors $\{w(t)\}_{t=1}^{\infty}$ are assumed to be independent, but are not required to be identically distributed (strict-sense stationary), and can even be heteroscedastic (wide-sense non-stationary). Note that the independence assumption on the noise vectors is not restrictive and can be replaced by assuming that they form a martingale difference sequence, without impacting the results established in this work.

To proceed, let $Q \in \mathbb{R}^{p \times p}$ and $R \in \mathbb{R}^{r \times r}$ be positive definite symmetric matrices. Then, the instantaneous cost of control law π at time t is defined as

$$c_t(\pi) = x(t)' Q x(t) + u(t)' R u(t). \quad (2)$$

Intuitively speaking, a desired control policy minimizes the long-run average cost. For a precise definition, we first specify the families of causal and adaptive policies, denoted by \mathcal{C} and \mathcal{A} , respectively. A *causal* policy $\pi \in \mathcal{C}$ determines the input based on the dynamics matrices A_0, B_0 , the matrices Q, R which are assessing the costs

of the state and control signals, and the record of the both signals thus far. That is,

$$u(t) = \pi \left(A_0, B_0, Q, R, \{x(i)\}_{i=0}^t, \{u(j)\}_{j=0}^{t-1} \right),$$

for all $t \geq 0$. Note that the above mapping can be a stochastic one.

An *adaptive* operator does not have access to the dynamical parameters A_0, B_0 . Therefore, $\hat{\pi} \in \mathcal{A}$ designs the input signal according to other available information:

$$u(t) = \hat{\pi} \left(Q, R, \{x(i)\}_{i=0}^t, \{u(j)\}_{j=0}^{t-1} \right).$$

Obviously, it holds that $\mathcal{A} \subsetneq \mathcal{C}$. The control objective is to minimize the average cost of the system, subject to the stochastic dynamics equation (1). Naturally, all $\pi \in \mathcal{C}$ are expected to target the same objective being determined by the context of the problem. Therefore, the cost matrices Q, R are assumed known for design of all $\hat{\pi} \in \mathcal{A}$. The best performance among causal control laws belongs to π^* , which attains the smallest average cost:

$$\limsup_{n \rightarrow \infty} n^{-1} \sum_{t=0}^{n-1} c_t(\pi^*) = \inf_{\pi \in \mathcal{C}} \limsup_{n \rightarrow \infty} n^{-1} \sum_{t=0}^{n-1} c_t(\pi).$$

In order to find π^* , one needs to solve a Riccati equation [17]. To do so, we assume that the following necessary condition of stabilizability is satisfied.

Assumption 1 *There exists a feedback gain matrix $L \in \mathbb{R}^{r \times p}$ such that $|\lambda_{\max}(A_0 + B_0 L)| < 1$.*

Remark 3 *Henceforth, we employ the shorthand $\theta \in \mathbb{R}^{p \times q}$ (where $q = p + r$), to denote the pair $[A, B]$ of dynamics matrices $A \in \mathbb{R}^{p \times p}$, $B \in \mathbb{R}^{p \times r}$.*

Then, for the dynamics parameter θ , define the Riccati operator $\Psi_\theta(\cdot) : \mathbb{R}^{p \times p} \rightarrow \mathbb{R}^{p \times p}$:

$$\Psi_\theta(K) = Q + A'KA - A'KB(B'KB + R)^{-1}B'KA.$$

Suppose that $K(\theta)$ solves the algebraic Riccati equation

$$K(\theta) = \Psi_\theta(K(\theta)). \quad (3)$$

It has been established that whenever θ is stabilizable, the positive definite $K(\theta)$ exists, and is unique [24]. In addition, the recursive Riccati equation $K_{t+1} = \Psi_\theta(K_t)$, is known to provide a sequence of matrices converging exponentially fast to $K(\theta)$ [24]. Using the matrix $K(\theta)$ in (3), let

$$L(\theta) = -(B'K(\theta)B + R)^{-1}B'K(\theta)A. \quad (4)$$

The linear time invariant feedback gain $L(\theta_0)$ provides an optimal causal regulator for a system of dynamics parameter θ_0 [17]. It is also well-known that $L(\theta_0)$ is a stabilizer; i.e. $|\lambda_{\max}(A_0 + B_0 L(\theta_0))| < 1$. Henceforth, let π^* be

$$\pi^* : u(t) = L(\theta_0)x(t), \quad t = 0, 1, 2, \dots \quad (5)$$

Optimality of π^* motivates us to consider adaptive regulators of the form $u(t) = L_t x(t)$. Note that since information on the unknown dynamics parameters is acquired over time, such adaptive regulators are time-varying. Further, to address identification of the unknown parameters, $\hat{\pi} \in \mathcal{A}$ is subject to exogenous randomness through additional perturbations. So, in general we study causal policies of the form

$$\pi : u(t) = L_t x(t) + v(t), \quad t = 0, 1, 2, \dots, \quad (6)$$

where L_t is a $r \times p$ feedback gain matrix (determined according to $A_0, B_0, Q, R, \{x(i)\}_{i=0}^t, \{u(j)\}_{j=0}^{t-1}$), and the r dimensional perturbation vectors are $\{v(t)\}_{t=0}^\infty$. We denote the above policy by $\pi = \{L_t, v(t)\}_{t=0}^\infty \in \mathcal{C}$. Similarly, the adaptive regulator $\hat{\pi} = \{\hat{L}_t, v(t)\}_{t=0}^\infty \in \mathcal{A}$ is defined according to the perturbation signal $v(t)$, together with the matrices $\{\hat{L}_t\}_{t=0}^\infty$, where \hat{L}_t is a function of $Q, R, \{x(i)\}_{i=0}^t, \{u(j)\}_{j=0}^{t-1}$.

Remark 4 *In this work, we study mean zero and statistically independent perturbation signals $\{v(t)\}_{t=0}^\infty$, which also are independent of the noise process $\{w(t)\}_{t=1}^\infty$.*

The efficiency of $\pi \in \mathcal{C}$ is determined by comparing the accumulative cost of π to that of π^* in (5). In fact, we investigate the following natural definition of *regret*:

$$\mathcal{R}_n(\pi) = \sum_{t=0}^{n-1} [c_t(\pi) - c_t(\pi^*)]. \quad (7)$$

Moreover, let $\hat{\theta}_n$ denote the LSE of the true dynamics parameter θ_0 , using the sample $\{x(i)\}_{i=0}^n, \{u(j)\}_{j=0}^{n-1}$. In order to measure the precision of the learning procedure, we consider the error $\|\hat{\theta}_n - \theta_0\|$. This study provides a non-asymptotic analysis of the performance of (causal and adaptive) control laws, addressing both aspects of learning and planning. The following two problems are rigorously addressed.

Problem 1 (Regulation) *For $\pi \in \mathcal{C}$, determine high probability uniform regret bounds. That is, specify the*

sequence $\{\tau_n(\delta)\}_{n=1}^\infty$ and the integer $n_0(\delta)$ such that

$$\sup_{n \geq n_0(\delta)} \tau_n(\delta) \mathcal{R}_n(\boldsymbol{\pi}) < \infty,$$

with probability at least $1 - \delta$.

Problem 2 (Identification) Find the high probability learning accuracy that holds uniformly over time. That is, specify $\{\tau_n(\delta)\}_{n=1}^\infty, n_0(\delta)$, so that with probability at least $1 - \delta$, the following holds:

$$\sup_{n \geq n_0(\delta)} \tau_n(\delta) \left\| \hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0 \right\| < \infty.$$

3 General Input Perturbations

In this section, we study arbitrary causal control laws with perturbed input signal. The regret analysis which considers the effects of both the linear feedback gains, as well as the perturbation signals will be presented in Subsection 3.1. Subsequently, we address the identification performance in Subsection 3.2, and discuss how the structure of a causal control law determines the finite sample accuracy of learning $\boldsymbol{\theta}_0$. The main results, Theorem 1 and Theorem 2, both provide high probability bounds which hold uniformly over time.

3.1 Regret Bound

We address Problem 1 for the regret of causal policy $\boldsymbol{\pi} = \{L_t, v(t)\}_{t=0}^\infty$. In absence of a perturbation sequence (i.e., $v(t) = 0$), it has been shown that $\mathcal{R}_n(\boldsymbol{\pi})$ scales as $\sum_{t=0}^{n-1} \|(L(\boldsymbol{\theta}_0) - L_t)x(t)\|^2$, as n grows [18]. Theorem 1 generalizes the above result along the following directions. First, the input signals provided by the linear feedback gains $\{L_t\}_{t=0}^\infty$ are perturbed with exogenous random signals $\{v(t)\}_{t=0}^\infty$. Second, the upper bounds for the regret presented in this subsection hold uniformly over time. Thus, the results in this subsection are still applicable, even if the time horizon of interacting with the system is not necessarily large. To proceed, we define the following energy quantities which will be used throughout the non-asymptotic analyses.

Definition 1 For $\boldsymbol{\pi} = \{L_t, v(t)\}_{t=0}^\infty \in \mathcal{C}$, suppose that $\{\bar{v}_t\}_{t=0}^\infty, \{\bar{\ell}_t\}_{t=0}^\infty$ are deterministic sequences satisfying $\|(L(\boldsymbol{\theta}_0) - L_t)x(t)\| \leq \bar{\ell}_t, \|v(t)\| \leq \bar{v}_t$, for all $t \geq 0$. Then, let $\mathcal{E}_v = \sum_{t=0}^{n-1} \bar{v}_t^2, \mathcal{E}_\ell = \sum_{t=0}^{n-1} \bar{\ell}_t^2$. Similarly, define \bar{w}_t, \bar{x}_t , as well as $\mathcal{E}_w, \mathcal{E}_x$, for $w(t), x(t)$, respectively. Fi-

nally, denote $\bar{w}^* = \max_{1 \leq t \leq n} \bar{w}_t$, and

$$\mathcal{E}_\pi = \left(\sum_{t=0}^{n-1} \bar{v}_t^2 (\bar{\ell}_t + \bar{x}_t)^2 \right)^{1/2} + \left(\bar{w}^{*2} \sum_{t=0}^{n-1} (\bar{\ell}_t + \bar{v}_t)^2 \right)^{1/2}.$$

The dependence of all terms \mathcal{E} on n is suppressed for notational convenience. Intuitively, \mathcal{E}_y corresponds to the deterministic upper bound for the energy of the signal $\{y(t)\}_{t=0}^\infty$ up to time n . Subsequently, we establish a high probability regret bound in terms of the energy quantities defined above.

Theorem 1 Using Definition 1 for $\boldsymbol{\pi} = \{L_t, v(t)\}_{t=0}^\infty$, there is a fixed constant $C_1 < \infty$ such that for all $\delta > 0$,

$$\mathbb{P} \left(\sup_{n \geq 1} \frac{\mathcal{R}_n(\boldsymbol{\pi})}{\mathcal{E}_v + \mathcal{E}_\ell + \mathcal{E}_\pi \log^{1/2}(n\delta^{-1})} \leq C_1 \right) \geq 1 - \delta.$$

One can explicitly calculate C_1 using the statements in the proof of Theorem 1 in the supplementary material [23]. To provide further intuition and outline the consequences of the above general result, we state the following corollary:

Corollary 1 Let $\mathbb{E}[\cdot]$ be the expectation with respect to the probability measures induced by $\boldsymbol{\pi}, \{w(t)\}_{t=1}^\infty$. Then, with the notation of Theorem 1, we have

$$\sup_{n \geq 1} \frac{\mathbb{E}[\mathcal{R}_n(\boldsymbol{\pi})]}{\mathcal{E}_v + \mathcal{E}_\ell} \leq C_1.$$

Note that the expressions in the above finite-time regret bounds provide a novel decomposition as follows. First, the previous high probability upper bounds for the regret consist of summation of the terms $\{\bar{\ell}_t, \bar{v}_t\}_{t=0}^{n-1}$ [17], while in Theorem 1 and Corollary 1, all the terms $\mathcal{E}_v, \mathcal{E}_\ell, \mathcal{E}_\pi$ are in the form of sum-of-squares of $\{\bar{\ell}_t, \bar{v}_t\}_{t=0}^{n-1}$. For an adaptive regulator, $\bar{\ell}_t$ shrinks as time goes by, since the decision-maker acquires more data to learn the dynamics. Hence, if the perturbation signal diminishes sufficiently fast, we get a significantly smaller regret bound compared to the existing results. For example, the magnitude of the lower bound for the regret changes from $n^{1/2}$ [17] to $\log n$ [18]. Further, this regret specification sets the stage for Theorem 3, which establishes the finite time efficiency of a perturbed greedy regulator. In the sequel, we discuss other aspects of the result of Theorem 1.

Considering Theorem 1 and Corollary 1, the only different term is $\mathcal{E}_\pi \log^{1/2}(n\delta^{-1})$. Thus, the stochastic behavior of $\boldsymbol{\pi}$ is reflected by \mathcal{E}_π , and the other terms reflect

the expected deviations from the optimal regulator π^* . In other words, the cumulative deviation from the desired optimal trajectory consists of the systematic long lasting portion, as well as the spontaneous fluctuations. According to Corollary 1, the former is reflected in the energy of the perturbation signal \mathcal{E}_v , and the energy of the instantaneous deviations in the feedback gains \mathcal{E}_ℓ . The magnitude of the spontaneous fluctuations is with high probability determined by \mathcal{E}_π , which essentially reflects the interaction between $x(t), v(t), L_t - L(\theta_0)$, as well as the interaction between $w(t+1), v(t), L_t - L(\theta_0)$. So, \mathcal{E}_π can be interpreted as the *fluctuation* energy. We will shortly compare the aforementioned involved quantities for different perturbation signals.

If there is no perturbation, $v(t) = 0$, \mathcal{E}_ℓ is shown to be a tight asymptotic bound for $\mathcal{R}_n(\pi)$ [18]. To compare with the result of Theorem 1, letting $\bar{v}_t = 0$ for all $t \geq 0$, suppose that the noise is uniformly bounded. Then, since $\sup_{t \geq 1} \bar{w}_t < \infty$, the fluctuation energy \mathcal{E}_π can

be replaced with $\mathcal{E}_\ell^{1/2}$. This leads to the non-asymptotic regret bound $\mathcal{E}_\ell + \mathcal{E}_\ell^{1/2} \log^{1/2}(n\delta^{-1})$, which is a finite time counterpart for the asymptotic bound \mathcal{E}_ℓ [18].

Next, we discuss a case of diminishing perturbations to compare the contributions of different factors toward the uniform bound of Theorem 1. Assuming $\sup_{t \geq 1} \bar{w}_t < \infty$, let $\{v(t)\}_{t=0}^\infty$ be a diminishing signal such that $\sup_{t \geq 0} t^{\alpha_1} \bar{v}_t < \infty$, for some $\alpha_1 > 0$. Further, let the deviation from the optimal feedback satisfy

$$\sup_{t \geq 0} t^{\alpha_2} \|L(\theta_0) - L_t\| < \infty,$$

for some $\alpha_2 > 0$. Since the above policy stabilizes the system [24], the state signal is uniformly bounded as well. Then, Theorem 1 provides the following uniform regret bound

$$n^{1-2\alpha_1 \wedge 2\alpha_2} + n^{1/2-\alpha_1 \wedge \alpha_2} \log^{1/2}(n\delta^{-1}).$$

Similarly, it is straightforward to show that uniform boundedness of the noise vectors and the state signal leads to the bound

$$\mathcal{E}_v + \mathcal{E}_\ell + (\mathcal{E}_v + \mathcal{E}_\ell)^{1/2} \log^{1/2}(n\delta^{-1}).$$

Therefore, the perturbation signal does not lead to an increase in the high probability regret bound as long as it diminishes as fast as $L(\theta_0) - L_t$. For the special case of $\alpha_1 \wedge \alpha_2 = 1/2$, the regret is of the order

$$\log n + \log^{1/2} n \log^{1/2} \delta^{-1},$$

which is also the information theoretic lower bound [18].

Finally, Theorem 1 shows that the non-asymptotic high probability uniform bound of the regret is memoryless [18]. That being said, if hypothetically the operator stops using perturbed inputs or deviated feedback gains, the regret bound of Theorem 1 freezes and does not grow anymore. Therefore, Theorem 1 and Corollary 2 indicate that the adaptive decision-maker can revoke the sub-optimal control inputs in the past time. In other words, the history of the non-optimal actions previously taken, does not significantly influence the future trajectory of the dynamical system (memorylessness).

3.2 Identification Bound

Next, we address Problem 2 by analyzing the effect of the input perturbation for learning the unknown dynamics parameter θ_0 . Based on the linear data generation mechanism in (1), a natural estimation procedure for θ_0 is least squares. That is, to regress every state sample $x(t+1)$ on the previous samples of input $u(t)$ and state $x(t)$, to get the following LSE:

$$\hat{\theta}_n = \arg \min_{\theta \in \mathbb{R}^{p \times q}} \sum_{t=0}^{n-1} \left\| x(t+1) - \theta \begin{bmatrix} x(t) \\ u(t) \end{bmatrix} \right\|^2. \quad (8)$$

In the statistical analyses of the accuracy of $\hat{\theta}_n$, the associated stochastic process $\{w(t)\}_{t=1}^\infty$, as well as the exogenous one $\{v(t)\}_{t=0}^\infty$, play an important role. So, we use the following definition which uses the second moments of the above stochastic processes.

Definition 2 Suppose that $\pi = \{L_t, v(t)\}_{t=0}^\infty \in \mathcal{C}$.

Then, let $\Sigma_w = \sum_{t=1}^{n-1} \mathbb{E}[w(t)w(t)']$, $\Sigma_v = \sum_{t=0}^{n-2} \mathbb{E}[v(t)v(t)']$, and $\lambda_n = |\lambda_{\min}(\Sigma_w)| \wedge |\lambda_{\min}(\Sigma_v)|$. Further, letting $D_0 = A_0 + B_0L(\theta_0)$, and

$$\bar{y}_t = \|B_0\| \bar{v}_t + \bar{w}_{t+1} + 2\|B_0\| \bar{\ell}_t + 2\|D_0\| \bar{x}_t,$$

define the following quantities similar to Definition 1:

$$\mathcal{E}_1 = \left(\sum_{t=0}^{n-2} \bar{y}_t^2 (\|B_0\| \bar{v}_t + \bar{w}_{t+1})^2 \right)^{1/2},$$

$$\mathcal{E}_2 = \left(\sum_{t=0}^{n-2} \bar{v}_t^2 (2[1 + \|L(\theta_0)\|] \bar{x}_t + 2\bar{\ell}_t + \bar{v}_t)^2 \right)^{1/2}.$$

Note that $\Sigma_w, \Sigma_v, \mathcal{E}_1, \mathcal{E}_2$ all depend on n . In order to interpret $\mathcal{E}_1, \mathcal{E}_2$, apply the causal policy $\pi = \{L_t, v(t)\}_{t=0}^\infty$ to the dynamical model in (1). Then, the closed-loop dynamics becomes $x(t+1) = D_t x(t) + B_0 v(t) + w(t+1)$, where $D_t = A_0 + B_0 L_t$. Thus, the noise in the evolution

of the system is $B_0v(t) + w(t+1)$, which indicates that \mathcal{E}_1 reflects a deterministic upper bound for the interaction energy between the observed signals $x(t)$, and the unobserved noise $B_0v(t) + w(t+1)$. Similarly, \mathcal{E}_2 bounds the interaction energy between the signal $x(t)$ and the perturbation $v(t)$.

In order to accurately identify the true parameter θ_0 , λ_n cannot be very small. Broadly speaking, the magnitude of λ_n determines the extent to which both the perturbation and the noise processes *excite* all coordinates of the matrix θ_0 . So, the precision of the learning procedure in (8) highly depends on λ_n . To establish a uniform upper bound for learning accuracy, we assume the followings for the quantities in Definition 2:

$$|\lambda_{\min}(\Sigma_w)| \geq 4\mathcal{E}_1 \log^{1/2}(8p\delta^{-1}), \quad (9)$$

$$|\lambda_{\min}(\Sigma_v)| \geq 4\mathcal{E}_2 \log^{1/2}(8q\delta^{-1}). \quad (10)$$

Intuitively, the above inequalities generalize the persistent excitation condition [25] to cases with diminishing (i.e., non-persistent) perturbations. Later on, we will show that one can easily satisfy (9), (10). Further, writing down the closed-loop dynamics, $|\lambda_{\min}(\Sigma_w)|$ denotes the minimum energy of the noise sequence $B_0v(t) + w(t)$. Going back to the interpretation of \mathcal{E}_1 after definition 2, the condition (9) indeed states that the noise energy is larger than the interaction energy (of the noise and the signal $x(t)$), with high probability. Therefore, the interaction is not powerful enough to prevent the noise from exciting all coordinates of the unknown transition matrix A_0 . Similar explanation for (10) states that the excitation of B_0 by the perturbation signal is not masked by its interaction with the state signal. These are formally presented in the following result.

Theorem 2 *Using the notation of Definitions 1, 2, with probability at least $1 - \delta$ we have*

$$\sup \frac{\lambda_n \|\hat{\theta}_n - \theta_0\|^2}{w^* \log[(\mathcal{E}_\ell + \mathcal{E}_x + \mathcal{E}_v)\delta^{-1}]} \leq C_2,$$

where the supremum is over all $n \geq 1$ satisfying (9), (10).

The constant C_2 depends on the true dynamics parameter θ_0 , and can be extracted from the proof in the supplementary material [23]. Next, to discuss Theorem 2, suppose that the covariance matrices of the noise vectors are bounded from below: $\inf_{n \geq 1} n^{-1} |\lambda_{\min}(\Sigma_w)| > 0$.

Further, for the perturbation process assume that

$$\inf_{t \geq 1} t^{\alpha_1} |\lambda_{\min}(\mathbb{E}[v(t)v(t)'])| > 0, \quad \sup_{t \geq 1} t^{\alpha_2} \bar{v}_t < \infty,$$

for some $0 \leq \alpha_1 < 1$, and $0 \leq \alpha_2 \leq 1$. Hence, we obtain

$\inf_{n \geq 1} n^{\alpha_1 - 1} \lambda_n > 0$. Further, similar to the discussion after

Theorem 1, suppose that both the noise and the state signal are uniformly bounded (e.g., $\bar{\ell}_t$ is diminishing). Assuming $\alpha_1 - \alpha_2 < 1/2$, since

$$\sup_{n \geq 1} n^{-1/2} \mathcal{E}_1 < \infty, \quad \sup_{n \geq 1} n^{\alpha_2 - 1/2} \mathcal{E}_2 < \infty,$$

the inequalities in (9), (10) are satisfied as long as n is up to a constant factor at least $(-\log \delta)^{1/(1-2\alpha_1+2\alpha_2)}$. Then, Theorem 2 implies the identification bound $n^{\alpha_1/2-1/2} \log^{1/2}(n\delta^{-1})$ for uniform learning accuracy of estimating the unknown parameter θ_0 . Hence, the identification accuracy is basically determined by the diminishing rate of the perturbation signal. Moreover, for the case of persistent perturbation $\alpha_1 = \alpha_2 = 0$, LSE in (8) achieves the optimal uniform learning rate of $n^{-1/2} \log^{1/2}(n\delta^{-1})$. Note that in comparison to the analysis presented in the previous subsection, a larger perturbation leads to more accurate learning together with larger deviations from the optimal cost, and vice versa. Balancing this trade-off is the main challenge for adaptive policies, and will be addressed in the sequel. Finally, for $\alpha_1 = 1/2$, the result extends the existing asymptotic identification rates [18].

4 Perturbed Greedy Regulator

Next, we analyze the non-asymptotic regret of perturbed greedy policies, as well as their identification errors. For this purpose, we state a fairly general condition about the probabilistic properties of the noise process, and present an adaptive regulator in Algorithm 1. The subsequent contents consist of establishing that Algorithm 1 addresses the main dilemma of reinforcement learning; balancing the exploration and the exploitation. In this section, we assume the following about the noise process $\{w(t)\}_{t=1}^\infty$ in (1).

Assumption 2 *The noise has a sub-Weibull distribution. That is, for some fixed $\tilde{\beta}, \beta, \alpha > 0$, and for all $t \geq 1$, and $\eta > 0$, we have*

$$\mathbb{P}(\|w(t)\| > \eta) \leq \tilde{\beta} \exp(-\beta^{-1}\eta^\alpha)$$

Further, we assume that the covariance matrices of the noise vectors satisfy

$$\inf_{n \geq 1} n^{-1} \sum_{t=1}^n |\lambda_{\min}(\mathbb{E}[w(t)w(t)'])| \geq \sigma_0 > 0.$$

Note that positive definiteness of the covariance matrices $\mathbb{E}[w(t)w(t)']$ is sufficient for the second part of Assumption 2, but is not necessary. Regarding the first

part of Assumption 2, sub-Weibull distributions are remarkably general in learning theory when analyzing the finite-time performance [8]. In fact, they encompass a wide range of distributions being commonly used in the literature, such as uniformly bounded ($\alpha = \infty$), sub-Gaussian ($\alpha = 2$), and sub-Exponential ($\alpha = 1$) random vectors, as well as heavy-tail distributions for which moment generating functions do not exist ($\alpha < 1$). Later on, we will see that the exponent α which determines the decay rate in the probability distributions of the noise vectors, plays a crucial role in the non-asymptotic analyses for both the regret and the identification error. Note that the noise distributions are not required to have densities or continuous cumulative distribution functions.

Algorithm 1 : Perturbed Greedy Regulator

Inputs: $\gamma > 1$, $0 < r\underline{C} < \overline{C} < \infty$
 Let $L(\hat{\theta}_0) \in \mathbb{R}^{r \times p}$ be a stabilizer
for $m = 0, 1, 2, \dots$ **do**
 while $t < \gamma^m$ **do**
 Draw perturbation $v(t)$ according to (11)
 Apply input $u(t) = L(\hat{\theta}_t)x(t) + v(t)$
 $\hat{\theta}_{t+1} = \hat{\theta}_t$
 end while
 Update the estimate $\hat{\theta}_t$ by (8)
end for

The pseudo-code of the perturbed greedy regulator is provided in Algorithm 1. It is an episodic algorithm; learning of the unknown parameters is deferred until sufficiently many observations have been collected. More precisely, Algorithm 1 updates the parameter estimates only at the end of epochs of exponentially growing length. This lets the solution of Riccati equation (3) be efficiently used by preventing unnecessary computations. Note that numerical computation of the optimal feedback gain in (4) is practically nontrivial [26].

Formally, the algorithm starts with an arbitrary initial stabilizing feedback matrix $L(\hat{\theta}_0)$. An algorithmic procedure for computing $L(\hat{\theta}_0)$ in finite time is available in the literature, together with theoretical performance guarantees for stabilization sets [24]. Therefore, we can assume that the closed-loop system remains stable during the execution of Algorithm 1. At each time $t = 0, 1, 2, \dots$, Algorithm 1 applies the adaptive control input $u(t) = L(\hat{\theta}_t)x(t) + v(t)$. The design of the perturbation signals $v(t)$ will be discussed shortly. Further, the lengths of the epochs are determined by the rate $\gamma > 1$ as follows. If t satisfies $t = \lfloor \gamma^m \rfloor$ for some $m = 0, 1, \dots$, Algorithm 1 tunes the regulator by finding $\hat{\theta}_t$ according to (8). Otherwise, for $t \neq \lfloor \gamma^m \rfloor$, the learning step will be skipped and no update occurs: $\hat{\theta}_t = \hat{\theta}_{t-1}$. The perturbation signals during each epoch $m \geq 1$ are wide-sense stationary (homoscedastic) bounded random

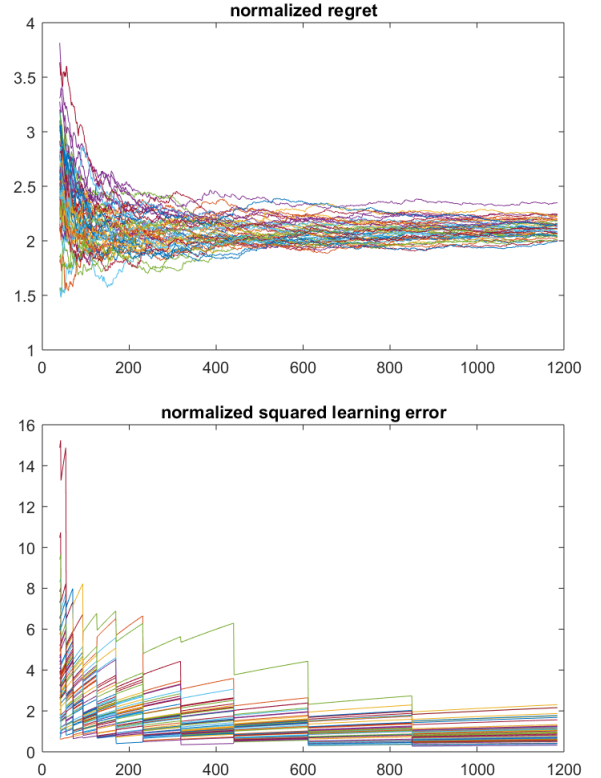


Figure 1. Performance of Algorithm 1: normalized regret $(n^{-1/2} \log^{-2} n) \mathcal{R}_n(\hat{\pi})$ vs n (top), and normalized squared learning error $n^{1/2} \|\hat{\theta}_n - \theta_0\|^2$ vs n (bottom).

signals with positive definite covariance matrices; for all $\gamma^m \leq t < \gamma^{m+1}$, we have $\mathbb{E}[v(t)v(t)'] = \Sigma_m$, and $\|v(t)\| < \overline{v}_m$, such that for some arbitrary universal constants $\underline{C} > 0$, $\overline{C} < \infty$,

$$\underline{C} < m^{-2} \gamma^{m/2} |\lambda_{\min}(\Sigma_m)| \leq m^{-2} \gamma^{m/2} \overline{v}_m^{-2} < \overline{C}. \quad (11)$$

In general, any sequence of mean-zero independent random vectors satisfying (11) can be used as the perturbation signal. Note that we need $r\underline{C} \leq \overline{C}$, since $r |\lambda_{\min}(\Sigma_m)| \leq \text{tr}(\mathbb{E}[v(t)v(t)']) = \mathbb{E}[v(t)'v(t)] \leq \overline{v}_m^2$.

Although larger values of γ make the epochs longer, and thus fewer updates of the parameter estimates occur, the rates of regulation and identification do not depend on the magnitude of γ . Intuitively, it is because Algorithm 1 utilizes all the random perturbation signals during each epoch when learning the dynamics parameter at the end of that epoch. The following theorem addresses the uniform finite time rates for the above regulator.

Theorem 3 *Suppose that $\hat{\pi}$ is the adaptive regulator of Algorithm 1. Let $\hat{\theta}_n$ be the parameter estimate at time*

n . Then, with probability at least $1 - \delta$ we have

$$\sup_{n \geq n_0} \frac{\mathcal{R}_n(\hat{\pi})}{n^{1/2} \log^{2V4/\alpha}(n\delta^{-1})} \leq C_3,$$

$$\sup_{n \geq n_0} \frac{\left\| \hat{\theta}_n - \theta_0 \right\|^2}{n^{-1/2} \log^{2/\alpha-1}(n\delta^{-1})} \leq C_4,$$

where $n_0 \geq C_0 \left(\log^{2+4/\alpha} \delta^{-1} \right) (\log \log \delta^{-1})$.

The fixed constants C_0, C_3, C_4 depend on the truth θ_0 , the noise parameters $\beta, \tilde{\beta}, \sigma_0$, and the constants $\gamma, \underline{C}, \overline{C}$ in (11). Using the notation of Theorem 2, λ_n scales linearly with \underline{C} . Thus, the constant C_4 in Theorem 3 is proportional to \underline{C}^{-1} , as well as the quantity \mathcal{E}_ℓ in Theorem 1 (see (14)). Further, with the notation of Theorem 1, \mathcal{E}_v scales linearly with \overline{C} . Therefore, the constant C_3 in Theorem 3 scales proportional to $\overline{C} + \underline{C}^{-1}$. Figure 1 depicts the performance of Algorithm 1 for the dynamics and cost matrices in (12), and $\gamma = 1.2, \underline{C} = 1, \overline{C} = 10$. The distribution of the random vectors $v(t)$ is Gaussian, truncated to satisfy (11). Finally, the initial stabilizer $L(\hat{\theta}_0)$ is provided by applying the stabilization algorithm of Faradonbeh et al. [24] to the first 17 samples; $\{x(t)\}_{t=0}^{16}, \{u(t)\}_{t=0}^{15}$. The curves in Figure 1 correspond to different replicates, and fully reflect the statement of Theorem 3.

Subsequently, we compare Theorem 3 with the relevant asymptotic results obtained without input perturbation. The comparisons focus on the performance of regulation and identification of Algorithm 1 vis-a-vis those of previously proposed adaptive regulators known as Randomized Certainty Equivalence (RCE) and Thompson Sampling (TS) [18]. The former policy consists of computing the LSE of the true parameter at the end of each epoch, and then adding a random matrix to it. The latter policy adopts a Bayesian approach, and approximates θ_0 with draws from Gaussian posteriors at the end of every epoch. Although the presented asymptotic bounds for regret and learning errors of RCE and TS are similar to those of Algorithm 1 (see [18]), their non-asymptotic performance deteriorates due to fluctuations in the state trajectory. Intuitively speaking, when the time period of interacting with the system is finite, polynomial functions of $n\delta^{-1}$ appear in the expressions of both regret and learning error. Therefore, the potential stochastic fluctuations are neither sufficiently *scarce* (because of δ^{-1}), nor *small* enough (because of n). Note that the results of Theorem 3 are stronger than the existing analyses also because of providing uniform bounds (the previous results address ‘lim sup’ instead of ‘sup’).

These improvements are visualized for the system of the previous numerical example. In Figure 2, we report the uniform upper bound of the normalized regret and

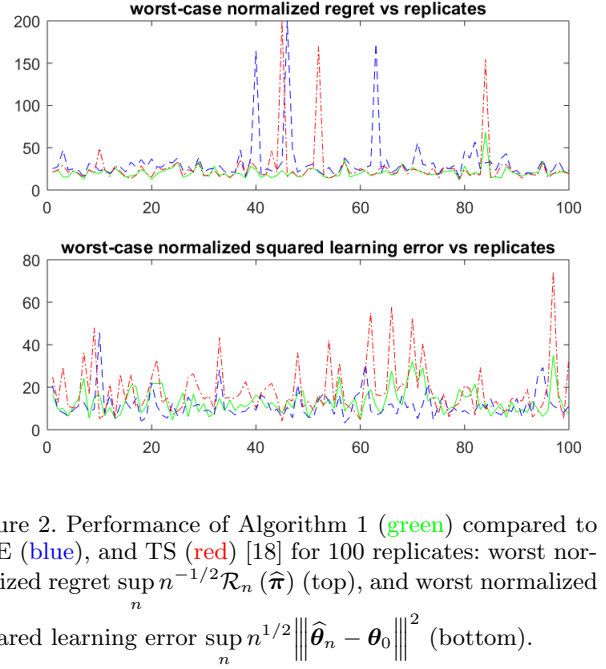


Figure 2. Performance of Algorithm 1 (green) compared to RCE (blue), and TS (red) [18] for 100 replicates: worst normalized regret $\sup_n n^{-1/2} \mathcal{R}_n(\hat{\pi})$ (top), and worst normalized squared learning error $\sup_n n^{1/2} \left\| \hat{\theta}_n - \theta_0 \right\|^2$ (bottom).




Figure 3. Upper bounds of regulation and identification obtained by 100 replicates of Algorithm 1 (green), RCE (blue), and TS (red) [18]: the largest values of the normalized regret $n^{-1/2} \mathcal{R}_n(\hat{\pi})$ (top), and the normalized squared learning error $n^{1/2} \left\| \hat{\theta}_n - \theta_0 \right\|^2$ (bottom), both as functions of n .

the normalized learning error for 100 replicates. Clearly, the fluctuations of Algorithm 1 are significantly smaller than the counterparts of RCE and TS. In Figure 3, the normalized performance quantities are calculated versus time, and the largest among the 100 replicates is being plotted. The non-asymptotic improvement of the perturbed greedy regulator over RCE and TS can be seen in both figures. A recent study provides an analogous regret bound for single dimensional systems, wherein at every time step the parameter estimate is being updated

$$A_0 = \begin{bmatrix} 0.13 & 0.35 & -0.26 \\ -1.34 & -0.30 & -1.75 \\ 1.18 & 0 & -0.29 \end{bmatrix}, \quad B_0 = \begin{bmatrix} -0.83 & -0.53 & 0.52 \\ -0.98 & -2.00 & 0 \\ -1.16 & 0.96 & -0.04 \end{bmatrix}, \quad Q = \begin{bmatrix} 0.79 & -0.15 & 0.09 \\ -0.15 & 0.60 & -0.04 \\ 0.09 & -0.04 & 0.61 \end{bmatrix}, \quad R = \begin{bmatrix} 0.52 & -0.06 & -0.07 \\ -0.06 & 0.39 & -0.04 \\ -0.07 & -0.04 & 0.67 \end{bmatrix}. \quad (12)$$

and perturbed by diminishing Gaussian random matrices [27]. The computational cost of the latter regulator is exponentially larger than that of Algorithm 1, since Riccati equations need to be repeatedly solved [27].

4.1 Partial Uncertainty of Parameters

Next, we study the effect of restricted uncertainty about the true dynamical parameter θ_0 on the performance of adaptive regulators. In fact, the context of the plant under consideration is capable of providing side information about θ_0 . Examples of such information include the structural connectivity of networks which determines the support of θ_0 , and artificially designed systems where the construction imposes restriction on possible sets of $p \times q$ dynamics matrices. Further, in the systems with longer memory where the current state depends on multiple previous time steps, the transition matrix A_0 has a particular structure once the dynamics is written in the form of (1) [11,8]. Therefore, the adaptive operator can restrict the learning of the unknown dynamics parameter to an specific subset of $\mathbb{R}^{p \times q}$.

In the sequel, we show that such additional information can lead to significant improvements in the regret bounds. Although such situations are investigated for the settings with finitely many possible control inputs, non-asymptotic analyses are not currently available for LQ systems. In fact, for general Markov Decision Processes (MDP) it is known that the regret is modulo a constant factor at least $n^{1/2}$ [28]. However, if one knows that the underlying MDP has a *distinguishability* property, the regret can be of order $\log n$ [29]. A similar situation holds for Multi-Armed Bandits (MAB) [30].

To proceed, we define the following identifiability condition as the analogue of the distinguishability condition in MDPs. Intuitively, in the identifiable systems, the optimal feedback gain $L(\theta_0)$ can be learned *as accurately as* the closed-loop transition matrix $A_0 + B_0L(\hat{\theta}_t)$.

Definition 3 [18] *Assuming $\theta_0 \in \Gamma_0$ for some $\Gamma_0 \subset \mathbb{R}^{p \times q}$, θ_0 is called identifiable if*

$$\sup_{\theta_1, \theta_2 \in \Gamma_0} \frac{\|L(\theta_2) - L(\theta_0)\|}{\|(\theta_2 - \theta_0) \tilde{L}(\theta_1)\|} < \infty, \quad (13)$$

where $\tilde{L}(\theta) = [I_p, L(\theta)']'$.

Remark 5 *It has been established that the feedback difference $L(\theta) - L(\theta_0)$ is dominated by the parameter difference $\theta - \theta_0$, modulo a constant factor [17]:*

$$\sup_{\theta: \|\mathcal{K}(\theta)\| \leq C_K} \frac{\|L(\theta) - L(\theta_0)\|}{\|\theta - \theta_0\|} \leq C_L, \quad (14)$$

where $C_L < \infty$ is a fixed constant determined by θ_0 and $C_K < \infty$. Note that (13) is stronger than (14) since in general, the identifiability condition does not need to hold. Indeed, one can have $L(\theta_2) \neq L(\theta_0)$, but at the same time $\theta_2 \tilde{L}(\theta_1) = \theta_0 \tilde{L}(\theta_1)$ [18].

An extensive discussion of concrete examples of Γ_0 is provided in the work of Faradonbeh et al. [18]. In general, Γ_0 can be a manifold (or a finite union of manifolds) induced by different conditions such as the rank or the sparsity of the dynamics matrices. In the sequel, we briefly discuss four examples. The first condition is the support of the dynamics matrices. Indeed, letting $\mathcal{I} \subset \{(i, j) : 1 \leq i \leq p, 1 \leq j \leq q\}$, assume that θ belongs to Γ_0 , if and only if the coordinate (i, j) of θ is zero for all $(i, j) \notin \mathcal{I}$. If the cardinality of \mathcal{I} is not larger than p^2 , Γ_0 can satisfy (13) [18]. The second condition is based on sparsity of the true parameter. That is, if θ_0 is known to have at most p^2 non-zero entries, Γ_0 is contained in the finite union of the subspaces satisfying the aforementioned support condition. Third, the rank of θ_0 can be used in the design of the adaptive policies as follows. The set of $p \times q$ matrices of rank at most d is a finite union of manifolds of dimension at most $d(p + q - d)$ [31], which can satisfy (13) if $d(p + q - d) \leq p^2$ [18]. Finally, one can utilize the knowledge about a linear subspace that θ_0 belongs to. Suppose that θ_0 is known to belong to a subspace of $\mathbb{R}^{p \times q}$. If the above subspace is of dimension at most p^2 , it can satisfy the identifiability condition (13) [18].

The quantity p^2 appearing in the above examples is actually determined by the learning invariant manifold \mathcal{N} . Indeed, \mathcal{N} contains the dynamics matrices θ for which the closed-loop matrix is indistinguishable from the truth θ_0 ; that is, $\theta \tilde{L}(\theta_0) = \theta_0 \tilde{L}(\theta_0)$. It has been established that if $\Gamma_0 \cap \mathcal{N} = \{\theta_0\}$, then Γ_0 satisfies (13) [18]. Since \mathcal{N} is a linear subspace of dimension pr , we restricted the dimension of Γ_0 to $pq - pr = p^2$. More technical details can be found in existing studies on learning invariant manifolds, optimality level-set manifolds, and their effect on the performance of adaptive LQRs [32,33,18].

In order to modify Algorithm 1 for the systems with the side information Γ_0 , we replace (11) with $\gamma^m \bar{v}_m^2 < \bar{C}$; i.e., the perturbation signals are smaller (and diminish faster) than (11), and the covariance matrix does not need to be positive definite. Further, the uncertain dynamics parameter θ_0 will be learned through the following least squares procedure on the identifiability set Γ_0 :

$$\hat{\theta}_n = \arg \min_{\theta \in \Gamma_0} \sum_{t=0}^{n-1} \left\| x(t+1) - \theta \begin{bmatrix} x(t) \\ u(t) \end{bmatrix} \right\|^2. \quad (15)$$

All other parts of the adaptive regulator, such as the lengths of the epochs, remain the same as in Algorithm 1. We conclude this section with the next result that provides a smaller regret bound vis-a-vis that of Theorem 3.

Theorem 4 *Let θ_0 be identifiable, and $\hat{\pi}$ be the above adaptive regulator. Then, with probability at least $1 - \delta$,*

$$\sup_{n \geq n_0} \frac{\mathcal{R}_n(\hat{\pi})}{\log n \log^{1+2/\alpha}(n\delta^{-1})} \leq C_5,$$

where n_0 is defined in Theorem 3, and the fixed constant $C_5 < \infty$ depends on θ_0 .

5 Concluding Remarks

We analyzed the performance of adaptive control policies with perturbed input for multidimensional LQ systems. Indeed, we established non-asymptotic results for the high probability regret bound (Theorem 1), as well as the learning accuracy (Theorem 2), which hold uniformly over time. Leveraging the developed general theory, we showed that a greedy adaptive regulator with suitably designed input perturbation provides (nearly) square root regret with respect to time (Theorem 3), and outperforms previously available reinforcement learning algorithms in a number of ways.

Specifically, the presented perturbed greedy regulator does not need to know the spectral properties of the closed-loop transition matrix, and the noise statistics. Further, the computationally intractable steps being used by optimism-based adaptive policies are not required. Finally, input perturbation efficiently regulates the trajectory of the system by preventing fluctuations due to the stochastic nature of the dynamics. We also discussed the situations where the regret is of a logarithmic magnitude, assuming that the adaptive operator has access to some information (such as the support or the rank) of the unknown dynamics matrices (Theorem 4).

Extending the presented framework to the case of imperfect observations where a linear transformation of the state signal is observed with some noise is an interesting direction for future work. Moreover, addressing the

problem of adaptive regulation for large scale systems (e.g., networks) in a possibly high dimensional setting such as sparse or low rank dynamical models is of interest for further investigation. Finally, settings involving time varying dynamics (e.g., switching systems and Markov Jump Processes) constitute topics for future studies.

Acknowledgements

The authors appreciate the helpful comments of the reviewers. The work of M. K. S. F. was partially supported by the University of Florida Informatics Institute. A. T. acknowledges the support of NSF via CAREER grant IIS-1452099. G.M.'s work was supported in part by NSF grants DMS-1821220 and DMS-1632730.

References

- [1] T. Kailath, *Linear systems*. Prentice-Hall Englewood Cliffs, NJ, 1980, vol. 156.
- [2] S. Meyn, *Control techniques for complex networks*. Cambridge University Press, 2008.
- [3] M. K. S. Faradonbeh, A. Tewari, and G. Michailidis, "Optimality of fast matching algorithms for random networks with applications to structural controllability," *IEEE Transactions on Control of Network Systems*, vol. 4, no. 4, pp. 770–780, 2017.
- [4] R. Marino and P. Tomei, *Nonlinear control design: geometric, adaptive and robust*. Prentice Hall London, 1995, vol. 1.
- [5] C. Li and J. Lam, "Stabilization of discrete-time nonlinear uncertain systems by feedback based on ls algorithm," *SIAM Journal on Control and Optimization*, vol. 51, no. 2, pp. 1128–1151, 2013.
- [6] N. Lazic, T. Lu, C. Boutilier, M. Ryu, E. Wong, B. Roy, and G. Imwalle, "Data center cooling using model-predictive control." NIPS, 2018.
- [7] M. H. Pesaran and A. Timmermann, "Small sample properties of forecasts from autoregressive models under structural breaks," *Journal of Econometrics*, vol. 129, no. 1-2, pp. 183–217, 2005.
- [8] M. K. S. Faradonbeh, A. Tewari, and G. Michailidis, "Finite time identification in unstable linear systems," *Automatica*, vol. 96, pp. 342–353, 2018.
- [9] D. P. Bertsekas, *Dynamic programming and optimal control*. Athena Scientific Belmont, MA, 1995, vol. 1, no. 2.
- [10] P. Dorato, C. T. Abdallah, V. Cerone, and D. H. Jacobson, *Linear-quadratic control: an introduction*. Prentice Hall Englewood Cliffs, NJ, 1995.
- [11] L. Guo and H.-F. Chen, "The åstrom-wittenmark self-tuning regulator revisited and els-based adaptive trackers," *IEEE Transactions on Automatic Control*, vol. 36, no. 7, pp. 802–812, 1991.
- [12] T. L. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Advances in applied mathematics*, vol. 6, no. 1, pp. 4–22, 1985.
- [13] M. C. Campi and P. Kumar, "Adaptive linear quadratic gaussian control: the cost-biased approach revisited," *SIAM Journal on Control and Optimization*, vol. 36, no. 6, pp. 1890–1907, 1998.

- [14] S. Bittanti and M. C. Campi, “Adaptive control of linear time invariant systems: the “bet on the best” principle,” *Communications in Information & Systems*, vol. 6, no. 4, pp. 299–320, 2006.
- [15] Y. Abbasi-Yadkori and C. Szepesvári, “Regret bounds for the adaptive control of linear quadratic systems.” in *COLT*, 2011, pp. 1–26.
- [16] M. Ibrahimi, A. Javanmard, and B. V. Roy, “Efficient reinforcement learning for high dimensional linear quadratic systems,” in *Advances in Neural Information Processing Systems*, 2012, pp. 2636–2644.
- [17] M. K. S. Faradonbeh, A. Tewari, and G. Michailidis, “Optimism-based adaptive regulation of linear-quadratic systems,” *IEEE Transactions on Automatic Control*, to appear, available online: *arXiv:1711.07230*.
- [18] —, “On optimality of adaptive linear-quadratic regulators,” *arXiv preprint arXiv:1806.10749*, 2018.
- [19] O. C. Imer, S. Yüksel, and T. Başar, “Optimal control of LTI systems over unreliable communication links,” *Automatica*, vol. 42, no. 9, pp. 1429–1439, 2006.
- [20] Y. Ouyang, S. M. Asghari, and A. Nayyar, “Optimal infinite horizon decentralized networked controllers with unreliable communication,” *arXiv preprint arXiv:1806.06497*, 2018.
- [21] Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári, “Improved algorithms for linear stochastic bandits,” *Advances in Neural Information Processing Systems*, pp. 2312–2320, 2011.
- [22] J. A. Tropp, “User-friendly tail bounds for sums of random matrices,” *Foundations of computational mathematics*, vol. 12, no. 4, pp. 389–434, 2012.
- [23] M. K. S. Faradonbeh, A. Tewari, and G. Michailidis, “Input perturbations for adaptive control and learning,” *arXiv preprint arXiv:1811.04258*, 2018.
- [24] —, “Finite time adaptive stabilization of linear systems,” *IEEE Transactions on Automatic Control*, vol. 64, no. 8, pp. 3498–3505, 2019.
- [25] M. Green and J. B. Moore, “Persistence of excitation in linear systems,” in *American Control Conference, 1985*. IEEE, 1985, pp. 412–417.
- [26] M. Fazel, R. Ge, S. Kakade, and M. Mesbahi, “Global convergence of policy gradient methods for the linear quadratic regulator,” in *International Conference on Machine Learning*, 2018, pp. 1466–1475.
- [27] M. Abeille and A. Lazaric, “Improved regret bounds for thompson sampling in linear quadratic control problems,” in *International Conference on Machine Learning*, 2018, pp. 1–9.
- [28] C. Jin, Z. Allen-Zhu, S. Bubeck, and M. I. Jordan, “Is q-learning provably efficient?” *arXiv preprint arXiv:1807.03765*, 2018.
- [29] A. Tewari and P. L. Bartlett, “Optimistic linear programming gives logarithmic regret for irreducible mdps,” in *Advances in Neural Information Processing Systems*, 2008, pp. 1505–1512.
- [30] H. Bastani, M. Bayati, and K. Khosravi, “Mostly exploration-free algorithms for contextual bandits,” *arXiv preprint arXiv:1704.09011*, 2017.
- [31] U. Shalit, D. Weinshall, and G. Chechik, “Online learning in the embedded manifold of low-rank matrices,” *Journal of Machine Learning Research*, vol. 13, no. Feb, pp. 429–458, 2012.
- [32] J. W. Polderman, “On the necessity of identifying the true parameter in adaptive LQ control,” *Systems & control letters*, vol. 8, no. 2, pp. 87–91, 1986.
- [33] —, “A note on the structure of two subsets of the parameter space in adaptive control problems,” *Systems & control letters*, vol. 7, no. 1, pp. 25–34, 1986.