

Non-stationary Reinforcement Learning without Prior Knowledge: An Optimal Black-box Approach

Chen-Yu Wei

Haipeng Luo

University of Southern California

CHENYU.WEI@USC.EDU

HAIPENGL@USC.EDU

Editors: Mikhail Belkin and Samory Kpotufe

Abstract

We propose a black-box reduction that turns a certain reinforcement learning algorithm with optimal regret in a (near-)stationary environment into another algorithm with optimal dynamic regret in a non-stationary environment, importantly without any prior knowledge on the degree of non-stationarity. By plugging different algorithms into our black-box, we provide a list of examples showing that our approach not only recovers recent results for (contextual) multi-armed bandits achieved by very specialized algorithms, but also significantly improves the state of the art for (generalized) linear bandits, episodic MDPs, and infinite-horizon MDPs in various ways. Specifically, in most cases our algorithm achieves the optimal dynamic regret $\tilde{O}(\min\{\sqrt{LT}, \Delta^{1/3}T^{2/3}\})$ where T is the number of rounds and L and Δ are the number and amount of changes of the world respectively, while previous works only obtain suboptimal bounds and/or require the knowledge of L and Δ .

1. Introduction

Most existing works on reinforcement learning consider a stationary environment and aim to find or be comparable to an optimal policy (known as having low *static regret*). In many applications, however, the environment is far from being stationary. In these cases, it is much more meaningful to minimize *dynamic regret*, the gap between the total reward of the optimal *sequence* of policies and that of the learner. Indeed, there is a surge of studies on this topic recently (Jaksch et al., 2010; Gajane et al., 2018; Li and Li, 2019; Ortner et al., 2020; Cheung et al., 2020; Fei et al., 2020; Domingues et al., 2021; Mao et al., 2021; Zhou et al., 2020; Touati and Vincent, 2020).

One common issue of all these works, however, is that their algorithms crucially rely on having some prior knowledge on the *degree of non-stationarity* of the world, such as how much or how many times the distribution changes, which is often unavailable in practice. Cheung et al. (2020) develop a Bandit-over-Reinforcement-Learning (BoRL) framework to relax this assumption, but it introduces extra overhead and leads to suboptimal regret. Indeed, as discussed in their work, there are multiple aspects (which they call *endogeneity*, *exogeneity*, *uncertainty*, and *bandit feedback*) combined in non-stationary reinforcement learning that make the problem highly challenging.

For bandit problems, the special case of reinforcement learning, the works of Auer et al. (2019) and Chen et al. (2019) are the first to achieve near-optimal dynamic regret without any prior knowledge on the degree of non-stationarity. The same technique has later been adopted by Chen et al. (2020) for the case of combinatorial semi-bandits. Their algorithms maintain a distribution over arms (or policies/super-arms in the contextual/combinatorial case (Chen et al., 2019, 2020)) with properly controlled variance for all reward estimators. This approach is generally incompatible with standard reinforcement learning algorithms, which are usually built upon the *optimism in the face*

of uncertainty principle and do not maintain a distribution over policies (see also (Lykouris et al., 2021; Wang et al., 2020) for related discussions). Another drawback is that their algorithms are very specialized to their problems, and it is highly unclear whether the ideas can be extended to other problems.

In this work, we address all these issues and make significant progress in this direction. Specifically, we propose a *general* approach that is applicable to various reinforcement learning settings (including bandits, episodic MDPs, infinite-horizon MDPs, etc.) and achieves *optimal dynamic regret without any prior knowledge* on the degree of non-stationarity. Our approach, called MASTER, is a *black-box reduction* that turns any algorithm with optimal performance in a (near-)stationary environment and additionally some mild requirements into another algorithm with optimal dynamic regret in a non-stationary environment, again, without the need of any prior knowledge. For example, all existing UCB-based algorithms satisfy the conditions of our reduction and are readily to be plugged into our black-box.

Applications and comparisons To showcase the versatility of our approach, we provide a list of examples by considering different settings and applying our reduction with different base algorithms. These examples, summarized in Table 1, recover the results of Auer et al. (2019) and Chen et al. (2019) for (contextual) multi-armed bandits, and more importantly, improve the best known results for (generalized) linear bandits, episodic MDPs, and infinite-horizon MDPs in various ways. More specifically, let L and Δ be the number and amount of changes of the environment respectively (see Section 2.1 for formal definition). For all settings except infinite-horizon MDPs, ignoring other parameters, our algorithms achieve dynamic regret $\min\{\text{Reg}_L^*, \text{Reg}_\Delta^*\}$ without knowing L and Δ , where $\text{Reg}_L^* = \sqrt{LT}$, $\text{Reg}_\Delta^* = \Delta^{1/3}T^{2/3} + \sqrt{T}$, and T is the number of rounds. These bounds are known to be optimal *even when L and Δ are known*, and they improve over (Cheung et al., 2019; Russac et al., 2019; Kim and Tewari, 2020; Zhao et al., 2020; Zhao and Zhang, 2021) for linear bandits, (Russac et al., 2020; Fauray et al., 2021) for generalized linear bandits, (Mao et al., 2021) for episodic tabular MDPs, and (Touati and Vincent, 2020; Zhou et al., 2020) for episodic linear MDPs. For infinite-horizon MDPs, we achieve the same optimal regret when the maximum diameter of the MDPs is known, or when L and Δ are known, improving over the best existing results by (Gajane et al., 2018) and (Cheung et al., 2020). When none of them is known, we can still adopt the BoRL technique (Cheung et al., 2020) with the price of paying extra $T^{3/4}$ regret, which is suboptimal but still outperforms best known results.

In particular, we emphasize that achieving dynamic regret Reg_L^* beyond (contextual) multi-armed bandits is one notable breakthrough we make. Indeed, *even when L is known*, previous approaches based on restarting after a fixed period, a sliding window with a fixed size, or discounting with a fixed discount factor, all lead to a suboptimal bound of $\tilde{O}(L^{1/3}T^{2/3})$ at best (Gajane et al., 2018). Since this bound is subsumed by Reg_Δ^* , related discussions are also often omitted in previous works.

We also emphasize that when dealing with problems with linear structures (including linear bandits, generalized linear bandits, and linear MDPs), our bounds Reg_Δ^* is also new *even when Δ is known*. Indeed, although several existing works on non-stationary linear bandits (Cheung et al., 2019; Russac et al., 2019; Kim and Tewari, 2020; Zhao et al., 2020) claim that their algorithms achieve the bound Reg_Δ^* , there is in fact a technical flaw in all of them, as explained and corrected recently in (Zhao and Zhang, 2021; Touati and Vincent, 2020). After the correction, their bounds all deteriorate to $\Delta^{1/4}T^{3/4} + \sqrt{T}$, which is no longer near-optimal. On the other hand, our approach not

Table 1: A summary of our results and comparisons with the state-of-the-art. Our algorithms are named in the form of “MASTER + X ” where X is the base algorithm used in our reduction. Here, $\text{Reg}_L^* = \sqrt{LT}$ and $\text{Reg}_\Delta^* = \Delta^{1/3}T^{2/3} + \sqrt{T}$, where T is the number of rounds and L and Δ are the number and amount of changes of the world respectively (dependence on other parameters is omitted). D_{\max} is the maximum diameters of the MDPs.

| Setting | Algorithm | Regret in $\tilde{\mathcal{O}}(\cdot)$ | Required knowledge |
|--|-----------------------------------|---|--------------------|
| Multi-armed bandits | (Auer et al., 2019) | Reg_L^* | |
| | MASTER + UCB1 | $\min\{\text{Reg}_L^*, \text{Reg}_\Delta^*\}$ | |
| Contextual bandits | (Chen et al., 2019) | | |
| | MASTER + ILTCB MASTER + FALCON | $\min\{\text{Reg}_L^*, \text{Reg}_\Delta^*\}$ | |
| Linear bandits | (Cheung et al., 2019) | $\Delta^{1/4}T^{3/4} + \sqrt{T}$ | Δ |
| | (Cheung et al., 2019) | $\Delta^{1/4}T^{3/4} + T^{3/4}$ | |
| | MASTER + OFUL | $\min\{\text{Reg}_L^*, \text{Reg}_\Delta^*\}$ | |
| Generalized linear bandits | (Russac et al., 2020) | $L^{1/3}T^{2/3}$ | L |
| | (Fauray et al., 2021) | $\Delta^{1/5}T^{4/5} + T^{3/4}$ | |
| | MASTER + GLM-UCB | $\min\{\text{Reg}_L^*, \text{Reg}_\Delta^*\}$ | |
| Episodic MDPs (tabular case) | (Mao et al., 2021) | Reg_Δ^* | Δ |
| | MASTER + Q-UCB | $\min\{\text{Reg}_L^*, \text{Reg}_\Delta^*\}$ | |
| Episodic MDPs (linear case) | (Touati and Vincent, 2020) | $\Delta^{1/4}T^{3/4} + \sqrt{T}$ | Δ |
| | MASTER + LSVI-UCB | $\min\{\text{Reg}_L^*, \text{Reg}_\Delta^*\}$ | |
| Infinite-horizon communicating MDPs (tabular case) | (Gajane et al., 2018) | $L^{1/3}T^{2/3}$ | L |
| | (Cheung et al., 2020) | $\Delta^{1/4}T^{3/4} + \sqrt{T}$ | Δ |
| | (Cheung et al., 2020) | $\Delta^{1/4}T^{3/4} + T^{3/4}$ | |
| | MASTER + UCRL | Reg_L^* or Reg_Δ^* | L or Δ |
| | MASTER + UCRL | $\min\{\text{Reg}_L^*, \text{Reg}_\Delta^*\}$ | D_{\max} |
| | MASTER + UCRL + BoRL | $\min\{\text{Reg}_L^*, \text{Reg}_\Delta^*\} + T^{3/4}$ | |

only sidesteps the difficulty of getting the optimal bound met in all these works, but also avoids the requirement of knowing Δ . Similar technical difficulties also appear in generalized linear bandits (Russac et al., 2020; Fauray et al., 2021) and linear MDPs (Touati and Vincent, 2020; Zhou et al., 2020), and our approach overcomes them similarly.

High-level ideas The high-level idea of our reduction is to schedule multiple instances of the base algorithm with different durations in a carefully-designed randomized scheme, which facilitates non-stationarity detection with little overhead. A related and well-known approach for non-stationary environments is to maintain multiple instances of a base algorithm with different parameter tunings or different starting points and to learn the best of them via another “expert” algorithm, which can be very successful when learning with full information (Hazan and Seshadhri, 2007; Luo

and Schapire, 2015; Daniely et al., 2015; Jun et al., 2017) but is suboptimal and has many limitations when learning with partial information (Luo et al., 2018; Cheung et al., 2019, 2020). Our approach is different as we do not try to learn the best instance; instead, we always follow the decision suggested by the instance with the currently shortest scheduled duration, and also only update this instance after receiving feedback from the environment. This is because base algorithms with shorter duration are responsible for detecting larger distribution changes, and always following the shortest one ensure that it is not blocked by the longer ones and thus every scale of distribution change is detected in a timely manner.

Another related approach is *regret balancing*, developed recently for model selection in bandit problems (Abbasi-Yadkori et al., 2020; Pacchiano et al., 2020). The idea is also to run multiple base algorithms in parallel, each with a putative regret upper bound. The learner executes one of them in each round which incurs the least regret so far, and also constantly compares the performance among base algorithms, eliminating those whose putative regret bounds are violated. While our algorithm resembles regret balancing in some aspects, the way it chooses the base algorithm in each round is clearly quite different, which is also crucial for our problem.

Other related work There are also a series of works on learning MDPs with adversarial rewards and a *fixed* transition (Even-Dar et al., 2009; Neu et al., 2010; Arora et al., 2012; Neu et al., 2012; Dekel and Hazan, 2013; Neu et al., 2013; Zimin and Neu, 2013; Dick et al., 2014; Rosenberg and Mansour, 2019; Cai et al., 2020; Jin et al., 2020a; Shani et al., 2020; Rosenberg and Mansour, 2020; Lee et al., 2020; Jin and Luo, 2020; Chen et al., 2021; Lancewicki et al., 2020). These models can potentially handle non-stationarity in the reward function but not the transition kernel (in fact, most of these works also only consider static regret). Lykouris et al. (2021) investigate an episodic MDP setting where an adversary can corrupt both the reward and the transition for up to L' episodes, and achieve dynamic regret $\tilde{\mathcal{O}}(\min\{L'\sqrt{T}, L'/\text{gap}\})$ without knowing L' , where *gap* is the minimal suboptimality gap and could be arbitrarily small. Since corruption of up to L' episodes implies that the world changes at most $L = 2L'$ times, our result improves theirs from $\tilde{\mathcal{O}}(L'\sqrt{T})$ to $\tilde{\mathcal{O}}(\sqrt{LT})$ when $1/\text{gap} > \sqrt{T}$. On the other hand, it is possible that L is much smaller than L' (e.g. $L = \Theta(1)$ while $L' = \Theta(T)$), in which case our results are also significantly better.

2. Problem Setting, Main Results, and High-level Ideas

Throughout the paper, we fix a probability parameter δ of order $1/\text{poly}(T)$, and write $h_1(x) = \tilde{\mathcal{O}}(h_2(x))$ or $h_2(x) = \tilde{\Omega}(h_1(x))$ if $h_1(x) = \mathcal{O}(\text{poly}(\log(T/\delta))h_2(x))$. We say “with high probability, $h_1 = \tilde{\mathcal{O}}(h_2(x))$ ” if “with probability $1 - \delta$, $h_1 = \mathcal{O}(h_2(x))$ ”. For an integer n , we denote the set $\{1, 2, \dots, n\}$ by $[n]$; and for integers s and e , we denote the set $\{s, s + 1, \dots, e\}$ by $[s, e]$.

2.1. Problem setting

We consider the following general reinforcement learning (RL) framework that covers a wide range of problems. Ahead of time, the learner is given a policy set Π , and the environment decides T reward functions $f_1, \dots, f_T : \Pi \rightarrow [0, 1]$ unknown to the learner. Then, in each round $t = 1, \dots, T$, the learner chooses a policy $\pi_t \in \Pi$ and receives a noisy reward $R_t \in [0, 1]$ whose mean is $f_t(\pi_t)$.¹ The dynamic regret of the learner is defined as $\text{D-REG} = \sum_{t=1}^T (f_t^* - R_t)$, where $f_t^* = \max_{\pi \in \Pi} f_t(\pi)$ is the expected reward of the optimal policy for round t .

1. The range $[0, 1]$ is only for simplicity. Our results can be directly extended to the case with sub-Gaussian noise.

Many heavily-studied problems fall into this framework. For example, in the classic multi-armed bandit problem (Lai and Robbins, 1985), it suffices to treat each arm as a policy; for finite-horizon episodic RL (e.g. (Jin et al., 2018)), each state-to-action mapping is considered as a policy, and $f_t(\pi)$ is the expected reward of executing π in the t -th episode’s MDP with some transition kernel and some reward function. See more examples in Appendix I. Note that our framework ignores many details of the actual problem we are trying to solve (e.g. not even mentioning the MDPs for RL). This is because our results only rely on certain guarantees provided by a base algorithm, making these details irrelevant to our presentation. There is also some technicality to fit the infinite-horizon RL problem into our framework, which we will discuss in detail in Section 4.

Non-stationarity measure A natural way to measure the distribution drift between rounds t and $t + 1$ is to see how much the expected reward of any policy could change, that is, $\max_{\pi \in \Pi} |f_t(\pi) - f_{t+1}(\pi)|$. However, to make our results more general, we take a slightly more abstract way to define non-stationarity whose exact form eventually depends on what guarantees the base algorithm can provide for a concrete problem. To this end, we define the following.

Definition 1 $\Delta : [T] \rightarrow \mathbb{R}$ is a non-stationarity measure if it satisfies $\Delta(t) \geq \max_{\pi \in \Pi} |f_t(\pi) - f_{t+1}(\pi)|$ for all t . Define for any interval $\mathcal{I} = [s, e]$, $\Delta_{\mathcal{I}} = \sum_{\tau=s}^{e-1} \Delta(\tau)$ (note $\Delta_{[s,s]} = 0$) and $L_{\mathcal{I}} = 1 + \sum_{\tau=s}^{e-1} \mathbb{1}[\Delta(\tau) \neq 0]$. With slight abuse of notation, we write $\Delta = \Delta_{[1,T]}$ and $L = L_{[1,T]}$.

Base algorithm and requirements As mentioned, our approach takes a base algorithm that tackles the problem when the environment is (near-)stationary, and turns it into another algorithm that can deal with non-stationary environments. Throughout the paper, we denote the base algorithm by ALG and assumes that it satisfies the following mild requirements when run alone.

Assumption 1 ALG outputs an auxiliary quantity $\tilde{f}_t \in [0, 1]$ at the beginning of each round t . There exist a non-stationarity measure Δ and a non-increasing function $\rho : [T] \rightarrow \mathbb{R}$ such that running ALG satisfies the following: for all $t \in [T]$, as long as $\Delta_{[1,t]} \leq \rho(t)$, without knowing $\Delta_{[1,t]}$ ALG ensures with probability at least $1 - \frac{\delta}{T}$:

$$\tilde{f}_t \geq \min_{\tau \in [1,t]} f_{\tau}^* - \Delta_{[1,t]} \quad \text{and} \quad \frac{1}{t} \sum_{\tau=1}^t (\tilde{f}_{\tau} - R_{\tau}) \leq \rho(t) + \Delta_{[1,t]}. \quad (1)$$

Furthermore, we assume that $\rho(t) \geq \frac{1}{\sqrt{t}}$ and $C(t) = t\rho(t)$ is a non-decreasing function.

We unpack the meaning of this assumption and explain why this is a mild requirement via a few remarks below, followed by examples of existing algorithms that do satisfy our assumption.

First, consider choosing $\Delta(t) = \max_{\pi \in \Pi} |f_t(\pi) - f_{t+1}(\pi)|$ and see what the assumption means for a stationary environment with $f_t = f$ and $\Delta(t) = 0$ for all t . In this case, Eq. (1) simply becomes $\tilde{f}_t \geq \max_{\pi \in \Pi} f(\pi)$ and $\sum_{\tau=1}^t (\tilde{f}_{\tau} - R_{\tau}) \leq C(t)$, which are standard properties of Upper-Confidence-Bound (UCB)-based algorithms, where \tilde{f}_t is an optimistic estimator of the optimal reward and $C(t)$ is the regret bound usually of order \sqrt{t} . In fact, even for non-UCB-based algorithms that do not explicitly maintain optimistic estimators, by looking into their analysis, it is still possible to extract a quantity \tilde{f}_t satisfying these two properties (see our example for contextual bandits in Appendix I). We also note that this requirement for the special stationary case is in fact all we need to achieve our claimed regret bound Reg_L^* .

Second, to simultaneously achieve the regret bound Reg_Δ^* as well, we require *slightly more* from the base algorithm: in a *near-stationary* environment with $\Delta_{[1,t]} \leq \rho(t)$, the two aforementioned properties still hold approximately with degradation $\Delta_{[1,t]}$ (that is, Eq. (1)).² We call this a near-stationary environment because $\Delta_{[1,t]}$ can be of order $\Theta(t)$ in a highly non-stationary environment, while here we restrict it to be at most $\rho(t)$, which is non-increasing in t (and in fact of order $1/\sqrt{t}$ in all our examples). To the best of our knowledge, all UCB-based algorithms satisfy [Assumption 1](#) with some suitable choice of Δ . The fact that we only require [Eq. \(1\)](#) to hold for near-stationary environments is the key to bypassing the technical difficulty of getting the optimal bound Reg_Δ^* met in ([Cheung et al., 2019](#); [Russac et al., 2019](#); [Zhao et al., 2020](#); [Russac et al., 2020](#); [Fauray et al., 2021](#); [Touati and Vincent, 2020](#); [Zhou et al., 2020](#)) for linear bandits, generalized linear bandits, and linear MDPs, as mentioned in [Section 1](#).

Finally, noting that $\rho(t)$ and $C(t)$ represent an average and an cumulative regret bound respectively, the monotonicity requirement on them is more than natural. The requirement $\rho(t) \geq \frac{1}{\sqrt{t}}$ is also usually unavoidable without further structures in the problem. Note that while we write ρ and C as a function of t only, they can depend on $\log(1/\delta)$, $\log T$, the complexity of Π , and other problem-dependent parameters such as the number of states/actions of an MDP.

Following the order in [Table 1](#), we now give a list of existing algorithms that satisfy [Assumption 1](#) in different problem settings with proper non-stationarity measure Δ and regret bound C . We defer the concrete form of \tilde{f}_t (which requires introducing other notations) and all the proofs to [Appendix I](#).

- UCB1 ([Auer et al., 2002a](#)): $C(t) = \tilde{O}(\sqrt{At} + A)$ and $\Delta(t) = \Theta(\|r_t - r_{t+1}\|_\infty)$, where A is the number of arms, and r_t is the expected reward vector at time t .
- ILTCB ([Agarwal et al., 2014](#), short for ILOVETOCONBANDITS): $C(t) = \tilde{O}(\sqrt{At \log |\Pi|} + A \log |\Pi|)$ and $\Delta(t) = \Theta(\int_r \int_x |\mathcal{D}_t(x, r) - \mathcal{D}_{t+1}(x, r)| dx dr)$, where A is the number of actions and \mathcal{D}_t is the joint distribution of the context-reward pair (x, r) at time t .
- FALCON ([Simchi-Levi and Xu, 2020](#)): $C(t) = \tilde{O}(\sqrt{At \log |\Phi|} + A \log |\Phi|)$ and $\Delta(t) = \Theta(\sqrt{A} \max_{x,a} |\phi_t^*(x, a) - \phi_{t+1}^*(x, a)| + \int_x |\mathcal{D}_t(x) - \mathcal{D}_{t+1}(x)| dx)$, where A is the number of actions, Φ is the set of regressors (each of which maps a context-action pair to a predicted reward), $\phi_t^* \in \Phi$ is the true regressor at time t , and \mathcal{D}_t is the distribution of contexts at time t .
- OFUL ([Abbasi-Yadkori et al., 2011](#)): $C(t) = \tilde{O}(d\sqrt{t})$ and $\Delta(t) = \tilde{\Theta}(d\|\theta_t - \theta_{t+1}\|_2)$, where d is the feature dimension and $\theta_t \in \mathbb{R}^d$ parameterizes the linear reward function at time t .
- GLM-UCB ([Filippi et al., 2010](#)): $C(t) = \tilde{O}\left(\frac{k_\mu d}{c_\mu} \sqrt{t}\right)$ and $\Delta(t) = \tilde{\Theta}\left(\frac{k_\mu d}{c_\mu} \|\theta_t - \theta_{t+1}\|_2\right)$, where d is the feature dimension, $\theta_t \in \mathbb{R}^d$ parameterizes the linear reward function at time t , and k_μ, c_μ are the upper and lower bounds of the gradient of the link function.
- Q-UCB ([Jin et al., 2018](#), short for Q-learning UCB-H):³ $C(t) = \tilde{O}(\sqrt{H^5 SA t} + H^3 SA)$ and $\Delta(t) = \Theta(H \sum_{h=1}^H \max_{s,a} |r_h^t(s, a) - r_h^{t+1}(s, a)| + H^2 \sum_{h=1}^H \max_{s,a} \|p_h^t(\cdot|s, a) - p_h^{t+1}(\cdot|s, a)\|_1)$, where H, S and A are the numbers of layers, states, and actions of the MDP respectively, and p_h^t and r_h^t are the transition and reward functions for layer h of episode t .

2. We use $\min_{\tau \in [1,t]} f_\tau^*$ instead of the more natural one f_t^* since the former is weaker and the difference between these two is at most $\Delta_{[1,t]}$ anyway.

3. For ease of comparison, here, the reward range is changed from $[0, 1]$ to the more common $[0, H]$.

- LSVI-UCB (Jin et al., 2020b):⁴ $C(t) = \tilde{\mathcal{O}}(\sqrt{d^3 H^4 t})$ and $\Delta(t) = \tilde{\mathcal{O}}(dH \sum_{h=1}^H \|\theta_h^t - \theta_h^{t+1}\|_2 + dH^2 \sum_{h=1}^H \|\mu_h^t - \mu_h^{t+1}\|_F)$, where d is the feature dimension, H is the number of layers, and θ_h^t and μ_h^t are the parameters of the linear MDP for layer h of episode t .

2.2. Main results

Our main result is that, with an algorithm satisfying [Assumption 1](#) at hand, our proposed black-box reduction, MASTER ([Algorithm 3](#)), ensures the following dynamic regret bound.

Theorem 2 *If [Assumption 1](#) holds with $C(t) = c_1 t^p + c_2$ for some $p \in [\frac{1}{2}, 1)$ and $c_1, c_2 > 0$, then MASTER ([Algorithm 3](#)), without knowing L and Δ , guarantees with high probability:*

$$\text{D-REG} = \tilde{\mathcal{O}} \left(\min \left\{ \left(c_1 + \frac{c_2}{c_1} \right) \sqrt{LT}, \left(c_1^{2/3} + c_2 c_1^{-4/3} \right) \Delta^{1/3} T^{2/3} + \left(c_1 + \frac{c_2}{c_1} \right) \sqrt{T} \right\} \right)$$

when $p = \frac{1}{2}$, and $\text{D-REG} = \tilde{\mathcal{O}} \left(\min \left\{ c_1 L^{1-p} T^p, \left(c_1 \Delta^{1-p} T \right)^{\frac{1}{2-p}} + c_1 T^p \right\} \right)$ when $p > \frac{1}{2}$ (omitting some lower-order terms).

For ease of presentation, in this theorem we assume that $C(\cdot)$ takes a certain form that is common in the literature and holds for all our examples with $p = \frac{1}{2}$. Applying this theorem to all the examples discussed earlier, we achieve all the optimal $\min\{\text{Reg}_L^*, \text{Reg}_\Delta^*\}$ bounds listed in [Table 1](#) (except for infinite-horizon MDPs which will be discussed in [Section 4](#)). Our definitions of L are the same as in previous works, and our definitions of Δ are sometimes larger by some problem-dependent factors (such as d and H) in order to fit [Assumption 1](#). More specifically, for (contextual) bandits, our MASTER combined with UCB1 and ILTCB recovers the same optimal bounds (in terms of all parameters) achieved by ([Auer et al., 2019](#); [Chen et al., 2019](#)). MASTER with FALCON obtains a similar bound as in ([Chen et al., 2019](#)) but with a different definition of Δ specific to the regressor setting. For other settings, we present our results in terms of the common definition of the non-stationarity measure (denoted by $\hat{\Delta}$) and compare them with the state of the art:

- MASTER + OFUL: $\text{D-REG} = \tilde{\mathcal{O}}(\min\{d\sqrt{LT}, d\hat{\Delta}^{1/3}T^{2/3} + d\sqrt{T}\})$, where $\hat{\Delta} = \sum_t \|\theta_t - \theta_{t+1}\|_2$. This improves ([Cheung et al., 2019](#); [Russac et al., 2019](#); [Kim and Tewari, 2020](#); [Zhao et al., 2020](#); [Zhao and Zhang, 2021](#)) which get $\tilde{\mathcal{O}}(d^{7/8} \hat{\Delta}^{1/4} T^{3/4} + d\sqrt{T})$ when $\hat{\Delta}$ is known.
- MASTER + GLM-UCB: $\text{D-REG} = \tilde{\mathcal{O}} \left(\min \left\{ \frac{k_\mu}{c_\mu} d\sqrt{LT}, \frac{k_\mu^{4/3}}{c_\mu} d\hat{\Delta}^{1/3} T^{2/3} + \frac{k_\mu}{c_\mu} d\sqrt{T} \right\} \right)$, where $\hat{\Delta} = \sum_t \|\theta_t - \theta_{t+1}\|_2$. This improves ([Russac et al., 2020](#)) which gets $\tilde{\mathcal{O}}(\frac{k_\mu}{c_\mu} d^{2/3} L^{1/3} T^{2/3})$ when L is known, and ([Fauray et al., 2021](#)) which gets $\tilde{\mathcal{O}}(\frac{k_\mu}{c_\mu} d^{9/10} \hat{\Delta}^{1/5} T^{4/5})$.
- MASTER + Q-UCB: $\text{D-REG} = \tilde{\mathcal{O}}(\min\{\sqrt{H^5 SALT}, (H^7 SA\hat{\Delta})^{1/3} T^{2/3} + \sqrt{H^5 SAT}\})$, where $\hat{\Delta} = \sum_{t,h} \max_{s,a} (|r_h^t(s,a) - r_h^{t+1}(s,a)| + \|p_h^t(\cdot|s,a) - p_h^{t+1}(\cdot|s,a)\|_1)$.⁵ ([Mao et al., 2021](#), Theorem 3) gets $\tilde{\mathcal{O}}((H^5 SA\hat{\Delta})^{1/3} T^{2/3} + \sqrt{H^3 SAT})$ when $\hat{\Delta}$ is known.⁶

4. Same as [Footnote 3](#).

5. Due to the scaling mentioned in [Footnote 3](#), here, we first scale down $C(\cdot)$ and Δ by an H factor, then apply [Theorem 2](#), and finally scale up the final bound by an H factor.

6. The bound reported in ([Mao et al., 2021](#)) is $(H^3 SA\hat{\Delta})^{1/3} T^{2/3} + \sqrt{H^2 SAT}$; however, their T is the total number of timesteps while our T is the number of episodes, and we have performed a proper translation between notations here. Their bound has a better H dependency thanks to the use of Freedman-style confidence bounds. The same idea unfortunately does not improve our bound due to the lower-order term c_2 in the definition of $C(t)$.

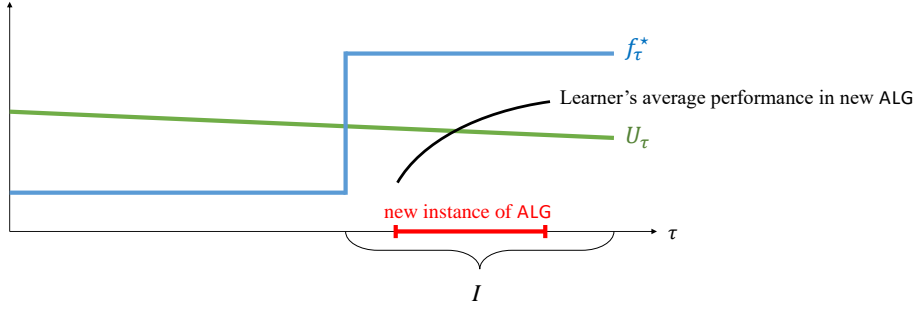


Figure 1: An illustration of how we detect non-stationarity via multiple instances of ALG

- MASTER + LSVI-UCB: D-REG = $\tilde{O}(\min\{\sqrt{d^3 H^4 L T}, (d^4 H^6 \hat{\Delta})^{1/3} T^{2/3} + \sqrt{d^3 H^4 T}\})$, where $\hat{\Delta} = \sum_{t,h} (\|\theta_h^t - \theta_h^{t+1}\|_2 + \|\mu_h^t - \mu_h^{t+1}\|_F)$. This improves (Zhou et al., 2020; Touati and Vincent, 2020) which get $\tilde{O}((d^5 H^8 \hat{\Delta})^{1/4} T^{3/4} + \sqrt{d^3 H^4 T})$ when $\hat{\Delta}$ is known.⁷

2.3. High-level ideas

To get a high-level idea of our approach, first consider what could go wrong when running ALG alone in a non-stationary environment and how to fix that intuitively. Decompose the dynamic regret as follows:

$$\underbrace{\sum_{\tau=1}^t (f_{\tau}^* - \tilde{f}_{\tau})}_{\mathbf{term}_1} + \underbrace{\sum_{\tau=1}^t (\tilde{f}_{\tau} - R_{\tau})}_{\mathbf{term}_2}. \quad (2)$$

As mentioned, in a stationary environment, ALG ensures that **term**₁ is simply non-positive and **term**₂ is bounded by $C(t)$ directly. In a non-stationary environment, however, both terms can be substantially larger. If we can detect the event that either of them is abnormally large, we know that the environment has changed substantially, and should just restart ALG. This detection can be easily done for **term**₂ since both \tilde{f}_{τ} and R_{τ} are observable, but not for **term**₁ since f_{τ}^* is of course unknown. Note that, a large **term**₁ implies that a policy, possibly suboptimal in the past, now becomes the optimal one with a much larger reward. A single instance of ALG run from the beginning thus cannot detect this because suboptimal policies are naturally selected very infrequently.

To address this issue, our main idea is to maintain different instances of ALG to facilitate non-stationarity detection, illustrated via an example in Figure 1. Here, there is one distribution change that happens in interval \mathcal{I} , making the value of f_{τ}^* (the blue curve) drastically increase. If within this interval, we start running another instance of ALG (the red interval), then its performance (the black curve) will gradually approach f_{τ}^* due to its regret guarantee in a stationary environment. Hypothetically, if another instance of ALG run from the beginning could coexist with this new instance, we would see that the latter significantly outperforms the former and infer that the environment has changed. The issue is that we cannot have multiple instances running and making decisions simultaneously, and here is where the optimistic estimators \tilde{f}_{τ} 's can help. Specifically, since the quantity $U_{\tau} = \min_{s \leq \tau} \tilde{f}_s$ (the green non-increasing curve) should always be an upper bound of the learner's

⁷ The same scaling as in Footnote 5 and Footnote 6 has been performed here.

Procedure 1: A procedure that randomly schedules ALG of different lengths within 2^n rounds

input: $n, \rho(\cdot)$
for $\tau = 0, \dots, 2^n - 1$ **do**
 for $m = n, n - 1, \dots, 0$ **do**
 If τ is a multiple of 2^m , with probability $\frac{\rho(2^n)}{\rho(2^m)}$, schedule a new instance alg of ALG that starts at $alg.s = \tau + 1$ and ends at $alg.e = \tau + 2^m$.
 end
end

Algorithm 2: MALG (Multi-scale ALG)

input: $n, \rho(\cdot)$
Initialization: run [Procedure 1](#) with inputs n and $\rho(\cdot)$.
At each time t , let the unique active instance be alg , output \tilde{g}_t (which is the \tilde{f}_t output by alg), follow alg 's decision π_t , and update alg after receiving feedback from the environment.

performance in a stationary environment, if we find that the new instance of ALG significantly outperforms this quantity at some point (as shown in [Figure 1](#)), we can also infer that the environment has changed, and prevent $\mathbf{term}_1 \leq \sum_{\tau=1}^t (f_\tau^* - U_\tau)$ from growing too large by restarting.

To formally implement the ideas above, we need to decide when to start a new instance, how long it should last, which instance should be active if multiple exist, and others. In [Section 3](#), we propose a randomized multi-scale scheme to do so, which is reminiscent of the ideas of sampling obligation in ([Auer et al., 2019](#)) and replay phase in ([Chen et al., 2019](#)), although their mechanisms are highly specific to their algorithms and problems.

3. Algorithm

In this section, we first introduce MALG, an algorithm that schedules and runs multiple instances of the base algorithm ALG in a multi-scale manner ([Section 3.1](#)). Then, equipping MALG with non-stationarity detection, we introduce our final black-box reduction MASTER ([Section 3.2](#)).

3.1. MALG: Running the Base Algorithm with Multiple Scales

We always run MALG for an interval of length 2^n , which we call a *block*, for some integer n (unless it is terminated by the non-stationarity detection mechanism). During initialization, MALG uses [Procedure 1](#) to schedule multiple instances of ALG within the block in the following way: for every $m = n, n - 1, \dots, 0$, partition the block equally into 2^{n-m} sub-intervals of length 2^m , and for each of these sub-intervals, with probability $\frac{\rho(2^n)}{\rho(2^m)} \leq 1$ schedule an instance of ALG (otherwise skip this sub-interval). We call these instances of length 2^m *order- m* instances.

Note that by definition there is always an order- n instance covering the entire block. We use alg to denote a particular instance of ALG, and use $alg.s$ and $alg.e$ to denote its start and end time.

After the initialization, MALG starts interacting with the environment as follows. In each time t , the unique instance covering this time step with the shortest length is considered as being *active*, while all others are *inactive*. MALG follows the decision of the active instance, and update it after receiving feedback from the environment. All inactive instances do not make any decisions or




Figure 2: An illustration of MALG with $n = 4$ (see detailed explanation in Section 3.1)

updates, that is, they are paused (they might be resumed at some point though). We use \tilde{g}_t to denote the scalar \tilde{f}_t output by the active instance. See Algorithm 2 for the pseudocode.

For better illustration, we give an example with $n = 4$ in Figure 2. Suppose that the realization of the random scheduling by Procedure 1 is: one order-4 instance (red), zero order-3 instance, two order-2 instances (green), two order-1 instances (blue), and five order-0 instances (purple). The bolder part of the segment indicates the period of time when the instances are active, while the thinner part indicates the inactive period. For example, the red order-4 instance is active for the first round, then paused for the next 8 rounds, and then resumed (from the frozen internal states) for another 3 rounds before becoming inactive again. The dashed black arrow marked with ① indicates that ALG is executed as if the two sides of the arrow are concatenated. On the other hand, as another example, the two purple instances on the two sides of the dashed line marked with ② are two *different* order-0 instances, so the second one should start from scratch even though they are consecutive. One can see that at any point of time, the active instance is always the one with the shortest length.

Regret analysis of MALG The multi-scale nature of MALG allows the learner's regret to also enjoy a multi-scale structure, as shown in the next lemma (proof deferred to Appendix B).

Lemma 3 *Let $\hat{n} = \log_2 T + 1$ and $\hat{\rho}(t) = 6\hat{n} \log(T/\delta)\rho(t)$. MALG with input $n \leq \log_2 T$ guarantees the following: for any instance alg that MALG maintains and any $t \in [alg.s, alg.e]$, as long as $\Delta_{[alg.s, t]} \leq \rho(t')$ where $t' = t - alg.s + 1$, we have with probability at least $1 - \frac{\delta}{T}$:*

$$\tilde{g}_t \geq \min_{\tau \in [alg.s, t]} f_{\tau}^* - \Delta_{[alg.s, t]}, \quad \frac{1}{t'} \sum_{\tau=alg.s}^t (\tilde{g}_{\tau} - R_{\tau}) \leq \hat{\rho}(t') + \hat{n} \Delta_{[alg.s, t]}, \quad (3)$$

and the number of instances started within $[alg.s, t]$ is upper bounded by $6\hat{n} \log(T/\delta) \frac{C(t')}{C(1)}$.

Note that Eq. (3) is essentially the analogue of Eq. (1) (up to logarithmic terms) with the starting time changed from 1 to $alg.s$. It shows that even if we have multiple instances interleaving in a complicated way, the regret for a specific interval is still almost the same as running ALG alone on this interval, thanks to the carefully chosen probability in Procedure 1. Recall that there is always an order- n instance starting from the beginning of the block, so MALG is always providing

Algorithm 3: MALG with Stationarity TEsts and Restarts (MASTER)

input: $\widehat{\rho}(\cdot)$ (defined in Lemma 3)

Initialize: $t \leftarrow 1$

for $n = 0, 1, \dots$ **do**

Set $t_n \leftarrow t$ and initialize an MALG (Algorithm 2) for the block $[t_n, t_n + 2^n - 1]$.

while $t < t_n + 2^n$ **do**

Receive \tilde{g}_t and π_t from MALG, execute π_t , and receive reward R_t .

Update MALG with any feedback from the environment, and set $U_t = \min_{\tau \in [t_n, t]} \tilde{g}_\tau$.

Perform **Test 1** and **Test 2** (see below). Increment $t \leftarrow t + 1$.

if either test returns fail then restart from Line 3.

end

end

Test 1: If $t = \text{alg.e}$ for some order- m alg and $\frac{1}{2^m} \sum_{\tau=\text{alg.s}}^{\text{alg.e}} R_\tau \geq U_t + 9\widehat{\rho}(2^m)$, return *fail*.

Test 2: If $\frac{1}{t-t_n+1} \sum_{\tau=t_n}^t (\tilde{g}_\tau - R_\tau) \geq 3\widehat{\rho}(t - t_n + 1)$, return *fail*.

a stronger multi-scale guarantee compared to running ALG alone. This richer guarantee facilitates non-stationarity detection as we show next.

3.2. MASTER: Equipping MALG with Stationarity Tests

We are now ready to present our final algorithm MASTER, short for **MALG with Stationarity TEsts and Restarts** (see Algorithm 3). MASTER runs MALG in a sequence of blocks with doubling lengths ($2^0, 2^1, \dots$). Within each block of length 2^n (with t_n being the starting time), MASTER simply runs a new instance of MALG and records the minimum optimistic predictor thus far for this block $U_t = \min_{\tau \in [t_n, t]} \tilde{g}_\tau$. At the end of each time, MASTER performs two tests (**Test 1** and **Test 2**), and if either of them returns *fail*, MASTER restarts from scratch.

The two tests exactly follow the ideas described in Section 2.3 (recall Figure 1). Following Eq. (2), we decompose the regret on $[t_n, t]$ as **term**₁ + **term**₂ where **term**₁ = $\sum_{\tau=t_n}^t (f_\tau^* - \tilde{g}_\tau)$ and **term**₂ = $\sum_{\tau=t_n}^t (\tilde{g}_\tau - R_\tau)$. **Test 1** prevents **term**₁ $\leq \sum_{\tau=t_n}^t (f_\tau^* - U_\tau)$ from growing too large by testing if there is some order- m instance's interval during which the learner's average performance $\frac{1}{2^m} \sum_{\tau=\text{alg.s}}^{\text{alg.e}} R_\tau$ is larger than the promised performance upper bound U_t by an amount of $9\widehat{\rho}(2^m)$. On the other hand, **Test 2** prevents **term**₂ from growing too large by directly testing if its average is large than something close to the promised regret bound $3\widehat{\rho}(t - t_n + 1)$.

It is now clear that MASTER indeed does not require the knowledge of L or Δ at all. To analyze MASTER, we prove the following key lemma that bounds the regret on a single block $[t_n, E_n]$ where E_n is either $t_n + 2^n - 1$ or something smaller in the case where a restart is triggered.

Lemma 4 *With high probability, the dynamic regret of MASTER on any block $\mathcal{J} = [t_n, E_n]$ where $E_n \leq t_n + 2^n - 1$ is bounded as*

$$\sum_{\tau \in \mathcal{J}} (f_\tau^* - R_\tau) \leq \tilde{\mathcal{O}} \left(\sum_{i=1}^{\ell} C(|\mathcal{I}'_i|) + \sum_{m=0}^n \frac{\rho(2^m)}{\rho(2^n)} C(2^m) \right). \quad (4)$$

where $\{\mathcal{I}'_1, \dots, \mathcal{I}'_\ell\}$ is any partition of \mathcal{J} such that $\Delta_{\mathcal{I}'_i} \leq \rho(|\mathcal{I}'_i|)$ for all i .

See [Appendix C](#) for the proof. When $\rho(t) = \Theta(1/\sqrt{t})$ (as in all our examples), the first term is $\tilde{O}(\sum_{i=1}^{\ell} \sqrt{|\mathcal{I}'_i|}) = \tilde{O}(\sqrt{\ell|\mathcal{J}|})$ by Cauchy-Schwarz; the second term is of order $\tilde{O}(\sqrt{2^n})$. To derive a bound in terms of L , we can simply choose the partition $\{\mathcal{I}'_1, \dots, \mathcal{I}'_{\ell}\}$ in a way such that $\Delta_{\mathcal{I}'_i} = 0$ and $\ell = L_{\mathcal{J}}$, while to derive a bound in terms of Δ , the partition needs to be chosen more carefully depending on the value of $\Delta_{\mathcal{J}}$. Noting that the number of blocks between two restarts is always at most $\log_2 T$, to finally prove [Theorem 2](#), it remains to bound the number of restarts, which intuitively should scale with L or Δ because by design a restart will not be triggered when the environment is stationary. The complete proof is deferred to [Appendix D–Appendix F](#).

4. Extension to Reinforcement Learning in Infinite-horizon Communicating MDPs

As mentioned, applying our results to infinite-horizon RL ([Jaksch et al., 2010](#)) requires some extra care and extensions. We refer the reader to ([Cheung et al., 2020](#)) for a thorough introduction on the problem setup of infinite-horizon RL in time-varying communicating MDPs. Here, we only highlight its difference compared to episode RL and explain how to fit it into our framework. Specifically, in episodic RL, we have treated each episode (consisting of multiple steps in an MDP) as one round of our framework, each state-to-action mapping as a policy π , and the expected reward of executing π in the MDP for round t as $f_t(\pi)$. In infinite-horizon RL, while the meaning of π and f_t remains the same, there is no episode any more and the learner interacts with the changing MDP from the start to the end without any reset on her state. In this case, we treat each step (that is, each state transition) in the MDP as one round in our framework, and the meaning of the reward feedback R_t has now changed from a noisy observation of the policy’s reward $f_t(\pi_t)$ to just the reward of π_t for this single step. With this change, the dynamic regret definition remains the same.

Due to the black-box nature of our approach, if one has a base algorithm that satisfies something close to [Assumption 1](#) within this setup, then it is not hard to imagine that the same idea of MASTER can be applied. In [Section 4.1](#), we provide such a base algorithm, and in [Section 4.2](#), we combine it with appropriate multi-scale scheduling and detection to obtain our final results.

4.1. UCRL with Adaptive Confidence Widening

Our base algorithm, UCRL-ACW, is an improvement of the standard UCRL algorithm ([Jaksch et al., 2010](#)) and its variant UCRL-CW ([Cheung et al., 2020](#)). The pseudocode is shown in [Algorithm 4](#) ([Appendix A](#)), where we highlight the differences compared to UCRL and UCRL-CW in blue.

The first difference is the explicit mention that the next state of the learner might sometime be *arbitrarily* assigned instead of following the transition of the current MDP ([Line 4](#)). This is necessary because of the multi-scale scheduling of MALG. Indeed, recall that in MALG, an instance of the base algorithm can sometimes be paused and then resumed later. In the infinite-horizon RL setup, this means that the instance can be resumed from an arbitrary state. Other than making this detail explicit, however, nothing really needs to be changed in the algorithm, since this happens infrequently and only incurs small additional regret due to the communicating property of the MDPs.

The second key difference is an adaptive version of the *Confidence Widening* technique of ([Cheung et al., 2020](#)) (see [Line 4–Line 4](#)). As pointed out in ([Cheung et al., 2020](#)), in non-stationary environments, the Extended Value Iteration (EVI) subroutine of UCRL might return a bias vector (\tilde{h}_k) with span much larger than D_{\max} , the maximum diameters of all the MDPs. To address this issue, their confidence widening technique adds a constant η , tuned based on Δ , to the confidence level of the confidence set \mathcal{P}_k , which eventually leads to sub-optimal regret $\Delta^{1/4}T^{3/4}$. Our *adaptive*

confidence widening, on the other hand, adaptively selects the value of η in a doubling manner, so that in a relatively stationary environment we only widen the confidence set slightly, while in a more non-stationary environment the widening is more significant. To avoid incurring too much additional regret in the latter case, we also monitor the cumulative widening amount and terminate the algorithm if it exceeds a certain threshold (Line 4–Line 4), because this implies that the environment is highly non-stationary. (This termination will also be a restart signal for MASTER.)

Finally, notice that our black-box approach requires knowing the regret bound $\rho(\cdot)$ of the base algorithm, which in this case depends on D_{\max} , a potentially unknown quantity. To deal with this issue, UCRL-ACW takes a guess \bar{D} on the value of D_{\max} as an additional input. In the next subsection, we discuss how MASTER decides the value of \bar{D} when D_{\max} is unknown.

With all these modifications, our base algorithm UCRL-ACW indeed provides a guarantee similar to Eq. (1) of Assumption 1; see Lemma 11.

4.2. Multi-scale UCRL-ACW and Its Combination with MASTER

Now, we use the same idea as in Section 3.1 to create a multi-scale version of UCRL-ACW, under a fixed input \bar{D} . The resulted algorithm is called Multi-scale UCRL-ACW or MUCRL for short (see Algorithm 5). MUCRL is basically identical to MALG with UCRL-ACW as the base algorithm, except that we let MUCRL terminate whenever the currently active UCRL-ACW instance makes a restart signal (due to having an abnormally large cumulative widening amount). The guarantee for MUCRL is provided in Lemma 12, which parallels Lemma 3.

Next, as in Section 3.2, we further combine MUCRL with non-stationarity tests, leading to MASTER-UCRL (see Algorithm 6). The only difference compared to MASTER is an additional condition to restart (highlighted in blue) — when MUCRL terminates due to a restart signal from an UCRL-ACW instance. We provide a single-block regret bound guarantee for MASTER-UCRL under a fixed \bar{D} in Lemma 13, which parallels Lemma 4. Finally, we discuss three different cases with knowledge of different parameters (if any), leading to the three results listed in Table 1.

Known D_{\max} When D_{\max} is known, we simply set $\bar{D} = D_{\max}$. In this case, all restarts of MASTER-UCRL are due to non-stationarity, and we can bound their number in terms of L or Δ . Together with the single-block regret guarantee from Lemma 13, we prove that MASTER-UCRL’s dynamic regret is $\tilde{O}(\min\{\text{Reg}_L^*, \text{Reg}_\Delta^*\})$; see Theorem 28 for the dependence on other parameters.

Unknown D_{\max} and Known L or Δ When D_{\max} is unknown, we unfortunately require the knowledge of L to get Reg_L^* and the knowledge of Δ to get Reg_Δ^* . However, as shown in Table 1, this still significantly improves over the best existing bounds $\tilde{O}(L^{1/3}T^{2/3})$ and $\tilde{O}(\Delta^{1/4}T^{3/4})$ when L and Δ are known. Specifically, we apply a doubling trick to set the value of \bar{D} following the strategy below, where we call the interval between two restarts an *epoch*:

1. Initialize $\bar{D} \leftarrow 1$.
2. Run MASTER-UCRL with \bar{D} . If the number of epochs exceeds \bar{N} , then double \bar{D} and repeat this step. Here, \bar{N} is set to L if L is known or $1 + 3(S^{-2}A^{-1}\Delta^2T)^{1/3}$ if Δ is known.

The rationale behind monitoring the number of epochs is that, when \bar{D} is too small, UCRL-ACW might have an abnormally large cumulative widening amount and signal a restart even in a fairly stationary environment. In Lemma 26, we show that if $\bar{D} \geq D_{\max}$, the number of epochs produced by MASTER-UCRL is upper bounded by the value of \bar{N} set above. Therefore, if it exceeds this

number, we can infer $\bar{D} < D_{\max}$ and double its value. This allows us to prove the regret bound Reg_L^* or Reg_Δ^* again; see [Theorem 29](#) for the details.

No prior knowledge at all When nothing is known, we apply the Bandit-over-Reinforcement-Learning (BoRL) framework of ([Cheung et al., 2019, 2020](#)) to get a suboptimal bound of order $\tilde{\mathcal{O}}(\min\{\text{Reg}_L^*, \text{Reg}_\Delta^*\} + T^{3/4})$. BoRL also serves as a black-box reduction to obtain parameter-free algorithms (albeit suboptimal), so applying it to our algorithm is straightforward. We omit the details and only give the concrete bound in [Appendix H](#). We leave the question of whether the optimal bound is achievable when L , Δ , and D_{\max} are all unknown as a future direction.

5. Conclusion and Future Directions

In this work, we study reinforcement learning in non-stationary environments. We propose a general black-box approach that can convert an algorithm with near-optimal regret in a (near-)stationary environment to another algorithm with near-optimal dynamic regret in a non-stationary environment. Prior to our work, the bound of $\tilde{\mathcal{O}}(\Delta^{1/3}T^{2/3})$ is only achievable with the knowledge of Δ , and no algorithm achieves the bound of $\tilde{\mathcal{O}}(\sqrt{LT})$ even with the knowledge of L . Our algorithm achieves both bounds simultaneously without any prior knowledge.

It would be interesting to see whether algorithms with data-dependent bounds work with our black-box approach. Previous work in this direction ([Wei et al., 2016](#)) achieves an improved dynamic regret bound for multi-armed bandits when the cumulative variance of the loss is small; however, their approach crucially relies on the knowledge on the degree of non-stationarity as well as the cumulative variance. On the other hand, there are some immediate difficulties in applying our black-box approach to data-dependent algorithms. For example, the monotonicity of the the average regret $\rho(\cdot)$ may not hold anymore, and it is unclear how to set the probability of initiating a new base algorithm. Therefore, the task of achieving data-dependent dynamic bounds without prior knowledge seems to be challenging and requires other innovations.

Another future direction is to study a class of contextual bandit problems where the context is adversarially generated ([Abbasi-Yadkori et al., 2011](#); [Cheung et al., 2019](#); [Foster and Rakhlin, 2020](#)). In this case, the expected reward of the optimal policy changes over time even if the environment is stationary, so our current algorithm cannot be directly applied. For linear contextual bandits with adversarial contexts ([Abbasi-Yadkori et al., 2011](#); [Cheung et al., 2019](#)), the fix is straightforward though: instead of requiring the base algorithm to generate a scalar \tilde{f}_t in each round, we let it generate a *confidence set* for the hidden parameter, and check the inconsistency of the confidence set over time. However, for general contextual bandits with adversarial contexts, where algorithms do not necessarily maintain a confident set for the hidden parameter ([Foster and Rakhlin, 2020](#)), the extension is less clear and is left for future investigation.

Finally, we are not aware of any near-optimal *convex bandit* algorithm satisfying our [Assumption 1](#), so achieving near-optimal dynamic regret bound in general convex bandits is also left open.

Acknowledgments We thank Peng Zhao for pointing out the technical flaw made in previous works on non-stationary linear bandits as well as a fix in ([Zhao and Zhang, 2021](#)). We also thank anonymous reviewers for pointing out the relation between our algorithm and regret balancing ([Abbasi-Yadkori et al., 2020](#); [Pacchiano et al., 2020](#)). This work is supported by NSF Award IIS-1943607 and a Google Faculty Research Award.

References

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24:2312–2320, 2011.
- Yasin Abbasi-Yadkori, Aldo Pacchiano, and My Phan. Regret balancing for bandit and rl model selection. *arXiv preprint arXiv:2006.05491*, 2020.
- Alekh Agarwal, Daniel Hsu, Satyen Kale, John Langford, Lihong Li, and Robert Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In *International Conference on Machine Learning*, pages 1638–1646, 2014.
- Raman Arora, Ofer Dekel, and Ambuj Tewari. Deterministic mdps with adversarial rewards and bandit feedback. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, pages 93–101, 2012.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002a.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multi-armed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002b.
- Peter Auer, Pratik Gajane, and Ronald Ortner. Adaptively tracking the best bandit arm with an unknown number of distribution changes. In *Conference on Learning Theory*, pages 138–158, 2019.
- Qi Cai, Zhuoran Yang, Chi Jin, and Zhaoran Wang. Provably efficient exploration in policy optimization. In *International Conference on Machine Learning*, pages 1283–1294. PMLR, 2020.
- Liyu Chen, Haipeng Luo, and Chen-Yu Wei. Minimax regret for stochastic shortest path with adversarial costs and known transition. In *Conference on Learning Theory*, 2021.
- Wei Chen, Liwei Wang, Haoyu Zhao, and Kai Zheng. Combinatorial semi-bandit in the non-stationary environment. *arXiv preprint arXiv:2002.03580*, 2020.
- Yifang Chen, Chung-Wei Lee, Haipeng Luo, and Chen-Yu Wei. A new algorithm for non-stationary contextual bandits: Efficient, optimal and parameter-free. In *Conference on Learning Theory*, pages 696–726, 2019.
- Wang Chi Cheung, David Simchi-Levi, and Ruihao Zhu. Learning to optimize under non-stationarity. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1079–1087. PMLR, 2019.
- Wang Chi Cheung, David Simchi-Levi, and Ruihao Zhu. Reinforcement learning for non-stationary markov decision processes: The blessing of (more) optimism. In *International Conference on Machine Learning*, pages 1843–1854. PMLR, 2020.
- Amit Daniely, Alon Gonen, and Shai Shalev-Shwartz. Strongly adaptive online learning. In *International Conference on Machine Learning*, pages 1405–1411, 2015.

- Ofer Dekel and Elad Hazan. Better rates for any adversarial deterministic mdp. In *International Conference on Machine Learning*, pages 675–683, 2013.
- Travis Dick, Andras Gyorgy, and Csaba Szepesvari. Online learning in markov decision processes with changing cost sequences. In *International Conference on Machine Learning*, pages 512–520, 2014.
- Omar Darwiche Domingues, Pierre Ménard, Matteo Pirotta, Emilie Kaufmann, and Michal Valko. A kernel-based approach to non-stationary reinforcement learning in metric spaces. In *International Conference on Artificial Intelligence and Statistics*, pages 3538–3546. PMLR, 2021.
- Eyal Even-Dar, Sham M Kakade, and Yishay Mansour. Online markov decision processes. *Mathematics of Operations Research*, 34(3):726–736, 2009.
- Louis Faury, Yoan Russac, Marc Abeille, and Clément Calauzènes. Regret bounds for generalized linear bandits under parameter drift. *arXiv preprint arXiv:2103.05750*, 2021.
- Yingjie Fei, Zhuoran Yang, Zhaoran Wang, and Qiaomin Xie. Dynamic regret of policy optimization in non-stationary environments. *Advances in Neural Information Processing Systems*, 33, 2020.
- Sarah Filippi, Olivier Cappé, Aurélien Garivier, and Csaba Szepesvári. Parametric bandits: the generalized linear case. In *Proceedings of the 23rd International Conference on Neural Information Processing Systems-Volume 1*, pages 586–594, 2010.
- Dylan Foster and Alexander Rakhlin. Beyond ucb: Optimal and efficient contextual bandits with regression oracles. In *International Conference on Machine Learning*, pages 3199–3210. PMLR, 2020.
- Pratik Gajane, Ronald Ortner, and Peter Auer. A sliding-window algorithm for markov decision processes with arbitrarily changing rewards and transitions. *arXiv preprint arXiv:1805.10066*, 2018.
- Elad Hazan and Comandur Seshadhri. Adaptive algorithms for online decision problems. In *Electronic colloquium on computational complexity (ECCC)*, volume 14, 2007.
- Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(4), 2010.
- Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is q-learning provably efficient? In *Advances in neural information processing systems*, pages 4863–4873, 2018.
- Chi Jin, Tiancheng Jin, Haipeng Luo, Suvrit Sra, and Tiancheng Yu. Learning adversarial markov decision processes with bandit feedback and unknown transition. In *International Conference on Machine Learning*, pages 4860–4869. PMLR, 2020a.
- Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pages 2137–2143. PMLR, 2020b.

- Tiancheng Jin and Haipeng Luo. Simultaneously learning stochastic and adversarial episodic mdps with known transition. *Advances in Neural Information Processing Systems*, 33, 2020.
- Kwang-Sung Jun, Francesco Orabona, Stephen Wright, and Rebecca Willett. Improved strongly adaptive online learning using coin betting. In *Artificial Intelligence and Statistics*, pages 943–951, 2017.
- Baekjin Kim and Ambuj Tewari. Randomized exploration for non-stationary stochastic linear bandits. In *Uncertainty in Artificial Intelligence*, 2020.
- Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- Tal Lencewicki, Aviv Rosenberg, and Yishay Mansour. Learning adversarial markov decision processes with delayed feedback. *arXiv preprint arXiv:2012.14843*, 2020.
- Chung-Wei Lee, Haipeng Luo, Chen-Yu Wei, and Mengxiao Zhang. Bias no more: high-probability data-dependent regret bounds for adversarial bandits and mdps. *Advances in Neural Information Processing Systems*, 33, 2020.
- Yingying Li and Na Li. Online learning for markov decision processes in nonstationary environments: A dynamic regret analysis. In *2019 American Control Conference (ACC)*, pages 1232–1237. IEEE, 2019.
- Haipeng Luo and Robert E Schapire. Achieving all with no parameters: Adanormalhedge. In *Conference on Learning Theory*, pages 1286–1304, 2015.
- Haipeng Luo, Chen-Yu Wei, Alekh Agarwal, and John Langford. Efficient contextual bandits in non-stationary worlds. In *Conference On Learning Theory*, pages 1739–1776. PMLR, 2018.
- Thodoris Lykouris, Max Simchowitz, Aleksandrs Slivkins, and Wen Sun. Corruption robust exploration in episodic reinforcement learning. In *Conference on Learning Theory*, 2021.
- Weichao Mao, Kaiqing Zhang, Ruihao Zhu, David Simchi-Levi, and Tamer Başar. Is model-free learning nearly optimal for non-stationary rl? In *International Conference on Machine Learning*, 2021.
- Gergely Neu, András György, and Csaba Szepesvári. The online loop-free stochastic shortest-path problem. In *COLT*, volume 2010, pages 231–243. Citeseer, 2010.
- Gergely Neu, Andras Gyorgy, and Csaba Szepesvári. The adversarial stochastic shortest path problem with unknown transition probabilities. In *Artificial Intelligence and Statistics*, pages 805–813, 2012.
- Gergely Neu, András György, Csaba Szepesvári, and András Antos. Online markov decision processes under bandit feedback. *IEEE Transactions on Automatic Control*, 59(3):676–691, 2013.
- Ronald Ortner, Pratik Gajane, and Peter Auer. Variational regret bounds for reinforcement learning. In *Uncertainty in Artificial Intelligence*, pages 81–90. PMLR, 2020.

- Aldo Pacchiano, Christoph Dann, Claudio Gentile, and Peter Bartlett. Regret bound balancing and elimination for model selection in bandits and rl. *arXiv preprint arXiv:2012.13045*, 2020.
- Aviv Rosenberg and Yishay Mansour. Online convex optimization in adversarial markov decision processes. In *International Conference on Machine Learning*, pages 5478–5486, 2019.
- Aviv Rosenberg and Yishay Mansour. Stochastic shortest path with adversarially changing costs. *arXiv preprint arXiv:2006.11561*, 2020.
- Yoan Russac, Claire Vernade, and Olivier Cappé. Weighted linear bandits for non-stationary environments. *Advances in Neural Information Processing Systems*, 2019.
- Yoan Russac, Olivier Cappé, and Aurélien Garivier. Algorithms for non-stationary generalized linear bandits. *arXiv preprint arXiv:2003.10113*, 2020.
- Lior Shani, Yonathan Efroni, Aviv Rosenberg, and Shie Mannor. Optimistic policy optimization with bandit feedback. In *International Conference on Machine Learning*, pages 8604–8613. PMLR, 2020.
- David Simchi-Levi and Yunzong Xu. Bypassing the monster: A faster and simpler optimal algorithm for contextual bandits under realizability. *Available at SSRN*, 2020.
- Ahmed Touati and Pascal Vincent. Efficient learning in non-stationary linear markov decision processes. *arXiv preprint arXiv:2010.12870*, 2020.
- Ruosong Wang, Simon S Du, Lin F Yang, and Sham M Kakade. Is long horizon reinforcement learning more difficult than short horizon reinforcement learning? *Advances in Neural Information Processing Systems*, 2020.
- Chen-Yu Wei, Yi-Te Hong, and Chi-Jen Lu. Tracking the best expert in non-stationary stochastic environments. *Advances in neural information processing systems*, 29:3972–3980, 2016.
- Peng Zhao and Lijun Zhang. Non-stationary linear bandits revisited. *arXiv preprint arXiv:2103.05324*, 2021.
- Peng Zhao, Lijun Zhang, Yuan Jiang, and Zhi-Hua Zhou. A simple approach for non-stationary linear bandits. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 746–755. PMLR, 26–28 Aug 2020.
- Huozhi Zhou, Jinglin Chen, Lav R Varshney, and Ashish Jagmohan. Nonstationary reinforcement learning with linear function approximation. *arXiv preprint arXiv:2010.04244*, 2020.
- Alexander Zimin and Gergely Neu. Online learning in episodic markovian decision processes by relative entropy policy search. *Advances in neural information processing systems*, 26:1583–1591, 2013.

Appendix A. Omitted Algorithms and Main Results in Section 4

Algorithm 4: UCRL with Adaptive Confidence Widening (UCRL-ACW)

input: $\bar{D} \geq 1$
 $t \leftarrow 1$. $N_1(s, a) \leftarrow 0$ for all s, a . $\Gamma \leftarrow 0$
for episode $k = 1, \dots$, **do**

 Set $t_k = t$, $\nu_k(s, a) = 0$ for all s, a .

 Define for all s, a :

$$\hat{p}_k(s'|s, a) = \frac{\sum_{\tau=1}^{t-1} \mathbb{1}[(s_\tau, a_\tau, s'_{\tau+1}) = (s, a, s')]}{N_k^+(s, a)},$$

$$\hat{r}_k(s, a) = \frac{\sum_{\tau=1}^{t-1} R_\tau \mathbb{1}[(s_\tau, a_\tau) = (s, a)]}{N_k^+(s, a)}.$$

 and for any η :

$$\mathcal{P}_k^\eta(s, a) = \left\{ \tilde{p}(\cdot|s, a) \in \Delta_S : \|\tilde{p}(\cdot|s, a) - \hat{p}_k(\cdot|s, a)\|_1 \leq \sqrt{S} \cdot \text{conf}_k(s, a) + \eta \right\}$$

$$\mathcal{R}_k(s, a) = \left\{ \tilde{r}(s, a) \in [0, 1] : |\tilde{r}(s, a) - \hat{r}_k(s, a)| \leq \text{conf}_k(s, a) \right\}$$

 where $\text{conf}_k(s, a) \triangleq 8\sqrt{\frac{\log(SAT/\delta)}{N_k^+(s, a)}}$ and $N_k^+(s, a) = \max\{1, N_k(s, a)\}$.

 $\eta \leftarrow \frac{1}{T}$
while true **do**

 Perform EVI on $(\mathcal{P}_k^\eta, \mathcal{R}_k)$ with error parameter $\epsilon_k = \sqrt{\frac{1}{t}}$, and obtain $\tilde{\pi}, \tilde{h}, \tilde{J}$.

if $sp(\tilde{h}) \leq 2\bar{D}$ **then break**
 $\eta \leftarrow 2\eta$
end
 $\pi_k \leftarrow \tilde{\pi}, \tilde{h}_k \leftarrow \tilde{h}, \tilde{J}_k \leftarrow \tilde{J}, \eta_k \leftarrow \eta$
 \triangleleft Adaptive confidence widening

while $\nu_k(s, a) < N_k^+(s, a)$ for all s, a **do**

 Choose action $a_t \sim \pi_k(s_t)$.

 $\nu_k(s_t, a_t) \leftarrow \nu_k(s_t, a_t) + 1$
 $\Gamma \leftarrow \Gamma + \eta_k$
if $\Gamma > 4S\sqrt{At \log(SAT/\delta)}$ **then terminate and signal restart** \triangleleft Early termination

 Observe the reward R_t with $\mathbb{E}[R_t] = r_t(s_t, a_t)$

 Observe $s'_{t+1} \sim p_t(\cdot|s_t, a_t)$.

 The next state s_{t+1} is either equal to s'_{t+1} , or re-assigned as an arbitrary state

 \triangleleft The next state might be re-assigned

 $t \leftarrow t + 1$
end
 $N_{k+1}(s, a) \leftarrow N_k(s, a) + \nu_k(s, a)$ for all s, a .

end

The following is the main result for the infinite-horizon MDP case. Its proof requires several lemmas

Algorithm 5: Multi-scale UCRL-ACW (MUCRL)

input: $n, \rho_{\text{UCRL}}(\cdot; \bar{D}), \bar{D}$ **Initialization:** run [Procedure 1](#) with base algorithm UCRL-ACW and inputs n and ρ_{UCRL} . At each time t , let the unique active instance be alg , output \tilde{g}_t (which is the quantity $\tilde{J}_{k(t)}$ of alg), follow alg 's decision, and update alg after receiving feedback from the environment. Additionally, [terminate if the \$alg\$ signals restart](#).

Algorithm 6: MASTER-UCRL

input: $\rho_{\text{UCRL}}(\cdot; \bar{D}), \bar{D}$ **Initialize:** $t \leftarrow 1$ **for** $n = 0, 1, \dots$ **do** Set $t_n \leftarrow t$ and initialize an MUCRL ([Algorithm 5](#)) for the block $[t_n, t_n + 2^n - 1]$. **while** $t < t_n + 2^n$ **do** Receive \tilde{g}_t and π_t from MUCRL, execute π_t , and receive reward R_t . Update MUCRL with any feedback from the environment, and set $U_t = \min_{\tau \in [t_n, t]} \tilde{g}_\tau$. Perform **Test 1** and **Test 2** (see below). Increment $t \leftarrow t + 1$. **if either test returns fail or MUCRL terminates then** restart from [Line 6](#). **end****end****Test 1:** If $t = alg.e$ for some order- m alg and $\frac{1}{2^m} \sum_{\tau=alg.s}^{alg.e} R_\tau \geq U_t + 9\hat{\rho}_{\text{UCRL}}(2^m; \bar{D})$, return *fail*.**Test 2:** If $\frac{1}{t-t_n+1} \sum_{\tau=t_n}^t (\tilde{g}_\tau - R_\tau) \geq 3\hat{\rho}_{\text{UCRL}}(t-t_n+1; \bar{D})$, return *fail*.

in the rest of this section, in addition to those from [Appendix B–Appendix E](#) whose ideas are mostly aligned with the standard setting. The final analysis is done in [Appendix G](#) and [Appendix H](#) (see [Theorem 28](#), [Theorem 29](#), and the discussions in [Appendix H](#)). Note that to be consistent with prior works in this setting, we adopt the notation $J_t(\pi)$, which is the expected average reward of executing π under the MDP for time t , and corresponds to the notation $f_t(\pi)$ we use in our general framework. Similarly, define $J_t^* = \max_{\pi} J_t(\pi)$.

Theorem 5 *Define non-stationarity measures*

$$\Delta = \sum_{t=1}^{T-1} \left(\max_{s,a} |r_t(s,a) - r_{t+1}(s,a)| + \max_{s,a} \|p_t(\cdot|s,a) - p_{t+1}(\cdot|s,a)\|_1 \right),$$

$$L = 1 + \sum_{t=1}^{T-1} \mathbb{1} \left\{ \max_{s,a} |r_t(s,a) - r_{t+1}(s,a)| + \max_{s,a} \|p_t(\cdot|s,a) - p_{t+1}(\cdot|s,a)\|_1 \neq 0 \right\}.$$

There exists an algorithm that takes D_{\max} as input and achieves

$$\sum_{t=1}^T (J_t^* - R_t) = \tilde{O} \left(\min \left\{ D_{\max} S \sqrt{ALT}, D_{\max} S^{\frac{2}{3}} A^{\frac{1}{3}} \Delta^{\frac{1}{3}} T^{\frac{2}{3}} + D_{\max} S \sqrt{AT} \right\} \right)$$

without knowing L or Δ . There is also an algorithm that takes L or Δ as input and achieves

$$\sum_{t=1}^T (J_t^* - R_t) = \tilde{\mathcal{O}}\left(D_{\max} S \sqrt{ALT}\right) \quad \text{or} \quad \sum_{t=1}^T (J_t^* - R_t) = \tilde{\mathcal{O}}\left(D_{\max} S^{\frac{2}{3}} A^{\frac{1}{3}} \Delta^{\frac{1}{3}} T^{\frac{2}{3}} + D_{\max} S \sqrt{AT}\right)$$

respectively, without knowing D_{\max} . Finally, there is an algorithm that achieves

$$\sum_{t=1}^T (J_t^* - R_t) = \tilde{\mathcal{O}}\left(D_{\max} (S^2 A)^{1/4} T^{3/4} + \min\left\{D_{\max} S \sqrt{ALT}, D_{\max} (S^2 A)^{\frac{1}{3}} \Delta^{\frac{1}{3}} T^{\frac{2}{3}}\right\}\right)$$

without knowing L , Δ , or D_{\max} .

A.1. Auxiliary Lemmas related to Extended Value Iteration and Bellman Equation

In this subsection, we provide auxiliary lemmas related to EVI and Bellman Equation. The results are extracted from (Jaksch et al., 2010; Cheung et al., 2020; Ortner et al., 2020). We restate them here for completeness.

Lemma 6 (Properties 1 and 2 in (Cheung et al., 2020)) *Let \tilde{J} , \tilde{h} , and $\tilde{\pi}$ be the set of solution obtained from EVI with confidence set \mathcal{R} and \mathcal{P} for reward and transition respectively, and error parameter ϵ . Then*

$$\tilde{J} + \tilde{h}(s) \geq \max_a \left(\max_{\tilde{r} \in \mathcal{R}(s,a)} \tilde{r}(s,a) + \max_{\tilde{p} \in \mathcal{P}(s,a)} \sum_{s'} \tilde{p}(s'|s,a) \tilde{h}(s') \right), \quad (5)$$

$$\tilde{J} + \tilde{h}(s) \leq \max_{\tilde{r} \in \mathcal{R}(s,\tilde{\pi}(s))} \tilde{r}(s,\tilde{\pi}(s)) + \max_{\tilde{p} \in \mathcal{P}(s,\tilde{\pi}(s))} \sum_{s'} \tilde{p}(s'|s,\tilde{\pi}(s)) \tilde{h}(s') + \epsilon. \quad (6)$$

Lemma 7 (Lemma 2 of (Cheung et al., 2020)) *Let \tilde{J} , \tilde{h} , and $\tilde{\pi}$ be the set of solution obtained from EVI with confidence set \mathcal{R} and \mathcal{P} for reward and transition respectively. If \mathcal{P} and \mathcal{R} contain an MDP with diameter upper bounded by D , then $\text{sp}(\tilde{h}) \leq 2D$.*

Lemma 8 (Eq. (16) of (Cheung et al., 2020)) *Let r, p define the reward function and the transition kernel for a communicating MDP, respectively. Let $\tilde{J} \in \mathbb{R}$, $\tilde{h} \in \mathbb{R}^S$ be bounded and satisfy*

$$\tilde{J} + \tilde{h}(s) \geq r(s,a) + \sum_{s'} p(s'|s,a) \tilde{h}(s')$$

for all s and a . Then $\tilde{J} \geq J^*$, where J^* is the average reward of the optimal policy under the MDP.

A.2. Guarantees of the UCRL-ACW Algorithm (when running alone with an input \bar{D})

Definition 9 *Define $\Delta^r(t) \triangleq \max_{s,a} |r_t(s,a) - r_{t+1}(s,a)|$, $\Delta^p(t) \triangleq \max_{s,a} \|p_t(\cdot|s,a) - p_{t+1}(\cdot|s,a)\|_1$, $\Delta^J(t) \triangleq \max_{\pi} |J_t(\pi) - J_{t+1}(\pi)|$. Similar to Definition 1, define $\Delta_{\mathcal{I}}^{\square} = \sum_{\tau=s}^{e-1} \Delta^{\square}(\tau)$ for interval $\mathcal{I} = [s, e]$, where $\square = r, p$, or J . Finally, we define $\Delta_{\mathcal{I}, \bar{D}} \triangleq \Delta_{\mathcal{I}}^r + 2\bar{D}\Delta_{\mathcal{I}}^p + \Delta_{\mathcal{I}}^J$.*

Lemma 10 (Theorem 1 of (Ortner et al., 2020)) $\Delta^J(t) \leq \Delta^r(t) + D_{\max} \Delta^p(t)$.

Lemma 11 (c.f. Assumption 1) *When run alone, Algorithm 4 with input \bar{D} guarantees for all t before it terminates:*

$$\begin{aligned} \tilde{J}_{k(t)} &\geq \min_{\tau \in [1, t]} J_\tau^* - \Delta_{[1, t]; \bar{D}} \\ \frac{1}{t} \sum_{\tau=1}^t \left(\tilde{J}_{k(\tau)} - R_\tau \right) &\leq \rho_{\text{UCRL}}(t; \bar{D}) + \frac{2}{t} \bar{D} \text{disc}_{[1, t]} + \Delta_{[1, t]; \bar{D}}, \end{aligned}$$

where $\rho_{\text{UCRL}}(t; \bar{D}) = \tilde{\Theta} \left(\min \left\{ \bar{D} S \sqrt{\frac{A}{t}} + \frac{\bar{D} S A}{t}, \bar{D} \right\} \right)$, $\text{disc}_{\mathcal{I}} \triangleq \sum_{t \in \mathcal{I}} \mathbb{1}[s_t \approx p_{t-1}(\cdot | s_{t-1}, a_{t-1})]$ is the number of state re-assignments within \mathcal{I} (Line 4 of Algorithm 4), and $k(t)$ is the index of the episode time t belongs to.

Proof Suppose that at time t the algorithm has not terminated. For any episode k that starts before t , we have

$$\bar{r}_k(s, a) = \frac{\sum_{\tau=1}^{t_k-1} r_\tau(s, a) \mathbb{1}[(s_\tau, a_\tau) = (s, a)]}{N_k^+(s, a)}, \quad \bar{p}_k(s' | s, a) = \frac{\sum_{\tau=1}^{t_k-1} p_\tau(s' | s, a) \mathbb{1}[(s_\tau, a_\tau) = (s, a)]}{N_k^+(s, a)}.$$

By Azuma's inequality, $\bar{r}_k(s, a) \in \mathcal{R}_k(s, a)$ and $\bar{p}_k(\cdot | s, a) \in \mathcal{P}_k^{\eta_k}(s, a)$ with high probability for all k, s, a .

To show the first part of the lemma, we lower bound the right-hand side of Eq. (5):

$$\begin{aligned} \tilde{J}_k + \tilde{h}_k(s) &\geq \bar{r}(s, a) + \sum_{s'} \bar{p}(s' | s, a) \tilde{h}_k(s') \\ &\geq r_\tau(s, a) + \sum_{s'} p_\tau(s' | s, a) \tilde{h}_k(s') - \left(\Delta_{[1, t]}^r + 2\bar{D} \Delta_{[1, t]}^p \right), \quad (\text{for any } \tau \in [1, t]) \end{aligned}$$

where in the last inequality we use $|\bar{r}(s, a) - r_\tau(s, a)| \leq \Delta_{[1, t]}^r$ and $\sum_{s'} |\bar{p}(s' | s, a) - p_\tau(s' | s, a)| \tilde{h}_k(s') \leq \|\bar{p}(\cdot | s, a) - p_\tau(\cdot | s, a)\|_1 \text{sp}(\tilde{h}_k) \leq 2\Delta_{[1, t]}^p \bar{D}$. Using Lemma 8, we get

$$\tilde{J}_k + \Delta_{[1, t]; \bar{D}} \geq \tilde{J}_k + \left(\Delta_{[1, t]}^r + 2\bar{D} \Delta_{[1, t]}^p \right) \geq J_\tau^*,$$

implying the first part of the lemma.

To show the second part of the lemma, starting from Eq. (6), we have with high probability

$$\begin{aligned} &\tilde{J}_k + \tilde{h}_k(s) \tag{7} \\ &\leq \bar{r}_k(s, \pi_k(s)) + \sum_{s'} \bar{p}_k(s' | s, \pi_k(s)) \tilde{h}_k(s') + 2\bar{D} \sqrt{S} \cdot \text{conf}_k(s, \pi_k(s)) + 2\bar{D} \eta_k + \epsilon_k \\ &\quad (\bar{r}_k(s, a) \in \mathcal{R}_k(s, a) \text{ and } \bar{p}_k(s, a) \in \mathcal{P}_k(s, a)) \\ &\leq r_\tau(s, \pi_k(s)) + \sum_{s'} p_\tau(s' | s, \pi_k(s)) \tilde{h}_k(s') + 2\bar{D} \sqrt{S} \cdot \text{conf}_k(s, \pi_k(s)) + \left(\Delta_{[1, t]}^r + 2\bar{D} \Delta_{[1, t]}^p \right) + 2\bar{D} \eta_k + \epsilon_k. \tag{8} \end{aligned}$$

Now, we apply Eq. (8) with $(k, \tau, s) = \{k(\tau), \tau, s_\tau\}_{\tau=1}^t$ respectively, and sum them up. Notice that $a_\tau = \pi_{k(\tau)}(s_\tau)$. Then we get

$$\sum_{\tau=1}^t \left(\tilde{J}_{k(\tau)} - R_\tau \right) \leq \sum_{\tau=1}^t \left(\sum_{s'} p_\tau(s' | s_\tau, a_\tau) \tilde{h}_{k(\tau)}(s') - \tilde{h}_{k(\tau)}(s_\tau) \right) + \sum_{\tau=1}^t (r_\tau(s_\tau, a_\tau) - R_\tau)$$

$$+ \sum_{\tau=1}^t 2\bar{D}\sqrt{S} \cdot \text{conf}_{k(\tau)}(s_\tau, a_\tau) + t\Delta_{[1,t];\bar{D}} + 2\bar{D} \sum_{\tau=1}^t \eta_{k(\tau)} + \sum_{\tau=1}^t \epsilon_{k(\tau)}.$$

We bound the terms on the right-hand side individually: for the first term, notice that when there is no state-reassignment at time $\tau + 1$, $\mathbb{E}_\tau[\tilde{h}_{k(\tau)}(s_{\tau+1})] = \sum_{s'} p_\tau(s'|s_\tau, a_\tau)\tilde{h}_{k(\tau)}(s')$. Therefore,

$$\begin{aligned} & \sum_{\tau=1}^t \left(\sum_{s'} p_\tau(s'|s_\tau, a_\tau)\tilde{h}_{k(\tau)}(s') - \tilde{h}_{k(\tau)}(s_\tau) \right) \\ & \leq \sum_{\tau=1}^t \left(\mathbb{E}_\tau \left[\tilde{h}_{k(\tau)}(s_{\tau+1}) \right] - \tilde{h}_{k(\tau)}(s_\tau) \right) + 2\bar{D}\text{disc}_{[1,t]} \\ & \leq 2\bar{D}\sqrt{t \log(SAT)} + 2\bar{D} \sum_{\tau=1}^t \mathbb{1} \left[\tilde{h}_{k(\tau)} \neq \tilde{h}_{k(\tau+1)} \right] + 2\bar{D}\text{disc}_{[1,t]} \quad (\text{by Azuma's inequality}) \\ & \leq 2\bar{D}\sqrt{t \log(SAT)} + 2\bar{D}SA \log_2 T + 2\bar{D}\text{disc}_{[1,t]}, \end{aligned}$$

where in the last inequality we use the fact that the number of episodes cannot exceed $SA \log_2 T$. For the other terms: $\sum_{\tau=1}^t (r_\tau(s_\tau, a_\tau) - R_\tau) \leq \tilde{\mathcal{O}}(\sqrt{t})$ by Azuma's inequality; $\sum_{\tau=1}^t 2\bar{D}\sqrt{S} \cdot \text{conf}_{k(\tau)}(s_\tau, a_\tau) = \tilde{\mathcal{O}}(\bar{D}S\sqrt{At})$ by the standard pigeonhole argument; $2\bar{D} \sum_{\tau=1}^t \eta_{k(\tau)} = \tilde{\mathcal{O}}(\bar{D}S\sqrt{At})$ by the termination condition specified in [Line 4](#); $\sum_{\tau=1}^t \epsilon_{k(\tau)}$ is also upper bounded by $\tilde{\mathcal{O}}(\bar{D}S\sqrt{At})$ by the way we choose the error parameter. Combining all the above arguments, we get

$$\sum_{\tau=1}^t \left(\tilde{J}_{k(\tau)} - R_\tau \right) \leq \tilde{\mathcal{O}}(\bar{D}S\sqrt{At} + \bar{D}SA) + 2\bar{D}\text{disc}_{[1,t]} + t\Delta_{[1,t];\bar{D}}$$

with high probability. On the other hand, $\sum_{\tau=1}^t (\tilde{J}_{k(\tau)} - R_\tau) \leq \bar{D}t$ is trivially true. Combining them we get the second claim of the lemma. \blacksquare

Appendix B. Analysis for the Multi-scale Algorithms

Proof of [Lemma 3](#) Below, we fix an *alg* and fix a $t \in [\text{alg}.s, \text{alg}.e]$, and consider the case $\Delta_{[\text{alg}.s,t]} \leq \rho(t')$ as specified in the lemma statement. For the first part of the lemma, note that \tilde{g}_t of MALG is defined as $\tilde{f}_t^{\text{alg}'}$ where *alg'* is the active instance of ALG at round t . By [Procedure 1](#), *alg'* can only be an instance that starts within $[\text{alg}.s, t]$ (i.e., $\text{alg}'.s \geq \text{alg}.s$). Therefore, the distribution drift undergone by *alg'* up to t is upper bounded by $\Delta_{[\text{alg}.s,t]} \leq \rho(t')$, which is further upper bounded by $\rho(t'')$ where t'' is the number of active rounds *alg'* runs within $[\text{alg}.s, t]$, because $\rho(\cdot)$ is a decreasing function. Therefore, the conditions in [Assumption 1](#) is satisfied for this *alg'*, and thus we have

$$\tilde{g}_t = \tilde{f}_t^{\text{alg}'} \geq \min_{\tau \leq t: \text{alg}' \text{ is active at } \tau} f_\tau^* - \Delta_{[\text{alg}'.s,t]} \geq \min_{\tau \in [\text{alg}.s,t]} f_\tau^* - \Delta_{[\text{alg}.s,t]},$$

proving the first part.

Next, we prove the second part of the lemma. We use S_m to denote the set of order- m instances which start within $[alg.s, t]$. Note that

$$\begin{aligned} \sum_{\tau=alg.s}^t (\tilde{g}_\tau - R_\tau) &= \sum_{\tau=alg.s}^t \sum_{m=0}^n \sum_{alg' \in S_m} \mathbb{1}[alg' \text{ is active at } \tau] \left(\tilde{f}_\tau^{alg'} - R_\tau \right) \\ &= \underbrace{\sum_{m=0}^n \sum_{alg' \in S_m} \sum_{\tau=alg.s}^t \mathbb{1}[alg' \text{ is active at } \tau] \left(\tilde{f}_\tau^{alg'} - R_\tau \right)}_{(*)}. \end{aligned} \quad (9)$$

The first equality holds because \tilde{g}_τ of MALG is defined as the \tilde{f}_τ of the active instance at round t .

Next, we focus on a specific m , and bound the $(*)$ term in Eq. (9). Let $|S_m| = \ell$ and $S_m = \{alg'_1, \dots, alg'_\ell\}$, and let $\mathcal{I}_i \triangleq [alg'_i.s, alg'_i.e] \cap [alg.s, t]$ for $i = 1, \dots, \ell$ (i.e., \mathcal{I}_i are the rounds within $[alg.s, t]$ where alg'_i is scheduled). Clearly, $|\mathcal{I}_i| \leq \min\{alg'_i.e - alg'_i.s + 1, t - alg.s + 1\} = \min\{2^m, t'\}$. By Assumption 1, we have

$$\begin{aligned} (*) &= \sum_{i=1}^{\ell} \sum_{\tau=alg.s}^t \mathbb{1}[alg'_i \text{ is active at } \tau] \left(\tilde{f}_\tau^{alg'_i} - R_\tau \right) \\ &\leq \sum_{i=1}^{\ell} (C(|\mathcal{I}_i|) + |\mathcal{I}_i| \Delta_{\mathcal{I}_i}) \\ &\leq \ell C(\min\{2^m, t'\}) + t' \Delta_{[alg.s, t]}, \end{aligned} \quad (10)$$

where in the first inequality we use Assumption 1, and that alg'_i updates for no more than $|\mathcal{I}_i|$ rounds in the interval $[alg.s, t]$ (also, the condition in Assumption 1 is satisfied because $\Delta_{\mathcal{I}_i} \leq \Delta_{[alg.s, t]} \leq \rho(t') \leq \rho(|\mathcal{I}_i|)$). In the last inequality, for the first term, we use that $C(\cdot)$ is increasing; for the second term, we use $|\mathcal{I}_i| \leq t'$, and that $\Delta_{\mathcal{I}_1} + \dots + \Delta_{\mathcal{I}_\ell} \leq \Delta_{[alg.s, t]}$ since $\mathcal{I}_1, \dots, \mathcal{I}_\ell$ are non-overlapping intervals lying within $[alg.s, t]$.

By Procedure 1, for every m , the expected number of order- m ALG's that starts within the interval $[alg.s, t]$ can be upper bounded as

$$\mathbb{E}[|S_m|] \leq \frac{\rho(2^n)}{\rho(2^m)} \left\lceil \frac{t'}{2^m} \right\rceil \leq \frac{\rho(2^n)}{\rho(2^m)} \left(\frac{t'}{2^m} + 1 \right) \leq \frac{\rho(2^n)}{\rho(2^m)} \frac{t'}{2^m} + 1 \quad (11)$$

By Bernstein's inequality, with probability $1 - \frac{\delta}{T}$, $|S_m| \leq \mathbb{E}[|S_m|] + \sqrt{2\mathbb{E}[|S_m|] \log(T/\delta)} + \log(T/\delta) \leq 2\mathbb{E}[|S_m|] + 2\log(T/\delta)$. Thus, continuing from Eq. (10), we have with probability at least $1 - \frac{\delta}{T}$,

$$\begin{aligned} (*) &\leq 2 \cdot \left(\frac{\rho(2^n)}{\rho(2^m)} \frac{t'}{2^m} + 1 \right) C(\min\{2^m, t'\}) + 2\log(T/\delta)C(\min\{2^m, t'\}) + t' \Delta_{[alg.s, t]} \\ &\leq 2 \left(\frac{C(t')}{C(2^m)} + 2 \right) \log(T/\delta)C(\min\{2^m, t'\}) + t' \Delta_{[alg.s, t]} \quad (\rho(2^n)t' \leq \rho(t')t' = C(t')) \\ &\leq 6C(t') \log(T/\delta) + t' \Delta_{[alg.s, t]} \quad (C(\cdot) \text{ is an increasing function}) \end{aligned} \quad (12)$$

Finally, using this in Eq. (9), we get the second claim of the lemma: with probability at least $1 - \frac{\delta}{T}$,

$$\sum_{\tau=alg.s}^t (\tilde{g}_\tau - R_\tau) \leq 6(n+1)C(t') \log(1/\delta) + t'(n+1)\Delta_{[alg.s,t]}. \quad (13)$$

For the third part of the lemma, as we calculated above, with probability at least $1 - \frac{\delta}{T}$, the number of instances started within $[alg.s, t]$ is upper bounded by

$$\sum_{m=0}^n 2 \cdot \left(\frac{\rho(2^n)}{\rho(2^m)} \frac{t'}{2^m} + 2 \right) \log(T/\delta) \leq 2\hat{n} \left(\frac{C(t')}{C(1)} + 2 \right) \log(T/\delta) \leq 6\hat{n} \frac{C(t')}{C(1)} \log(T/\delta)$$

where we use $\rho(2^m)2^m = C(2^m) \geq C(1)$ and $\rho(2^n)t' \leq \rho(t')t' = C(t')$. \blacksquare

Lemma 12 (c.f. Lemma 3) *Before MUCRL terminates, for every alg and $t \in [alg.s, alg.e]$, MUCRL guarantees with high probability*

$$\begin{aligned} \tilde{g}_t &\geq \min_{\tau \in [alg.s,t]} J_\tau^* - \Delta_{[alg.s,t];\bar{D}} \\ \frac{1}{t'} \sum_{\tau=alg.s}^t (\tilde{g}_\tau - R_\tau) &\leq \hat{\rho}_{\text{UCRL}}(t'; \bar{D}) + \hat{n} \Delta_{[alg.s,t];\bar{D}} \end{aligned}$$

where $t' = t - alg.s + 1$, $\hat{n} = \log_2 T + 1$, and $\hat{\rho}_{\text{UCRL}}(t; \bar{D}) = 18\hat{n} \log(T/\delta) \rho_{\text{UCRL}}(t; \bar{D})$.

Proof This proof is similar to that of Lemma 3. For the first part of the lemma, we can simply follow the proof of the first part of Lemma 3, with \tilde{f}_t replaced by $\tilde{J}_{k(t)}$, and $\Delta_{[alg.s,t]}$ by $\Delta_{[alg.s,t];\bar{D}}$.

For the second part, the analysis still tightly follows that of Lemma 3, but we need to add the additional cost caused by state re-assignment (i.e., the $\bar{D}disc_{[1,t]}$ term in Lemma 11). Following the same arguments as in proof as in Eq. (9) and Eq. (10), we get

$$\begin{aligned} \sum_{\tau=alg.s}^t (\tilde{g}_\tau - R_\tau) &= \sum_{m=0}^n \sum_{alg' \in S_m} \sum_{\tau=alg.s}^t \mathbb{1}[alg' \text{ is active at } \tau] \left(\tilde{f}_\tau^{alg'} - R_\tau \right) \\ &\quad (S_m \triangleq \text{the set of order-}m \text{ ALG initiated within } [alg.s, t]) \\ &\leq \sum_{m=0}^n \sum_{i=1}^{|S_m|} \sum_{\tau=alg.s}^t \mathbb{1}[alg'_{m,i} \text{ is active at } \tau] \left(\tilde{f}_\tau^{alg'_{m,i}} - R_\tau \right) \\ &\quad (\text{Let } S_m = \{alg'_{m,1}, alg'_{m,2}, \dots\}) \\ &\leq \sum_{m=0}^n \sum_{i=1}^{|S_m|} \left(C_{\text{UCRL}}(|\mathcal{I}_{m,i}|; \bar{D}) + |\mathcal{I}_{m,i}| \Delta_{\mathcal{I}_{m,i};\bar{D}} + 2\bar{D}disc_{\mathcal{I}_{m,i}}^{alg'_{m,i}} \right) \\ &\quad (\text{by Lemma 11}) \end{aligned} \quad (14)$$

where in the last expression, we denote $\mathcal{I}_{m,i} = [alg'_{m,i}.s, alg'_{m,i}.e] \cap [alg.s, t]$ (the time within $[alg.s, t]$ where $alg'_{m,i}$ is scheduled), and $disc_{\mathcal{I}}^{alg'}$ is the total number of times within \mathcal{I} when alg' encounters state-reassignments.

For a fixed m , observe that all order- m instances are non-overlapping. Also, the aggregated number of state re-assignment for all order- m instances started within $[alg.s, t]$ is upper bounded by the total number of new instances of order not larger than $m - 1$ started within $[alg.s, t]$. The latter is further upper bounded by $6\hat{n} \log(T/\delta) \frac{C_{UCRL}(t'; \bar{D})}{C_{UCRL}(1; \bar{D})}$ according to the last claim of [Lemma 3](#). In other words, for every m , with probability $1 - \frac{\delta}{T}$,

$$\sum_{i=1}^{|S_m|} disc_{\mathcal{I}_{m,i}}^{alg'_{m,i}} = 6\hat{n} \log(T/\delta) \frac{C_{UCRL}(t'; \bar{D})}{C_{UCRL}(1; \bar{D})}.$$

Following the same calculation as in [Eq. \(10\)](#), [Eq. \(11\)](#) and [Eq. \(12\)](#), we also have that for every m , with probability $1 - \frac{\delta}{T}$,

$$\sum_{i=1}^{|S_m|} \left(C_{UCRL}(|\mathcal{I}_{m,i}|; \bar{D}) + |\mathcal{I}_{m,i}| \Delta_{\mathcal{I}_{m,i}; \bar{D}} \right) \leq 6\hat{n} \log(T/\delta) C_{UCRL}(t'; \bar{D}) + t' \Delta_{[alg.s, t]; \bar{D}}.$$

Using the above two bounds in [Eq. \(14\)](#), we get

$$\begin{aligned} \sum_{\tau=alg.s}^t (\tilde{g}_\tau - R_\tau) &= 6\hat{n} \log(T/\delta) C_{UCRL}(t'; \bar{D}) + \hat{n} t' \Delta_{[alg.s, t]; \bar{D}} + 2\bar{D} \times 6\hat{n} \log(T/\delta) \frac{C_{UCRL}(t'; \bar{D})}{C_{UCRL}(1; \bar{D})} \\ &= 18\hat{n} \log(T/\delta) C_{UCRL}(t'; \bar{D}) + \hat{n} t' \Delta_{[alg.s, t]; \bar{D}} \end{aligned}$$

where we use $C_{UCRL}(1; \bar{D}) \geq \bar{D}$ (by the definition of $\rho_{UCRL}(\cdot, \bar{D})$ in [Lemma 11](#)). Dividing both sides by t' finishes the proof. \blacksquare

Appendix C. Single-block Regret Analysis I

In this section, we focus on the regret in a block of index n . The analysis applies to both the standard case ([Lemma 4](#)), and the infinite-horizon RL case summarized in the following lemma.

Lemma 13 (c.f. [Lemma 4](#)) *In a block of index n that starts from t_n and ends on E_n (E_n could be equal to $t_n + 2^n - 1$, or smaller, if any stationarity test fails or MUCRL terminates), we have*

$$\sum_{\tau=t_n}^{E_n} (f_\tau^* - R_\tau) \leq \tilde{O} \left(\sum_{i=1}^{\ell} C_{UCRL}(|\mathcal{I}'_i|; \bar{D}) + \sum_{m=0}^n \frac{\rho(2^m; \bar{D})}{\rho(2^n; \bar{D})} C_{UCRL}(2^m; \bar{D}) \right)$$

where $\mathcal{I}'_1, \dots, \mathcal{I}'_\ell$ are intervals that partition $[t_n, E_n]$ such that $\Delta_{\mathcal{I}'_i; \bar{D}} \leq \rho_{UCRL}(|\mathcal{I}'_i|; \bar{D})$ for all i .

Throughout this section, if infinite-horizon RL is considered, $\rho(\cdot) \triangleq \rho_{UCRL}(\cdot; \bar{D})$, $\hat{\rho}(\cdot) \triangleq \hat{\rho}_{UCRL}(\cdot; \bar{D})$, $\Delta_{\mathcal{I}} \triangleq \Delta_{\mathcal{I}; \bar{D}} = \Delta_{\mathcal{I}}^r + 2\bar{D}\Delta_{\mathcal{I}}^p + \Delta_{\mathcal{I}}^J$ with a fixed \bar{D} , and $f_t^* \triangleq J_t^*$.

For the purpose of conducting analysis, we divide $[t_n, t_n + 2^n - 1]$ into consecutive intervals $\mathcal{I}_1 = [s_1, e_1], \mathcal{I}_2 = [s_2, e_2], \dots, \mathcal{I}_K = [s_K, e_K]$ ($s_1 = t_n, e_i + 1 = s_{i+1}, e_K = t_n + 2^n - 1$) in a way such that for all i :

$$\Delta_{\mathcal{I}_i} \leq \rho(|\mathcal{I}_i|) \tag{15}$$

One simple way to divide the intervals is to let $\Delta_{\mathcal{I}_i} = 0$ in each \mathcal{I}_i . Then the number of intervals K would be upper bounded by the number of stationary intervals within $[t_n, t_n + 2^n - 1]$. Intuitively, the number of intervals can also be related to $\Delta_{[t_n, t_n + 2^n - 1]}$. We defer the calculation of the required number of intervals to [Lemma 19](#). For now, we only need the fact that the partition satisfies [Eq. \(15\)](#). From a high level, this partition makes the distribution in each interval close to stationary. Notice that this partition is independent of the learner's behavior in block n .

For convenience, we further define the following quantities that depend on the learner's behavior in block n :

Definition 14 Define E_n as the index of the last round in block n . Since the block might terminate earlier than planned, we have $E_n \leq t_n + 2^n - 1$. Let $\ell \in [K]$ be such that $E_n \in \mathcal{I}_\ell$ (that is, ℓ is the index of the interval where block n ends). Define $e'_i = \min\{e_i, E_n\}$ and $\mathcal{I}'_i = [s_i, e'_i]$ (therefore, $\mathcal{I}'_i = \emptyset$ for $i > \ell$).

Recall the definition of \hat{n} and $\hat{\rho}(t)$ from [Lemma 3](#) (or [Lemma 12](#)). For simplicity, we define $\alpha_m \triangleq \rho(2^m)$, $\hat{\alpha}_m \triangleq \hat{\rho}(2^m)$, and also $\hat{C}(t) \triangleq t\hat{\rho}(t)$. Furthermore, we define the following technical quantities.

Definition 15 For every $i \in \{1, \dots, K\}$, and every $m \in \{0, 1, \dots, n\}$, define

$$\tau_i(m) = \min \{ \tau \in \mathcal{I}'_i : f_\tau^* - \tilde{g}_\tau \geq 12\hat{\alpha}_m \};$$

that is, $\tau_i(m)$ is the first time τ in $\mathcal{I}'_i = \mathcal{I}_i \cap [t_n, E_n]$ such that $f_\tau^* - \tilde{g}_\tau$ exceeds $12\hat{\alpha}_m$. If such τ does not exist or \mathcal{I}'_i is empty, we let $\tau_i(m) = \infty$.

Besides, we define $\xi_i(m) = [e'_i - \tau_i(m) + 1]_+$ where $[a]_+ = \max\{0, a\}$ (which is the length of the interval $[\tau_i(m), e'_i]$ when $\tau_i(m)$ is not ∞).

The intuition for $\tau_i(m)$ and $\xi_i(m)$ is as follows. Suppose that block n has not ended at τ . If there exists some $\tau \in \mathcal{I}_i$ such that $f_\tau^* - \tilde{g}_\tau \geq 12\hat{\alpha}_m$ (which first happens at $\tau_i(m)$), and if \mathcal{I}_i is long enough (i.e., $\xi_i(m)$ is large enough) so that after $\tau_i(m)$, an order- m instance of ALG can run entirely within \mathcal{I}_i , then the learner is able to discover the fact that $f_\tau^* - \tilde{g}_\tau$ is large, and then restart. This coincides with our explanation in [Figure 1](#). The derivation in this section will formalize this intuition.

Lemma 16 Let the high-probability events described in [Lemma 3](#) (or [Lemma 12](#)) hold. Then with high probability,

$$\begin{aligned} \sum_{\tau=t_n}^{E_n} (\tilde{g}_\tau - R_\tau) &\leq 4\hat{C}(2^n), \\ \sum_{\tau=t_n}^{E_n} (f_\tau^* - \tilde{g}_\tau) &\leq 96\hat{n} \sum_{i=1}^{\ell} \hat{C}(|\mathcal{I}'_i|) + 60 \sum_{m=0}^n \frac{\alpha_m}{\alpha_n} \hat{C}(2^m) \log(T/\delta) \end{aligned}$$

(notations are defined at the beginning of this section).

Proof $\sum_{\tau=t_n}^{E_n} (\tilde{g}_\tau - R_\tau)$ is trivially upper bounded by $3\hat{C}(E_n - t_n + 1) + 1 \leq 4\hat{C}(2^n)$ because it is guarded by [Test 2](#). Below we focus on the second claim.

Note that we can write for all $i = 1, \dots, K$,

$$\begin{aligned}
 & \sum_{\tau \in \mathcal{I}'_i} (f_\tau^* - \tilde{g}_\tau) \\
 & \leq 12 \sum_{\tau \in \mathcal{I}'_i} \left(\mathbb{1} \left[f_\tau^* - \tilde{g}_\tau \leq 12\hat{\alpha}_n \right] \hat{\alpha}_n + \sum_{m=1}^n \mathbb{1} \left[12\hat{\alpha}_m < f_\tau^* - \tilde{g}_\tau \leq 12\hat{\alpha}_{m-1} \right] \hat{\alpha}_{m-1} + \mathbb{1} \left[f_\tau^* - \tilde{g}_\tau > 12\hat{\alpha}_0 \right] 1 \right) \\
 & \leq 12 \left(|\mathcal{I}'_i| \hat{\alpha}_n + \sum_{m=1}^n \hat{\alpha}_{m-1} \xi_i(m) + \rho(1) \xi_i(0) \right) \quad (\rho(1) \geq 1 \text{ by Assumption 1}) \\
 & \leq 12 |\mathcal{I}'_i| \hat{\alpha}_n + 24 \sum_{m=0}^n \hat{\alpha}_m \xi_i(m) \quad (\hat{\alpha}_m = \frac{\hat{C}(2^m)}{2^m} \leq \frac{\hat{C}(2^{m+1})}{2^m} = 2\hat{\alpha}_{m+1})
 \end{aligned}$$

where in the second-to-last inequality we use $\sum_{\tau \in \mathcal{I}'_i} \mathbb{1} \left[f_\tau^* - \tilde{g}_\tau \geq 12\hat{\alpha}_m \right] = \sum_{\tau \in [\tau_i(m), e'_i]} \mathbb{1} \left[f_\tau^* - \tilde{g}_\tau \geq 12\hat{\alpha}_m \right] \leq \xi_i(m)$ by the definition of $\tau_i(m)$.

Summing the above over intervals i and notice that $\sum_{i=1}^\ell |\mathcal{I}'_i| \leq 2^n$, we get

$$\sum_{\tau=t_n}^{E_n} (f_\tau^* - \tilde{g}_\tau) \leq 12 \cdot 2^n \hat{\alpha}_n + 24 \sum_{m=0}^n \sum_{i=1}^\ell \hat{\alpha}_m \xi_i(m) = 12\hat{C}(2^n) + 24 \sum_{m=0}^n \sum_{i=1}^\ell \hat{\alpha}_m \xi_i(m). \quad (16)$$

Next, we upper bound $\sum_{i=1}^\ell \hat{\alpha}_m \xi_i(m)$ for each m .

$$\begin{aligned}
 \sum_{i=1}^\ell \hat{\alpha}_m \xi_i(m) &= \sum_{i=1}^\ell \hat{\alpha}_m \min \{ \xi_i(m), 4 \cdot 2^m \} + \sum_{i=1}^\ell \hat{\alpha}_m [\xi_i(m) - 4 \cdot 2^m]_+. \quad (17) \\
 & \quad \text{(using } a = \min\{a, b\} + [a - b]_+)
 \end{aligned}$$

The first term on the right-hand side of Eq. (17) can be bounded as below:

$$\begin{aligned}
 \sum_{i=1}^\ell \hat{\alpha}_m \min \{ \xi_i(m), 4 \cdot 2^m \} &\leq 4 \sum_{i=1}^\ell \hat{\rho}(2^m) \times \min \{ \xi_i(m), 2^m \} \\
 &\leq 4 \sum_{i=1}^\ell \hat{\rho}(\min \{ \xi_i(m), 2^m \}) \times \min \{ \xi_i(m), 2^m \} \\
 & \quad (\hat{\rho}(\cdot) \text{ is a decreasing function}) \\
 &= 4 \sum_{i=1}^\ell \hat{C}(\min \{ \xi_i(m), 2^m \}) \\
 &\leq 4 \sum_{i=1}^\ell \hat{C}(|\mathcal{I}'_i|). \quad (\hat{C}(\cdot) \text{ is an increasing function})
 \end{aligned}$$

The second term on the right-hand side of Eq. (17) is bounded using Lemma 17 below. Combining them into Eq. (16) finishes the proof. \blacksquare

Lemma 17 *Let the high probability events described in Lemma 3 (or Lemma 12) hold. Then with high probability,*

$$\sum_{i=1}^{\ell} \hat{\alpha}_m [\xi_i(m) - 4 \cdot 2^m]_+ \leq \frac{2\alpha_m}{\alpha_n} \hat{C}(2^m) \log(T/\delta).$$

Proof Using the fact that $[[a]_+ - b]_+ = [a - b]_+$ when $b \geq 0$, we have

$$[\xi_i(m) - 4 \cdot 2^m]_+ = [e'_i - \tau_i(m) + 1 - 4 \cdot 2^m]_+. \quad (18)$$

Next, we consider the following quantity: “the number of rounds in the interval $[\tau_i(m), e'_i - 2 \cdot 2^m B]$ which are candidate starting points of an order- m ALG”. By Procedure 1, this quantity can be written and lower bounded as

$$A_i \triangleq \sum_{t \in \mathcal{I}_i} \mathbb{1} \left[t \in [\tau_i(m), e'_i - 2 \cdot 2^m], \quad (t - t_n) \bmod 2^m = 0 \right] \geq \frac{[e'_i - \tau_i(m) + 1 - 4 \cdot 2^m]_+}{2^m}$$

where we use the fact in an interval of length w , there are at least $\frac{w+2-2u}{u}$ points whose indices are multiples of u . Notice that the right-hand side is related to what we want to upper bound in the lemma according to Eq. (18). Thus we continue to upper bound the left-hand side above. We define the following events:

$$\begin{aligned} W_t &= \{ \tau_i(m) \leq t \leq e_i - 2 \cdot 2^m \text{ where } i \text{ is such that } t \in \mathcal{I}_i \}, \\ X_t &= \{ t \leq E_n - 2 \cdot 2^m \}, \\ Y_t &= \{ t \leq E_n \text{ and } (t - t_n) \bmod 2^m = 0 \}, \\ Z_t &= \{ \exists \text{ order-}m \text{ alg such that } \text{alg.s} = t \}, \\ V_t &= \{ \exists \tau \in [t_n, t] \text{ such that } W_\tau \cap Y_\tau \cap Z_\tau \}. \end{aligned}$$

Then we can write (recall the definition of K in the beginning of this section)

$$\sum_{i=1}^{\ell} A_i = \sum_{i=1}^K A_i = \sum_{t=t_n}^{t_n+2^n-1} \mathbb{1}[W_t, X_t, Y_t] \leq \underbrace{\sum_{t=t_n}^{t_n+2^n-1} \mathbb{1}[W_t, Y_t, \bar{V}_t]}_{\text{term}_3} + \underbrace{\sum_{t=t_n}^{t_n+2^n-1} \mathbb{1}[X_t, V_t]}_{\text{term}_4}$$

For **term**₃, notice that conditioned on $W_t \cap Y_t$, the event Z_t happens with a constant probability $\frac{\alpha_n}{\alpha_m}$ (by Procedure 1). Therefore, **term**₃ counts the number of trials up to the first success in a repeated trial with success probability $\frac{\alpha_n}{\alpha_m}$. Therefore, with probability $1 - \frac{\delta}{T}$, **term**₃ $\leq 1 + \frac{\log(T/\delta)}{-\log(1 - \frac{\alpha_n}{\alpha_m})} \leq \frac{2\alpha_m}{\alpha_n} \log(T/\delta)$.

Next, we deal with **term**₄. Below we show that **term**₄ = 0. The event V_t implies that there exists some order- m alg which starts at $\text{alg.s} = t^*$, where $t^* \leq t$ and $\tau_i(m) \leq t^* \leq e_i - 2 \cdot 2^m$. Therefore, we have $\text{alg.e} = \text{alg.s} + 2^m - 1 = t^* + 2^m - 1 \leq e_i - 2^m - 1 < e_i$, and thus $[\text{alg.s}, \text{alg.e}] \subseteq \mathcal{I}_i$. Together with X_t , the event $V_t \cap X_t$ implies that $\text{alg.e} = \text{alg.s} + 2^m - 1 \leq t + 2^m - 1 < E_n$, and therefore, and time alg.e , block n has not ended.

Since at time $alg.e$, block n is still on-going, the learner performs **Test 1**. By [Lemma 3](#) (or [Lemma 12](#) for the infinite-horizon RL case), with high probability, we have

$$\begin{aligned}
 \frac{1}{2^m} \sum_{\tau=alg.s}^{alg.e} R_\tau &\geq \frac{1}{2^m} \sum_{\tau=alg.s}^{alg.e} \tilde{g}_\tau - \hat{\alpha}_m - \hat{n} \Delta_{[alg.s, alg.e]} && \text{(Lemma 3 or Lemma 12)} \\
 &\geq \min_{\tau \in \mathcal{I}_i} f_\tau^* - \hat{\alpha}_m - (\hat{n} + 1) \Delta_{\mathcal{I}_i} && \text{(because } [alg.s, alg.e] \subseteq \mathcal{I}_i \text{)} \\
 &\geq f_{\tau_i(m)}^* - \hat{\alpha}_m - (\hat{n} + 3) \Delta_{\mathcal{I}_i} && (|\min_{\tau \in \mathcal{I}_i} f_\tau^* - f_{\tau_i(m)}^*| \leq 2\Delta_{\mathcal{I}_i}) \\
 &\geq \tilde{g}_{\tau_i(m)} + 12\hat{\alpha}_m - 2\hat{\alpha}_m && \text{(by the definition of } \tau_i(m) \text{ and } \Delta_{\mathcal{I}_i} \leq \rho(|\mathcal{I}_i|) \leq \rho(2^m) \leq \frac{\hat{\alpha}_m}{6\hat{n}} \text{)} \\
 &\geq U_{alg.e} + 10\hat{\alpha}_m && \text{(Because } alg.e \geq \tau_i(m), U_{alg.e} \leq \tilde{g}_{\tau_i(m)} \text{ by the algorithm)}
 \end{aligned}$$

This should trigger the restart at time $alg.e < E_n$, contradicting the definition of E_n . Therefore, $\mathbb{1}[X_t, V_t] = 0$.

Finally, combining all previous arguments, we have that with high probability,

$$\begin{aligned}
 \sum_{i=1}^{\ell} \hat{\alpha}_m [\xi_i(m) - 4 \cdot 2^m]_+ &= \sum_{i=1}^{\ell} \hat{\alpha}_m [e'_i - \tau_i(m) + 1 - 4 \cdot 2^m]_+ \leq \hat{\alpha}_m 2^m \sum_{i=1}^{\ell} A_i \\
 &= \hat{C}(2^m) \sum_{i=1}^{\ell} A_i \leq \frac{2\alpha_m}{\alpha_n} \hat{C}(2^m) \log(T/\delta),
 \end{aligned}$$

finishing the proof. ■

Appendix D. Single-block Regret Analysis II (under a Special Form of $C(\cdot)$)

In [Appendix C](#), we have derived the regret bound in a single block for both the standard setting and the infinite-horizon MDP setting ([Lemma 4](#) and [Lemma 13](#)). They are both of the form

$$\sum_{\tau \in \mathcal{J}} (f_\tau^* - R_\tau) = \tilde{\mathcal{O}} \left(\sum_{i=1}^{\ell} C(|\mathcal{I}'_i|) + \sum_{m=0}^n \frac{\rho(2^m)}{\rho(2^m)} C(2^m) \right). \quad (19)$$

(replacing $C(\cdot)$ and $\rho(\cdot)$ by $C_{UCRL}(\cdot; \bar{D})$ and $\rho(\cdot; \bar{D})$ for the case of infinite-horizon MDP).

In this section, we further derive more concrete dynamic regret bounds for both cases by assuming that $C(\cdot)$ is of some specific form. The form of $C(\cdot)$ we consider in this section is defined as follows:

Definition 18 *We define a form of $C(t)$ as $C(t) = \min\{c_1 t^p + c_2, c_3 t\}$ for some $p \in [\frac{1}{2}, 1)$ and some c_1, c_2, c_3 ($c_3 \geq 1$) that capture dependencies on $\log(T/\delta)$ and other problem-dependent constants.*

In fact, usually, a regret bound is only written in the form of $c_1 t^p + c_2$. However, since the reward is bounded between 0 and 1, the regret bound of $\min\{c_1 t^p + c_2, t\}$ is also trivially correct. [Definition 18](#) is slightly more general than this by allowing a coefficient $c_3 \geq 1$ (the regret bound would still be

trivially correct). In some cases, we make c_1, c_2, c_3 larger than their tightest possible values to make the final regret bound better — notice that the choice of c_1, c_2, c_3 affects the probability specified in [Procedure 1](#), and thus smaller c_1, c_2, c_3 does not necessarily make the final regret bound smaller. This subtle issue can be observed from the analysis.

To get a concrete bound, we also need to decide the number ℓ in the single-block regret bound above. In [Appendix C](#), we have stated the condition (i.e., [Eq. \(15\)](#)) that should be satisfied by $\mathcal{I}'_1, \dots, \mathcal{I}'_\ell$ (or $\mathcal{I}_1, \dots, \mathcal{I}_K$). In the next lemma, we upper bound the value of ℓ that is required to fulfill the condition.

Lemma 19 *Let $\mathcal{J} = [t_n, E_n]$. Then we have $\ell \leq L_{\mathcal{J}}$. Furthermore, if $C(t)$ is in the form specified in [Definition 18](#), we also have $\ell \leq 1 + 2 \left(c_1^{-1} \Delta_{\mathcal{J}} |\mathcal{J}|^{1-p} \right)^{\frac{1}{2-p}} + c_3^{-1} \Delta_{\mathcal{J}}$.*

Proof The fact that $\ell \leq L_{\mathcal{J}}$ is straightforward to see (and has been explained in [Appendix C](#)): to satisfy the condition [Eq. \(15\)](#), one way to divide the block is to make each \mathcal{I}_i a stationary interval, which makes $\Delta_{\mathcal{I}_i} = 0$ for all $i \in [K]$. This way of division leads to $\ell \leq L_{\mathcal{J}}$.

For the second claim, we follow the same procedure as described in the proof of [Lemma 5](#) in [\(Chen et al., 2019\)](#). Basically, the procedure divides $[t_n, t_n + 2^n - 1]$ in a greedy way, making all $\mathcal{I}_i = [s_i, e_i]$ satisfy $\Delta_{[s_i, e_i]} \leq \rho(e_i - s_i + 1)$ and $\Delta_{[s_i, e_i+1]} > \rho(e_i - s_i + 2)$ for all $i \in [K - 1]$ (i.e., except for the last interval). Then we have

$$\begin{aligned}
 \Delta_{\mathcal{J}} &\geq \sum_{i=1}^{\ell-1} \Delta_{[s_i, e_i+1]} && \text{(by the definition of } \Delta_{[\cdot, \cdot]}\text{)} \\
 &> \sum_{i=1}^{\ell-1} \rho(e_i - s_i + 2) \\
 &\geq \sum_{i=1}^{\ell-1} \min \{ c_1 (e_i - s_i + 2)^{p-1}, c_3 \} && \text{(by Definition 18)} \\
 &\geq \sum_{i=1}^{\ell-1} \min \left\{ \frac{1}{2} c_1 (e_i - s_i + 1)^{p-1}, c_3 \right\} \\
 &\quad \left((x+2)^{p-1} \geq (2(x+1))^{p-1} \geq \frac{1}{2} (x+1)^{p-1} \text{ for any } x \geq 0 \text{ and } p \leq 1 \right) \\
 &= \frac{1}{2} \sum_{i=1}^{\ell_1} c_1 (e_i - s_i + 1)^{p-1} + \sum_{i=1}^{\ell_2} c_3
 \end{aligned}$$

where in the last equality we separate the intervals where $\min \left\{ \frac{1}{2} c_1 (e_i - s_i + 1)^{p-1}, c_3 \right\}$ takes the former or the latter value. Note that $\ell_1 + \ell_2 = \ell - 1$.

The above inequality implies that $\Delta_{\mathcal{J}}$ upper bounds both $\frac{1}{2} \sum_{i=1}^{\ell_1} c_1 (e_i - s_i + 1)^{p-1}$ and $\sum_{i=1}^{\ell_2} c_3$. Thus, $\ell_2 \leq c_3^{-1} \Delta_{\mathcal{J}}$, and by Hölder's inequality,

$$\ell_1 \leq \left(\sum_{i=1}^{\ell_1} (e_i - s_i + 1)^{p-1} \right)^{\frac{1}{2-p}} \left(\sum_{i=1}^{\ell_1} (e_i - s_i + 1) \right)^{\frac{1-p}{2-p}} \leq \left(\frac{2\Delta_{\mathcal{J}}}{c_1} \right)^{\frac{1}{2-p}} |\mathcal{J}|^{\frac{1-p}{2-p}}.$$

Combining them finishes the proof. ■

In the following [Lemma 20](#), we bound the regret within a block by combining [Eq. \(19\)](#) and [Lemma 19](#). We will frequently use the following two properties: let $\{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_K\}$ be a partition of the interval \mathcal{S} . Then

$$\sum_{i=1}^K L_{\mathcal{S}_i} \leq L_{\mathcal{S}} + (K - 1), \quad (20)$$

$$\sum_{i=1}^K \Delta_{\mathcal{S}_i} \leq \Delta_{\mathcal{S}}. \quad (21)$$

They can be derived using the definitions of $L_{[\cdot, \cdot]}$ and $\Delta_{[\cdot, \cdot]}$.

Lemma 20 *If $C(t)$ is of the form specified in [Definition 18](#), then*

$$\sum_{\tau=t_n}^{E_n} (f_{\tau}^* - R_{\tau}) \leq \tilde{\mathcal{O}} \left(\min \left\{ \text{Reg}_L(\mathcal{J}), \text{Reg}_{\Delta}(\mathcal{J}) \right\} + c_1 2^{np} + \frac{c_2 c_3}{c_1} 2^{n(1-p)} + \frac{c_2^2}{c_3} \right),$$

where $\text{Reg}_L(\mathcal{J}) \triangleq c_1 L_{\mathcal{J}}^{1-p} |\mathcal{J}|^p + c_2 L_{\mathcal{J}}$ and

$$\text{Reg}_{\Delta}(\mathcal{J}) \triangleq \left(c_1 \Delta_{\mathcal{J}}^{1-p} |\mathcal{J}| \right)^{\frac{1}{2-p}} + c_1 |\mathcal{J}|^p + c_1 (c_3^{-1} \Delta_{\mathcal{J}})^{1-p} |\mathcal{J}|^p + c_2 \left(c_1^{-1} \Delta_{\mathcal{J}} |\mathcal{J}|^{1-p} \right)^{\frac{1}{2-p}} + c_2 + c_2 c_3^{-1} \Delta_{\mathcal{J}}.$$

Proof We bound each term in [Eq. \(19\)](#) using [Definition 18](#). First, notice that

$$\begin{aligned} \tilde{\mathcal{O}} \left(\sum_{i=1}^{\ell} C(|\mathcal{I}'_i|) \right) &= \tilde{\mathcal{O}} \left(\sum_{i=1}^{\ell} \min \{ c_1 |\mathcal{I}'_i|^p + c_2, c_3 t \} \right) \\ &\leq \tilde{\mathcal{O}} \left(\sum_{i=1}^{\ell} (c_1 |\mathcal{I}'_i|^p + c_2) \right) \leq \tilde{\mathcal{O}} (c_1 \ell^{1-p} |\mathcal{J}|^p + c_2 \ell). \end{aligned} \quad (22)$$

Using the first upper bound for ℓ given in [Lemma 19](#), [Eq. \(22\)](#) can be bounded by $\tilde{\mathcal{O}}(\text{Reg}_L(\mathcal{J}))$; using the second upper bound, [Eq. \(22\)](#) can be bounded by $\tilde{\mathcal{O}}(\text{Reg}_{\Delta}(\mathcal{J}))$. Next, we have

$$\tilde{\mathcal{O}} \left(\frac{\rho(2^m)}{\rho(2^n)} C(2^m) \right) = \tilde{\mathcal{O}} \left(c_1 2^{np} + \frac{c_2 c_3}{c_1} 2^{n(1-p)} + \frac{c_1^2}{c_3} 2^{m(2p-1)} + \frac{c_2^2}{c_3} 2^{-m} \right).$$

by [Lemma 21](#) below. Notice that because $c_3 \geq 1$ and $p \geq \frac{1}{2}$, $\frac{c_1^2}{c_3} 2^{m(2p-1)} \leq c_1^2 2^{n(2p-1)} \leq c_1 2^{np}$ when $c_1 \leq 2^{n(1-p)}$. This is indeed the regime we care about since if $c_1 > 2^{n(1-p)}$ then the first term $c_1 2^{np} > 2^n$, which is a vacuous bound for the regret of block n . Therefore, we can drop this term. Thus, the dynamic regret in block n can be summarized as the following based on [Eq. \(19\)](#):

$$\tilde{\mathcal{O}} \left(\min \left\{ \text{Reg}_L(\mathcal{J}), \text{Reg}_{\Delta}(\mathcal{J}) \right\} + c_1 2^{np} + \frac{c_2 c_3}{c_1} 2^{n(1-p)} + \frac{c_2^2}{c_3} \right), \quad (23)$$

finishing the proof. ■

Lemma 21 *Let $C(t)$ be of the form in [Definition 18](#). Then*

$$\frac{\rho(2^m)}{\rho(2^n)} C(2^m) = \mathcal{O} \left(c_1 2^{np} + \frac{c_2 c_3}{c_1} 2^{n(1-p)} + \frac{c_1^2}{c_3} 2^{m(2p-1)} + \frac{c_2^2}{c_3} 2^{-m} \right).$$

Proof This is by direct calculation:

$$\begin{aligned} \frac{\rho(2^m)}{\rho(2^n)} C(2^m) &= \frac{C(2^m)^2}{C(2^n)} 2^{n-m} \\ &= \mathcal{O} \left(\frac{\min\{c_1^2 2^{2mp} + c_2^2, c_3^2 2^{2m}\}}{c_1 2^{np} + c_2} 2^{n-m} + \frac{\min\{c_1^2 2^{2mp} + c_2^2, c_3^2 2^{2m}\}}{c_3 2^n} 2^{n-m} \right) \\ &= \mathcal{O} \left(\min \left\{ c_1 2^{np} 2^{(n-m)(1-2p)} + \frac{c_2^2}{c_1} 2^{n(1-p)-m}, \frac{c_3^2}{c_1} 2^{n(1-p)+m} \right\} + \frac{c_1^2}{c_3} 2^{m(2p-1)} + \frac{c_2^2}{c_3} 2^{-m} \right) \\ &= \mathcal{O} \left(c_1 2^{np} + \min \left\{ \frac{c_2^2}{c_1} 2^{n(1-p)-m}, \frac{c_3^2}{c_1} 2^{n(1-p)+m} \right\} + \frac{c_1^2}{c_3} 2^{m(2p-1)} + \frac{c_2^2}{c_3} 2^{-m} \right) \\ &= \mathcal{O} \left(c_1 2^{np} + \frac{c_2 c_3}{c_1} 2^{n(1-p)} + \frac{c_1^2}{c_3} 2^{m(2p-1)} + \frac{c_2^2}{c_3} 2^{-m} \right). \end{aligned}$$

■

Appendix E. Single-epoch Regret Analysis

We call $[t_0, E]$ an *epoch* if t_0 is the first step after restart (or $t_0 = 1$), and E is the first time after round t_0 when the restart is triggered. In this section, we continue the discussion in [Appendix D](#) and bound the regret in a single epoch. Recall that we consider cases where the single-block regret can be written as [Eq. \(19\)](#) and $C(\cdot)$ is in the form of [Definition 18](#). This holds both for the case of the standard setting and the infinite-horizon MDP setting.

Lemma 22 *Let \mathcal{E} be an epoch. Then*

$$\sum_{\tau \in \mathcal{E}} \leq \tilde{\mathcal{O}} \left(\min \{ \text{Reg}_L(\mathcal{E}), \text{Reg}_\Delta(\mathcal{E}) \} + \frac{c_2 c_3}{c_1} |\mathcal{E}|^{1-p} + \frac{c_2^2}{c_3} \right)$$

($\text{Reg}_L(\cdot)$ and $\text{Reg}_\Delta(\cdot)$ are defined in [Lemma 20](#))

Proof Let \mathcal{E} be an epoch whose last block is indexed by n . Then $|\mathcal{E}| = \Theta(2^n)$. Let $\mathcal{J}_1, \dots, \mathcal{J}_n$ be blocks in \mathcal{E} . Then by [Lemma 20](#), the dynamic regret in \mathcal{E} is upper bounded by

$$\tilde{\mathcal{O}} \left(\min \left\{ \sum_{m=0}^n \text{Reg}_L(\mathcal{J}_m), \sum_{m=0}^n \text{Reg}_\Delta(\mathcal{J}_m) \right\} + c_1 \sum_{m=0}^n 2^{mp} + \frac{c_2 c_3}{c_1} \sum_{m=0}^n 2^{m(1-p)} + \sum_{m=0}^n \frac{c_2^2}{c_3} \right).$$

By Hölder's inequality,

$$\sum_{m=0}^n \text{Reg}_L(\mathcal{J}_m) = c_1 \left(\sum_{m=0}^n L_{\mathcal{J}_m} \right)^{1-p} \left(\sum_{m=0}^n |\mathcal{J}_m| \right)^p + c_2 \sum_{m=0}^n L_{\mathcal{J}_m}$$

$$\begin{aligned}
 &\leq c_1 (L_{\mathcal{E}} + n)^{1-p} |\mathcal{E}|^p + c_2 (L_{\mathcal{E}} + n) && \text{(using Eq. (20))} \\
 &\leq \tilde{\mathcal{O}} \left(c_1 L_{\mathcal{E}}^{1-p} |\mathcal{E}|^p + c_2 L_{\mathcal{E}} \right) = \tilde{\mathcal{O}} (\text{Reg}_L(\mathcal{E})) \quad \text{(because } n = \mathcal{O}(\log T) = \tilde{\mathcal{O}}(1))
 \end{aligned}$$

Similarly, $\sum_{m=0}^n \text{Reg}_{\Delta}(\mathcal{J}_m) = \tilde{\mathcal{O}}(\text{Reg}_{\Delta}(\mathcal{E}))$. On the other hand, $c_1 \sum_{m=0}^n 2^{mp} + \frac{c_2 c_3}{c_1} \sum_{m=0}^n 2^{m(1-p)} + \sum_{m=0}^n \frac{c_2^2}{c_3} = \tilde{\mathcal{O}} \left(c_1 2^{np} + \frac{c_2 c_3}{c_1} 2^{n(1-p)} + \frac{c_2^2}{c_3} \right) = \tilde{\mathcal{O}} \left(c_1 |\mathcal{E}|^p + \frac{c_2 c_3}{c_1} |\mathcal{E}|^{1-p} + \frac{c_2^2}{c_3} \right)$. In summary, the dynamic regret within an epoch is of order

$$\tilde{\mathcal{O}} \left(\min \{ \text{Reg}_L(\mathcal{E}), \text{Reg}_{\Delta}(\mathcal{E}) \} + \frac{c_2 c_3}{c_1} |\mathcal{E}|^{1-p} + \frac{c_2^2}{c_3} \right) \quad (24)$$

(the $c_1 |\mathcal{E}|^p$ term is absorbed into $\min \{ \text{Reg}_L(\mathcal{E}), \text{Reg}_{\Delta}(\mathcal{E}) \}$). \blacksquare

Appendix F. Proof of Theorem 2

We are now ready to prove Theorem 2 after showing the following two lemmas.

Lemma 23 *Let t be in an epoch starting from t_0 . If $\Delta_{[t_0, t]} \leq \rho(t - t_0 + 1)$, then with high probability, no restart would be triggered at time t .*

Proof We first verify that **Test 1** would not fail with high probability. Let $t = \text{alg.e}$ where *alg* is any order- m ALG in block n . Then with high probability,

$$\begin{aligned}
 U_t &= \min_{\tau \in [t_n, t]} \tilde{g}_{\tau} \\
 &\geq \min_{\tau \in [t_n, t]} f_{\tau}^* - \Delta_{[t_n, t]} && \text{(by Lemma 3)} \\
 &\geq \frac{1}{2^m} \sum_{\tau \in [\text{alg.s}, t]} f_{\tau}^* - 3\Delta_{[t_n, t]} && ([\text{alg.s}, t] \subseteq [t_n, t]) \\
 &\geq \frac{1}{2^m} \sum_{\tau \in [\text{alg.s}, t]} R_{\tau} - 2\sqrt{\frac{\log(T/\delta)}{2^m}} - 3\rho(t - t_0 + 1) \\
 &\quad (\mathbb{E}[R_{\tau}] = \mathbb{E}[f_{\tau}(\pi_t)] \leq f_{\tau}^* \text{ and we use Azuma's inequality}) \\
 &\geq \frac{1}{2^m} \sum_{\tau \in [\text{alg.s}, t]} R_{\tau} - \hat{\rho}(2^m) - 3\rho(t - t_0 + 1) \\
 &\quad \text{(By Assumption 1, } \hat{\rho}(2^m) \geq 6 \log(T/\delta) \rho(2^m) \geq 6 \log(T/\delta) \sqrt{\frac{1}{2^m}}) \\
 &\geq \frac{1}{2^m} \sum_{\tau \in [\text{alg.s}, t]} R_{\tau} - 2\hat{\rho}(2^m). \quad (\rho(t - t_0 + 1) \leq \rho(2^m) \text{ because } \rho(\cdot) \text{ is decreasing})
 \end{aligned}$$

So with high probability, **Test 1** will not return fail.

Furthermore, by Lemma 3, with high probability,

$$\frac{1}{t - t_n + 1} \sum_{\tau=t_n}^t (\tilde{g}_{\tau} - R_{\tau}) \leq \hat{\rho}(t - t_n + 1) + \Delta_{[t_n, t]} \leq 2\hat{\rho}(t - t_n + 1).$$

Therefore, with high probability, **Test 2** will not return fail either. \blacksquare

Lemma 24 *With high probability, the number of epochs is upper bounded by L . If $C(\cdot)$ is in the form of [Definition 18](#), the number of epochs is also upper bounded by $1 + 2 \left(c_1^{-1} \Delta T^{1-p} \right)^{\frac{1}{2-p}} + c_3^{-1} \Delta$.*

Proof By [Lemma 23](#), if $[t_0, E]$ is not the last epoch, then $\Delta_{[t_0, E]} > \rho(E - t_0 + 1)$ with high probability. Then following the exact same arguments as in [Lemma 19](#) proves the lemma. \blacksquare

Proof [Proof of [Theorem 2](#)] If $C(t) = c_1 t^p + c_2$ satisfies [Assumption 1](#), then $C(t) = \min\{c_1 t^p + c_2, t\}$ also satisfies it (since the reward is bounded in $[0, 1]$). Below we use $C(t) = \min\{c_1 t^p + c_2, t\}$ as the input to our algorithm. Notice that this is in the form of [Definition 18](#) with $c_3 = 1$. Let $\mathcal{E}_1, \dots, \mathcal{E}_N$ be epochs in $[1, T]$. Then by [Lemma 22](#), the dynamic regret in $[1, T]$ is upper bounded by

$$\tilde{\mathcal{O}} \left(\min \left\{ \sum_{i=1}^N \text{Reg}_L(\mathcal{E}_i), \sum_{i=1}^N \text{Reg}_\Delta(\mathcal{E}_i) \right\} + \frac{c_2}{c_1} \sum_{i=1}^N |\mathcal{E}_i|^{1-p} + c_2^2 N \right). \quad (25)$$

By Hölder's inequality and [Eq. \(20\)](#),

$$\sum_{i=1}^N \text{Reg}_L(\mathcal{E}_i) \leq \tilde{\mathcal{O}} \left(c_1 (L + N - 1)^{1-p} T^p + c_2 (L + N - 1) \right) \leq \tilde{\mathcal{O}} (c_1 L^{1-p} T^p + c_2 L),$$

where in the last inequality we use [Lemma 24](#) to bound N .

Similarly,

$$\begin{aligned} & \sum_{i=1}^N \text{Reg}_\Delta(\mathcal{E}_i) \\ & \leq \tilde{\mathcal{O}} \left((c_1 \Delta^{1-p} T)^{\frac{1}{2-p}} + c_1 N^{1-p} T^p + c_1 \Delta^{1-p} T^p + c_2 (c_1^{-1} \Delta T^{1-p})^{\frac{1}{2-p}} + c_2 N + c_2 \Delta \right) \\ & \leq \tilde{\mathcal{O}} \left((c_1 \Delta^{1-p} T)^{\frac{1}{2-p}} + c_1 T^p + c_1 \Delta^{1-p} T^p + c_2 (c_1^{-1} \Delta T^{1-p})^{\frac{1}{2-p}} + c_2 + c_2 \Delta \right). \end{aligned}$$

(using [Lemma 24](#) to bound N)

Then we deal with the second term in [Eq. \(25\)](#):

$$\frac{c_2}{c_1} \sum_{i=1}^N |\mathcal{E}_i|^{1-p} \leq \frac{c_2}{c_1} N^p T^{1-p},$$

which can be either bounded by $\tilde{\mathcal{O}} \left(\frac{c_2}{c_1} L^p T^{1-p} \right)$ or

$$\tilde{\mathcal{O}} \left(\frac{c_2}{c_1} T^{1-p} + \frac{c_2}{c_1} \left(c_1^{-p} \Delta^p T^{2-2p} \right)^{\frac{1}{2-p}} + \frac{c_2}{c_1} \Delta^p T^{1-p} \right)$$

using the upper bound for N in [Lemma 24](#). Finally, the third term in [Eq. \(25\)](#) can be upper bounded either by $\tilde{\mathcal{O}}(c_2^2 L)$ or

$$\tilde{\mathcal{O}}\left(c_2^2 + c_2^2 (c_1^{-1} \Delta T^{1-p})^{\frac{1}{2-p}} + c_2^2 \Delta\right).$$

With all terms expanded, below, we collect the dominant terms for the cases of $p = \frac{1}{2}$ and $p > \frac{1}{2}$. We say term $a(T)$ is dominated by $b(T)$ if $\lim_{T \rightarrow \infty} a(T)/b(T) = 0$ under any sublinear growth rate of L or Δ (e.g., $\sqrt{\Delta T}$ is dominated by $\Delta^{1/3} T^{2/3}$ and L is dominated by \sqrt{LT}). And below we only write down terms that are not dominated by other terms.

The case for $p = \frac{1}{2}$:

$$\tilde{\mathcal{O}}\left(\min\left\{\left(c_1 + \frac{c_2}{c_1}\right)\sqrt{LT}, \left(c_1^{2/3} + c_2 c_1^{-4/3}\right)\Delta^{1/3} T^{2/3} + \left(c_1 + \frac{c_2}{c_1}\right)\sqrt{T}\right\}\right);$$

The case for $p > \frac{1}{2}$:

$$\tilde{\mathcal{O}}\left(\min\left\{c_1 L^{1-p} T^p, (c_1 \Delta^{1-p} T)^{\frac{1}{2-p}} + c_1 T^p\right\}\right).$$

This finishes the proof. ■

Appendix G. Main Results for Infinite-horizon MDP

Lemma 25 (c.f. [Lemma 23](#)) *Let t be in an epoch started from round t_0 . If $\Delta_{[t_0, t]; \bar{D}} < \bar{D} S \sqrt{\frac{A}{t-t_0+1}}$ and $\bar{D} \geq D_{\max}$, then with high probability, no restart will be triggered at time t .*

Proof To verify that **Test 1** will not fail with high probability, we follow very similar steps as in [Lemma 23](#). Let $t = \text{alg.e}$ where *alg* is an order- m ALG in block n . Then with high probability (the following calculation is same as that in the proof of [Lemma 23](#) except for the third inequality),

$$\begin{aligned} U_t &= \min_{\tau \in [t_n, t]} \tilde{g}_\tau \\ &\geq \min_{\tau \in [t_n, t]} J_\tau^* - \Delta_{[t_n, t]; \bar{D}} && \text{(by [Lemma 12](#))} \\ &\geq \frac{1}{2^m} \sum_{\tau \in [\text{alg.s}, t]} J_\tau^* - 3\Delta_{[t_n, t]; \bar{D}} && ([\text{alg.s}, t] \subseteq [t_0, t]) \\ &\geq \frac{1}{2^m} \sum_{\tau \in [\text{alg.s}, t]} R_\tau - 4\bar{D} \sqrt{\frac{\log(T/\delta)}{2^m}} - 4\Delta_{[t_n, t]; \bar{D}} && \text{(explained below)} \\ &\geq \frac{1}{2^m} \sum_{\tau \in [\text{alg.s}, t]} R_\tau - 4\hat{\rho}_{\text{UCRL}}(2^m; \bar{D}) - 4\rho_{\text{UCRL}}(t - t_0 + 1; \bar{D}) \\ &\geq \frac{1}{2^m} \sum_{\tau \in [\text{alg.s}, t]} R_\tau - 8\hat{\rho}_{\text{UCRL}}(2^m; \bar{D}). \quad (\rho(t - t_0 + 1) \leq \rho(2^m) \text{ because } \rho(\cdot) \text{ is decreasing}) \end{aligned}$$

where the third inequality is based on the following calculation: for all $\tau \in [\text{alg}.s, t]$,

$$\begin{aligned}
 J_t^* &= r_t(s_\tau, a_\tau) + \sum_{s'} p_t(s'|s_\tau, a_\tau) h_t^*(s') - h_t^*(s_\tau) \\
 &\geq r_\tau(s_\tau, a_\tau) + \sum_{s'} p_\tau(s'|s_\tau, a_\tau) h_t^*(s') - h_t^*(s_\tau) - \left(\Delta_{[\text{alg}.s, t]}^r + D_{\max} \Delta_{[\text{alg}.s, t]}^p \right) \\
 &\geq r_\tau(s_\tau, a_\tau) + \sum_{s'} p_\tau(s'|s_\tau, a_\tau) h_t^*(s') - h_t^*(s_\tau) - \left(\Delta_{[\text{alg}.s, t]}^r + \bar{D} \Delta_{[\text{alg}.s, t]}^p \right) \\
 &\hspace{15em} \text{(by the assumption } \bar{D} \geq D_{\max} \text{)}
 \end{aligned}$$

and thus

$$\begin{aligned}
 \frac{1}{2^m} \sum_{\tau \in [\text{alg}.s, t]} J_\tau^* &\geq J_t^* - \Delta_{[\text{alg}.s, t]}^J \\
 &\geq \frac{1}{2^m} \sum_{\tau \in [\text{alg}.s, t]} \left(r_\tau(s_\tau, a_\tau) + \sum_{s'} p_\tau(s'|s_\tau, a_\tau) h_t^*(s') - h_t^*(s_\tau) \right) - \Delta_{[\text{alg}.s, t]; \bar{D}} \\
 &\geq \frac{1}{2^m} \sum_{\tau \in [\text{alg}.s, t]} \left(R_\tau + h_t^*(s_{\tau+1}) - h_t^*(s_\tau) \right) - 2D_{\max} \sqrt{\frac{2 \log(SAT/\delta)}{2^m}} - \Delta_{[\text{alg}.s, t]; \bar{D}} \\
 &\hspace{15em} \text{(Azuma's inequality)} \\
 &\geq \frac{1}{2^m} \sum_{\tau \in [\text{alg}.s, t]} R_\tau - 4\bar{D} \sqrt{\frac{\log(SAT/\delta)}{2^m}} - \Delta_{[\text{alg}.s, t]; \bar{D}}. \quad (D_{\max} \leq \bar{D})
 \end{aligned}$$

So with high probability, **Test 1** will not return fail.

Furthremore, by [Lemma 12](#), with high probability,

$$\frac{1}{t - t_n + 1} \sum_{\tau=t_n}^t (\tilde{g}_\tau - R_\tau) \leq \hat{\rho}_{\text{UCRL}}(t - t_n + 1; \bar{D}) + \Delta_{[t_n, t]; \bar{D}} \leq 2\hat{\rho}_{\text{UCRL}}(t - t_n + 1; \bar{D})$$

where the last inequality is by the condition on $\Delta_{[t_0, t]; \bar{D}}$. Therefore, with high probability, **Test 2** will not return fail either.

It remains to show that the UCRL-ACW will not terminate and call for restart under the specified condition. By [Lemma 7](#), if \mathcal{P}_k^η contains an MDP whose diameter is upper bounded by \bar{D} , then the span of the output bias vector is upper bounded by $2\bar{D}$, and then the if-statement in [Line 4](#) of [Algorithm 4](#) will be triggered. Therefore, to show that UCRL-ACW will not terminate, we upper bound the η_k that needs to be added to \mathcal{P}_k in order to make at least one true MDP (whose diameter is upper bounded by $D_{\max} \leq \bar{D}$) lie in $\mathcal{P}_k^{\eta_k}$. Then we further argue that $\sum_{\tau=t_0}^t \eta_k(\tau)$ is not large enough to reach the condition in [Line 4](#) of [Algorithm 4](#).

For all episode k that starts before t , by Azuma's inequality,

$$\|\bar{p}_k(\cdot|s, a) - \hat{p}_k(\cdot|s, a)\|_1 \leq 2\sqrt{\frac{S \log(1/\delta)}{N_k^+(s, a)}}.$$

By the condition on $\Delta_{[t_0, t]; \bar{D}}$, we have

$$\|p_{t_0}(\cdot|s, a) - \bar{p}_k(\cdot|s, a)\|_1 \leq \frac{1}{\bar{D}} \Delta_{[t_0, t]; \bar{D}} \leq S \sqrt{\frac{A}{t - t_0 + 1}}.$$

Combining them, we get

$$\|p_{t_0}(\cdot|s, a) - \hat{p}_k(\cdot|s, a)\|_1 \leq 2 \sqrt{\frac{S \log(1/\delta)}{N_k^+(s, a)}} + S \sqrt{\frac{A}{t - t_0 + 1}}.$$

Therefore, we see that in [Line 4](#) of [Algorithm 4](#), as long as $\eta \geq S \sqrt{\frac{A}{t - t_0 + 1}}$, p_{t_0} is contained in \mathcal{P}_k^η . Then we have $\text{sp}(\tilde{h}) \leq 2D_{\max} \leq 2\bar{D}$ by [Lemma 7](#), and the for-loop will be broken at this η .

Thus we conclude that $\eta_k \leq 2S \sqrt{\frac{A}{t - t_0 + 1}}$ for all episode k started before t . Thus, $\sum_{\tau=t_0}^t \eta_k(\tau) \leq (t - t_0 + 1) \times 2S \sqrt{\frac{A}{t - t_0 + 1}} = 2S \sqrt{A(t - t_0 + 1)}$, and thus the algorithm will not terminate and call for restart at time t . \blacksquare

Lemma 26 (c.f. [Lemma 24](#)) *If $\bar{D} \geq D_{\max}$, then the number of epochs is upper bounded by $\min \left\{ L, 1 + 3 \left(\frac{\Delta^r + \Delta^p}{S\sqrt{A}} \right)^{\frac{2}{3}} T^{\frac{1}{3}} \right\}$.*

Proof Let $\mathcal{E}_1, \dots, \mathcal{E}_N$ be the epochs. By [Lemma 25](#), for $i \leq N - 1$, we must have $\Delta_{\mathcal{E}_i; \bar{D}} \geq \bar{D} S \sqrt{\frac{A}{|\mathcal{E}_i|}}$. By Hölder's inequality,

$$N - 1 \leq \left(\sum_{i=1}^{N-1} \frac{1}{\sqrt{|\mathcal{E}_i|}} \right)^{\frac{2}{3}} \left(\sum_{i=1}^{N-1} |\mathcal{E}_i| \right)^{\frac{1}{3}} \leq \left(\frac{\Delta_{[1, T]; \bar{D}}}{\bar{D} S \sqrt{A}} \right)^{\frac{2}{3}} T^{\frac{1}{3}}.$$

We can further upper bound the term $\frac{1}{\bar{D}} \Delta_{[1, T]; \bar{D}}$ as follows:

$$\begin{aligned} \frac{1}{\bar{D}} \Delta_{[1, T]; \bar{D}} &= \frac{1}{\bar{D}} (\Delta^r + 2\bar{D} \Delta^p + \Delta^J) \\ &\leq \frac{1}{\bar{D}} (2\Delta^r + (2\bar{D} + D_{\max}) \Delta^p) && \text{(by [Lemma 10](#))} \\ &\leq \frac{3}{\bar{D}} (\Delta^r + \bar{D} \Delta^p) \\ &\leq 3(\Delta^r + \Delta^p). \end{aligned}$$

Thus we get

$$N \leq 1 + 3 \left(\frac{\Delta^r + \Delta^p}{S\sqrt{A}} \right)^{\frac{2}{3}} T^{\frac{1}{3}}.$$

Also, by [Lemma 25](#), when $\bar{D} \geq D_{\max}$, an epoch is created only when the reward function or the transition function changes. Thus the number of epochs is also upper bounded by L . \blacksquare

Lemma 27 *In every epoch \mathcal{E} , the dynamic regret of MASTER-UCRL is upper bounded by*

$$\tilde{\mathcal{O}} \left(X + \overline{DS} \sqrt{A|\mathcal{E}|} + \overline{DS}^2 A^2 \right),$$

where X is the minimum of the following two terms:

$$\overline{DS} \sqrt{AL_{\mathcal{E}}|\mathcal{E}|} + \overline{DS} AL_{\mathcal{E}}$$

and

$$\left(\overline{D}^2 S^2 A \Delta_{\mathcal{E}; \overline{D}} |\mathcal{E}|^2 \right)^{1/3} + SA \sqrt{\overline{D} \Delta_{\mathcal{E}; \overline{D}} |\mathcal{E}|} + \left(\overline{D} SA^2 \Delta_{\mathcal{E}; \overline{D}}^2 |\mathcal{E}| \right)^{1/3} + SA \Delta_{\mathcal{E}; \overline{D}}.$$

Proof Let $\mathcal{E}_1, \dots, \mathcal{E}_N$ be the epochs. By [Lemma 22](#), we know that the regret within an epoch \mathcal{E} is $\tilde{\mathcal{O}} \left(\min \{ \text{Reg}_L(\mathcal{E}), \text{Reg}_{\Delta}(\mathcal{E}) \} + \frac{c_2 c_3}{c_1} |\mathcal{E}|^{1-p} + \frac{c_2^2}{c_3} \right)$ with

$$\text{Reg}_L(\mathcal{E}) = c_1 L_{\mathcal{E}}^{1-p} |\mathcal{E}|^p + c_2 L_{\mathcal{E}},$$

$$\text{Reg}_{\Delta}(\mathcal{E}) = \left(c_1 \Delta_{\mathcal{E}; \overline{D}}^{1-p} |\mathcal{E}| \right)^{\frac{1}{2-p}} + c_1 |\mathcal{E}|^p + c_1 (c_3^{-1} \Delta_{\mathcal{E}; \overline{D}})^{1-p} |\mathcal{E}|^p + c_2 \left(c_1^{-1} \Delta_{\mathcal{E}; \overline{D}} |\mathcal{E}|^{1-p} \right)^{\frac{1}{2-p}} + c_2 + c_2 c_3^{-1} \Delta_{\mathcal{E}; \overline{D}}$$

when $C(t)$ is in the form of [Definition 18](#). In our case $C_{\text{UCRL}}(t; \overline{D})$ is in this form with $c_1 = \overline{DS} \sqrt{A}$, $c_2 = \overline{D} SA$, $c_3 = \overline{D}$, and $p = \frac{1}{2}$. Using them in the bound above, we get that in an epoch, the dynamic regret is upper bounded by

$$\tilde{\mathcal{O}} \left(\min \{ \text{Reg}_L(\mathcal{E}), \text{Reg}_{\Delta}(\mathcal{E}) \} + \overline{D} \sqrt{A|\mathcal{E}|} + \overline{D} S^2 A^2 \right)$$

where

$$\text{Reg}_L(\mathcal{E}) = \overline{DS} \sqrt{AL_{\mathcal{E}}|\mathcal{E}|} + \overline{DS} AL_{\mathcal{E}}$$

$$\text{Reg}_{\Delta}(\mathcal{E}) = \left(\overline{D}^2 S^2 A \Delta_{\mathcal{E}; \overline{D}} |\mathcal{E}|^2 \right)^{1/3} + \overline{DS} \sqrt{A|\mathcal{E}|} + SA \sqrt{\overline{D} \Delta_{\mathcal{E}; \overline{D}} |\mathcal{E}|} + \left(\overline{D} SA^2 \Delta_{\mathcal{E}; \overline{D}}^2 |\mathcal{E}| \right)^{1/3} + \overline{D} SA + SA \Delta_{\mathcal{E}; \overline{D}}.$$

Collecting terms finishes the proof. ■

Theorem 28 *If $D_{\max} \leq \overline{D} \leq 2D_{\max}$, then MASTER-UCRL guarantees the following dynamic regret bound:*

$$\tilde{\mathcal{O}} \left(\min \left\{ D_{\max} S \sqrt{ALT}, D_{\max} (S^2 A)^{1/3} (\Delta^r + \Delta^r)^{1/3} T^{2/3} + D_{\max} S \sqrt{AT} \right\} \right).$$

Proof Let $\mathcal{E}_1, \dots, \mathcal{E}_N$ be the epochs. The per epoch dynamic regret is given by [Lemma 27](#). Combining them with Hölder's inequality and [Eq. \(20\)](#), [Eq. \(21\)](#), the dynamic regret in $[1, T]$ can be upper bounded by

$$\tilde{\mathcal{O}} \left(\min \{ \text{Reg}_L, \text{Reg}_{\Delta} \} + \overline{DS} \sqrt{ANT} + \overline{D} S^2 A^2 N \right) \tag{26}$$

where

$$\text{Reg}_L = \bar{D}S\sqrt{A(L+N)T} + \bar{D}SA(L+N) \quad (27)$$

and

$$\text{Reg}_\Delta = \left(\bar{D}^2 S^2 A \Delta_{[1,T];\bar{D}} T^2\right)^{1/3} + SA\sqrt{\bar{D}\Delta_{[1,T];\bar{D}}T} + \left(\bar{D}SA^2\Delta_{[1,T];\bar{D}}^2\right)^{1/3} + SA\Delta_{[1,T];\bar{D}}. \quad (28)$$

Since $\bar{D} \geq D_{\max}$, the number of epochs can be bounded using [Lemma 26](#):

$$N \leq \min \left\{ L, 1 + 3 \left(\frac{\Delta^r + \Delta^p}{S\sqrt{A}} \right)^{\frac{2}{3}} T^{\frac{1}{3}} \right\}.$$

With $N \leq L$, [Eq. \(26\)](#), and [Eq. \(27\)](#), the dynamic regret in $[1, T]$ can be bounded by (omitting lower order terms)

$$\tilde{\mathcal{O}} \left(\bar{D}S\sqrt{ALT} \right). \quad (29)$$

With $N \leq 1 + 3 \left(\frac{\Delta^r + \Delta^p}{S\sqrt{A}} \right)^{\frac{2}{3}} T^{\frac{1}{3}}$, [Eq. \(26\)](#), and [Eq. \(28\)](#), the regret can alternatively be upper bounded by (omitting lower order terms)

$$\tilde{\mathcal{O}} \left(\left(\bar{D}^2 S^2 A \Delta_{[1,T];\bar{D}} T^2 \right)^{1/3} + \bar{D}S\sqrt{AT} + \bar{D}(S^2 A)^{1/3}(\Delta^r + \Delta^p)^{1/3} T^{2/3} \right). \quad (30)$$

Then notice that $\bar{D} \leq 2D_{\max}$ and thus $\Delta_{[1,T];\bar{D}} = \Delta^r + 2\bar{D}\Delta^p + \Delta^J = \mathcal{O}(\Delta^r + D_{\max}\Delta^p)$ where we use [Lemma 10](#). Using these in [Eq. \(29\)](#) and [Eq. \(30\)](#) finishes the proof. \blacksquare

Theorem 29 *The doubling trick strategy described in [Section 4.2](#) for the unknown D_{\max} and known L case has a dynamic regret bound of $\tilde{\mathcal{O}} \left(D_{\max}S\sqrt{ALT} \right)$; for the unknown D_{\max} and known Δ case, the bound is*

$$\tilde{\mathcal{O}} \left(D_{\max}S\sqrt{AT} + D_{\max}(S^2 A)^{1/3}(\Delta^r + \Delta^p)^{1/3} T^{2/3} \right).$$

Proof For the known L case, when $\bar{D} \leq D_{\max}$, recall that the number of epochs is forced to be $N \leq L$. Similar to the proof of [Theorem 28](#), the regret in any of these epochs is upper bounded by

$$\tilde{\mathcal{O}} \left(\bar{D}S\sqrt{ANT} \right) = \tilde{\mathcal{O}} \left(\bar{D}S\sqrt{ALT} \right).$$

Summing the above over $\bar{D} = 1, 2, 4, \dots, D_{\max}$, we get $\tilde{\mathcal{O}} \left(D_{\max}S\sqrt{ALT} \right)$. When \bar{D} first enters $[D_{\max}, 2D_{\max}]$, we use [Theorem 28](#) to bound the regret in the rest of the rounds, which is still of order $\tilde{\mathcal{O}} \left(D_{\max}S\sqrt{ALT} \right)$.

For the case of known $\Delta = \Delta^r + \Delta^p$, the analysis is similar: when $\bar{D} \leq D_{\max}$, we force $N = 1 + 3(S^{-2}A^{-1}\Delta^2T)^{1/3}$, and thus the regret within any of these epochs is upper bounded by (similarly to the proof of [Theorem 28](#))

$$\tilde{\mathcal{O}} \left(\left(\bar{D}^2 S^2 A \Delta_{[1,T];\bar{D}} T^2 \right)^{1/3} + \bar{D} S \sqrt{AT} + \bar{D} (S^2 A)^{1/3} (\Delta^r + \Delta^p)^{1/3} T^{2/3} \right).$$

Summing this over $\bar{D} = 1, 2, \dots, D_{\max}$ and using $\Delta_{[1,T];\bar{D}} = \Delta^r + 2\bar{D}\Delta^p + \Delta^J = \mathcal{O}(\Delta^r + D_{\max}\Delta^p)$ for $\bar{D} = \mathcal{O}(D_{\max})$, we get $\tilde{\mathcal{O}} \left(D_{\max} (S^2 A \Delta T^2)^{1/3} + D_{\max} S \sqrt{AT} \right)$. When \bar{D} first enters $[D_{\max}, 2D_{\max}]$, we use [Theorem 28](#) to bound the regret in the rest of the rounds, which is still of the same order. \blacksquare

Appendix H. Bandit-over-Reinforcement-Learning Approach

The idea of the BoRL framework is to run a multi-armed bandit algorithm over a set of sub-algorithms each using a different parameter. In our case, each sub-algorithm is a MASTER-UCRL with a different guess on D_{\max} . The set of \bar{D} only needs to span the range of $[1, \sqrt{T}]$, since if $D_{\max} = \Omega(\sqrt{T})$, the regret bound would be vacuous.

We divide the horizon into $\frac{T}{B}$ equal-length intervals each of length $B = S\sqrt{AT}$. In each interval, sub-algorithm i restarts a MASTER-UCRL with $\bar{D} = 2^{i-1}$. The reward of sub-algorithm i in interval $b \in [\frac{T}{B}]$ is its total reward gained in the MDP for this interval. We denote i^* as the sub-algorithm that uses $\bar{D} \in [D_{\max}, 2D_{\max}]$.

On top of these sub-algorithms, we run the EXP3.P algorithm ([Auer et al., 2002b](#)). The ‘‘arms’’ are the sub-algorithms. From the above description, for this EXP3.P, there are $M = \lceil \log_2 \sqrt{T} \rceil$ arms, the algorithm proceeds for $\frac{T}{B}$ rounds, and in each round the reward range is B . By the standard regret bound of EXP3.P, the learner’s regret against sub-algorithm i^* is of order

$$\tilde{\mathcal{O}} \left(B \sqrt{M \frac{T}{B}} + BM \right) = \tilde{\mathcal{O}} \left(\sqrt{BT} \right).$$

with high probability.

On the other hand, in each interval $b \in [\frac{T}{B}]$, since sub-algorithm i^* uses a correct guess of \bar{D} , by [Theorem 28](#), its regret against the best sequence of policy in that interval is

$$\tilde{\mathcal{O}} \left(\min \left\{ D_{\max} S \sqrt{AL_b B}, D_{\max} (S^2 A)^{\frac{1}{3}} (\Delta_b)^{\frac{1}{3}} B^{\frac{2}{3}} + D_{\max} S \sqrt{AB} \right\} \right)$$

where we abuse notations and denote $L_b = L_{[(b-1)B+1, bB]}$, $\Delta_b = \Delta_{[(b-1)B+1, bB]}$.

Combining the two bounds above, we get that the regret of the learner against the best sequence of policies in $[1, T]$ is

$$\begin{aligned} & \tilde{\mathcal{O}} \left(\sqrt{BT} + \sum_{b=1}^{\frac{T}{B}} \min \left\{ D_{\max} S \sqrt{AL_b B}, D_{\max} (S^2 A)^{\frac{1}{3}} (\Delta_b)^{\frac{1}{3}} B^{\frac{2}{3}} + D_{\max} S \sqrt{AB} \right\} \right) \\ &= \tilde{\mathcal{O}} \left(\sqrt{BT} + \min \left\{ D_{\max} S \sqrt{A \left(L + \frac{T}{B} \right) T}, D_{\max} (S^2 A)^{\frac{1}{3}} (\Delta)^{\frac{1}{3}} T^{\frac{2}{3}} + D_{\max} S \sqrt{AB} \times \frac{T}{B} \right\} \right) \\ & \hspace{15em} \text{(using Eq. (20) and Eq. (21))} \end{aligned}$$

$$= \tilde{\mathcal{O}} \left(\sqrt{BT} + D_{\max} S \sqrt{\frac{A}{B}} T + \min \left\{ D_{\max} S \sqrt{ALT}, D_{\max} (S^2 A)^{\frac{1}{3}} \Delta^{\frac{1}{3}} T^{\frac{2}{3}} \right\} \right).$$

Using the B that we specified above, we get

$$\tilde{\mathcal{O}} \left(D_{\max} (S^2 A)^{1/4} T^{3/4} + \min \left\{ D_{\max} S \sqrt{ALT}, D_{\max} (S^2 A)^{\frac{1}{3}} \Delta^{\frac{1}{3}} T^{\frac{2}{3}} \right\} \right).$$

Appendix I. Verifying Assumption 1 for Several Algorithms

To prove Eq. (1), it suffices to prove the following.

Assumption 1' *There exist universal constants $c_1, c_2, c_3, c_4, c_5, c_6 > 0$ such that for all $t = 1, 2, \dots$, as long as $\Delta_{[1,t]} \leq c_1 \rho(t)$, the following holds with probability $1 - \frac{\delta}{T}$:*

$$\tilde{f}_t \geq \min_{\tau \in [1,t]} f_\tau^* - c_2 \Delta_{[1,t]} \quad (31)$$

$$\frac{1}{t} \sum_{\tau=1}^t (\tilde{f}_\tau - R_\tau) \leq c_3 \rho(t) + c_4 \Delta_{[1,t]}. \quad (32)$$

Furthermore, $\rho(t) \geq \frac{c_5}{\sqrt{t}}$, $\Delta(t) \geq c_6 \max_{\pi} |f_t(\pi) - f_{t+1}(\pi)|$.

This is because for an algorithm satisfying Assumption 1', we can redefine $\Delta(t) \leftarrow (c_3/c_1 + c_2 + c_4 + 1/(c_1 c_5) + 1/c_6) \Delta(t)$ and $\rho(t) \leftarrow (c_3 + c_1 c_2 + c_1 c_4 + 1/c_5 + c_1/c_6) \rho(t)$. Then Eq. (1) is satisfied. Our verification below is thus mostly based on Assumption 1' for simplicity.

The following proofs are brief (some of them are just sketches) since they follow standard analysis and mostly appear in previous works. Please find more details in the references. We sometimes make minor modifications to the original algorithm to make them more aligned with our framework.

I.1. UCB1 for Multi-armed Bandits

Algorithm 7: UCB1 for multi-armed bandits

input: A (number of arms), T, δ .

for $t = 1, \dots, T$ **do**

Choose $a_t = \operatorname{argmax}_{a \in [A]} \left(\hat{r}_{t,a} + c \sqrt{\frac{\log(T/\delta)}{N_{t,a}^+}} \right)$ $c > 0$ is some universal constant

where

$$\hat{r}_{t,a} = \frac{\sum_{\tau=1}^{t-1} R_\tau \mathbb{1}[a_\tau = a]}{N_{t,a}^+}, \quad N_{t,a}^+ = \max \left\{ 1, \sum_{\tau=1}^{t-1} \mathbb{1}[a_\tau = a] \right\}. \quad (33)$$

Receive R_t with $\mathbb{E}[R_t] = r_{t,a_t}$.

end

In this subsection, we consider the multi-armed bandit problem and the UCB1 algorithm by Auer et al. (2002a). Suppose there are A arms, and let $r_{t,a}$ denote the expected reward of arm a at

time t . Then the multi-armed bandit problem fits in our framework with $\Pi = [A]$ and $f_t(a) = r_{t,a}$. Below, we show that the UCB1 algorithm satisfies Assumption 1'.

The pseudocode of UCB1 is presented in [Algorithm 7](#). At time t , UCB1 chooses the arm that has the highest optimistic reward estimator $\tilde{r}_{t,a} \triangleq \hat{r}_{t,a} + c\sqrt{\frac{\log(T/\delta)}{N_{t,a}^+}}$, where $\hat{r}_{t,a}$ is the empirical mean of the reward of arm a up to time $t - 1$, $N_{t,a}$ is the cumulative number of pulls of arm a up to time $t - 1$ and $N_{t,a}^+ = \max\{1, N_{t,a}\}$, all defined in [Eq. \(33\)](#); $c > 0$ is some universal constant that is determined by Azuma's inequality.

To see that UCB1 satisfies Assumption 1', we define

$$\Delta(t) = \max_a |r_{t,a} - r_{t+1,a}|, \quad \tilde{f}_t = \max_a \tilde{r}_{t,a}, \quad \rho(t) = \sqrt{\frac{A \log(T/\delta)}{t}} + \frac{A \log(T/\delta)}{t}. \quad (34)$$

Furthermore, denote $\bar{r}_{t,a} = \frac{\sum_{\tau=1}^{t-1} r_{\tau,a} \mathbb{1}[a_t=a]}{N_{t,a}}$ (define $\bar{r}_{t,a} = 1$ if $N_{t,a} = 0$ for simplicity). Note that with high probability,

$$\tilde{f}_t \geq \max_a \bar{r}_{t,a} \geq \max_a \max_{\tau \leq t} r_{\tau,a} - \Delta_{[1,t]} \geq \min_{\tau \leq t} \max_a r_{\tau,a} - \Delta_{[1,t]},$$

where the first inequality is because with high probability, $\tilde{r}_{t,a} \geq \bar{r}_{t,a}$ by Azuma's inequality. This verifies [Eq. \(31\)](#).

On the other hand, by the selection rule $a_t = \operatorname{argmax}_a \tilde{r}_{t,a}$, we have with probability $1 - \delta$,

$$\begin{aligned} \sum_{\tau=1}^t (\tilde{f}_\tau - R_\tau) &\leq \sum_{\tau=1}^t (\tilde{r}_{\tau,a_\tau} - r_{\tau,a_\tau}) + \sum_{\tau=1}^t (r_{\tau,a_\tau} - R_\tau) \\ &= \sum_{\tau=1}^t \left(\bar{r}_{\tau,a_\tau} - r_{\tau,a_\tau} + c\sqrt{\frac{\log(T/\delta)}{N_{\tau,a}^+}} \right) + \sum_{\tau=1}^t (r_{\tau,a_\tau} - R_\tau) \\ &\leq t\Delta_{[1,t]} + \mathcal{O}\left(\sqrt{At \log(T/\delta)} + A \log(T/\delta)\right) \end{aligned}$$

where in the last inequality we use $\bar{r}_{\tau,a_\tau} - r_{\tau,a_\tau} \leq \Delta_{[1,t]}$ and the standard pigeonhole argument, and use Azuma's inequality to bound $\sum_{\tau=1}^t (r_{\tau,a_\tau} - R_\tau)$. This proves [Eq. \(32\)](#). Note that the condition $\Delta_{[1,t]} = \mathcal{O}(\rho(t))$ in Assumption 1' is even not needed.

I.2. OFUL for Linear Bandits

In this subsection, we consider linear bandits with a fixed action set, and the OFUL algorithm by [Abbasi-Yadkori et al. \(2011\)](#). The original OFUL algorithm handles the case where the action set can change over time (also known as the linear contextual bandit setting), but this is beyond the main focus of this paper. Let \mathcal{A} be the action set, and θ_t be the reward vector at time t . Then the linear bandit problem fits in our framework with $\Pi = \mathcal{A}$ and $f_t(a) = a^\top \theta_t$.

The pseudocode of OFUL (with a fixed action set) is presented in [Algorithm 8](#). For simplicity, assume that for all actions $a \in \mathcal{A}$, $\|a\|_2 \leq 1$, and for all t , the reward vector θ_t satisfies $\|\theta_t\|_2 \leq 1$. The OFUL algorithm chooses the action $a_t = \operatorname{argmax}_a a^\top \hat{\theta}_t + 2\beta \|a\|_{\Lambda_t^{-1}}$ at time t , where β , $\Lambda_t = I + \sum_{\tau=1}^{t-1} a_\tau a_\tau^\top$, and $\hat{\theta}_t$ are defined in [Eq. \(35\)](#).

Algorithm 8: OFUL for linear bandits**input:** $\mathcal{A} \subset \mathbb{R}^d$ (action set), T, δ .**for** $t = 1, 2, \dots, T$ **do** Choose $a_t = \operatorname{argmax}_{a \in \mathcal{A}} \left(a^\top \hat{\theta}_t + 2\beta \|a\|_{\Lambda_t^{-1}} \right)$,

where

$$\beta = 4\sqrt{d \log(T/\delta)}, \quad \Lambda_t = I + \sum_{\tau=1}^{t-1} a_\tau a_\tau^\top, \quad \hat{\theta}_t = \Lambda_t^{-1} \sum_{\tau=1}^{t-1} R_\tau a_\tau. \quad (35)$$

 Receive R_t with $\mathbb{E}[R_t] = a_t^\top \theta_t$.**end**

Then we define

$$\Delta(t) = d\sqrt{\log(T/\delta)} \|\theta_t - \theta_{t+1}\|_2, \quad \tilde{f}_t = \max_{a \in \mathcal{A}} \left(a^\top \hat{\theta}_t + 2\beta \|a\|_{\Lambda_t^{-1}} \right), \quad \rho(t) = \beta \sqrt{\frac{d \log(T/\delta)}{t}}. \quad (36)$$

Below, we verify that OFUL satisfies Assumption 1' with the choices in Eq. (36). Under the assumption that $\Delta_{[1,t]} \leq \rho(t)$, for any action a , by similar arguments as in (Zhao et al., 2020, Lemma 1),

$$\begin{aligned} \left| a^\top (\theta_t - \hat{\theta}_t) \right| &\leq \left| a^\top \Lambda_t^{-1} \sum_{\tau=1}^{t-1} a_\tau a_\tau^\top (\theta_s - \theta_t) \right| + \beta \|a\|_{\Lambda_t^{-1}} \\ &\leq \sum_{\tau=1}^{t-1} \left| a^\top \Lambda_t^{-1} a_\tau \right| \left| a_\tau^\top (\theta_\tau - \theta_t) \right| + \beta \|a\|_{\Lambda_t^{-1}} \\ &\leq \frac{\Delta_{[1,t]}}{d\sqrt{\log(T/\delta)}} \times \left(\sum_{\tau=1}^{t-1} \|a\|_{\Lambda_t^{-1}} \|a_\tau\|_{\Lambda_t^{-1}} \right) + \beta \|a\|_{\Lambda_t^{-1}} \\ &\quad \left(a_\tau^\top (\theta_\tau - \theta_t) \leq \|\theta_\tau - \theta_t\|_2 \leq \frac{\Delta_{[1,t]}}{d\sqrt{\log(T/\delta)}} \right) \\ &\leq \frac{\Delta_{[1,t]}}{d\sqrt{\log(T/\delta)}} \times \|a\|_{\Lambda_t^{-1}} \times \sqrt{(t-1) \sum_{\tau=1}^{t-1} \|a_\tau\|_{\Lambda_t^{-1}}^2} + \beta \|a\|_{\Lambda_t^{-1}} \\ &\quad \text{(Cauchy-Schwarz)} \\ &\leq \left(\beta + \Delta_{[1,t]} \sqrt{\frac{t}{d \log(T/\delta)}} \right) \|a\|_{\Lambda_t^{-1}} \\ &\quad \left(\sum_{\tau=1}^{t-1} \|a_\tau\|_{V_{t-1}^{-1}}^2 = \operatorname{tr}(\Lambda_t^{-1} \sum_{\tau=1}^{t-1} a_\tau a_\tau^\top) \leq d \right) \\ &\leq 2\beta \|a\|_{\Lambda_t^{-1}}. \quad \text{(by the assumption } \Delta_{[1,t]} \leq \rho(t)) \end{aligned} \quad (37)$$

Thus,

$$\begin{aligned}
 \sum_{\tau=1}^t (\tilde{f}_\tau - R_\tau) &= \sum_{\tau=1}^t (\tilde{f}_\tau - a_\tau^\top \theta_\tau) + \sum_{\tau=1}^t (a_\tau^\top \theta_\tau - R_\tau) \\
 &= \sum_{\tau=1}^t a_\tau^\top (\hat{\theta}_\tau - \theta_\tau) + 2 \sum_{\tau=1}^t \beta \|a_\tau\|_{\Lambda_\tau^{-1}} + \mathcal{O}\left(\sqrt{t \log(T/\delta)}\right) \\
 &\quad (\text{by the definition of } \tilde{f}_\tau \text{ and that OFUL chooses } a_\tau = \operatorname{argmax}_a (a^\top \hat{\theta}_\tau + 2\beta \|a\|_{\Lambda_\tau^{-1}})) \\
 &= \mathcal{O}\left(\sum_{s=1}^t \beta \|a_s\|_{\Lambda_s^{-1}}\right) + \mathcal{O}\left(\sqrt{t \log(T/\delta)}\right) \quad (\text{by Eq. (37)}) \\
 &= \mathcal{O}\left(\beta \sqrt{dt \log t}\right) = \mathcal{O}(t\rho(t)) \leq \mathcal{O}(t\rho(t) + t\Delta_{[1,t]}).
 \end{aligned}$$

This verifies Eq. (32). Also, by Eq. (37),

$$\tilde{f}_t = \max_a (a^\top \hat{\theta}_t + 2\beta \|a\|_{\Lambda_t^{-1}}) \geq \max_a a^\top \theta_t = f_t^* \geq \min_{\tau \in [1,t]} f_\tau^* - \Delta_{[1,t]}.$$

This verifies Eq. (31).

I.3. GLM-UCB for Generalized Linear Bandits

Algorithm 9: GLM-UCB for generalized linear bandits

input: $\mathcal{A} \subset \mathbb{R}^d, T, \delta, \mu$ (link function), λ .

define: $k_\mu = \sup_{x \in [0,1]} \frac{d\mu(x)}{dx}$, $c_\mu = \inf_{x \in [0,1]} \frac{d\mu(x)}{dx} > 0$.

for $t = 1, \dots, T$ **do**

Choose $a_t = \operatorname{argmax}_{a \in \mathcal{A}} (\mu(a^\top \hat{\theta}_t) + 2\beta \|a\|_{\Lambda_t^{-1}})$

where

$$\beta = \frac{4k_\mu}{c_\mu} \left(\sqrt{d \log(c_\mu T / (\lambda \delta))} + c_\mu \sqrt{\lambda} \right), \quad \Lambda_t = \lambda I + \sum_{\tau=1}^{t-1} a_\tau a_\tau^\top,$$

and $\hat{\theta}_t$ is the unique solution of the following set of equations (define

$$g_t(x) \triangleq \lambda c_\mu x + \sum_{\tau=1}^{t-1} \mu(a_\tau^\top x) a_\tau:$$

$$g_t(\theta_t) = \sum_{\tau=1}^{t-1} R_\tau a_\tau, \quad \hat{\theta}_t = \operatorname{argmin}_{\theta: \|\theta\|_2 \leq 1} \|g_t(\theta_t) - g_t(\theta)\|_{\Lambda_t^{-1}}.$$

Receive R_t with $\mathbb{E}[R_t] = \mu(a_t^\top \theta_t)$.

end

Generalized linear bandit is proposed by [Filippi et al. \(2010\)](#) and extended to the non-stationary case by [Cheung et al. \(2019\)](#); [Zhao et al. \(2020\)](#); [Russac et al. \(2020\)](#); [Faury et al. \(2021\)](#). We refer

the readers to these papers for the introduction of the setting. Again, we consider the special case where the action set is fixed over time, and for simplicity, we assume that the action set \mathcal{A} is a subset of $\{a \in \mathbb{R}^d : \|a\|_2 \leq 1\}$ and the hidden parameter θ_t satisfies $\|\theta_t\|_2 \leq 1$. The generalized linear bandit problem is accompanied with an increasing link function $\mu : \mathbb{R} \rightarrow \mathbb{R}$. It fits in our framework with $\Pi = \mathcal{A}$ and $f_t(a) = \mu(a^\top \theta_t)$.

The standard GLM-UCB is presented in [Algorithm 9](#). Below we show that GLM-UCB satisfies Assumption 1' with the following definitions:

$$\Delta(t) = \frac{k_\mu^2 d}{c_\mu} \sqrt{\log(T/\delta)} \|\theta_t - \theta_{t+1}\|_2, \quad \tilde{f}_t = \max_{a \in \mathcal{A}} \left(\mu(a^\top \hat{\theta}_t) + 2\beta \|a\|_{\Lambda_t^{-1}} \right), \quad \rho(t) = \beta \sqrt{\frac{d \log(T/\delta)}{t}},$$

where c_μ and k_μ are the infimum and supremum of the derivative of μ (defined in [Algorithm 9](#)). Define $G_t \triangleq \sum_{\tau=1}^{t-1} \left[\int_{v=0}^1 \mu \left(\langle a_\tau, (1-v)\hat{\theta}_t + v\theta_t \rangle dv \right) a_\tau a_\tau^\top + \lambda c_\tau I \right] \succeq c_\mu \Lambda_t$. Under the assumption that $\Delta(t) \leq \rho(t)$, for all $a \in \mathcal{A}$,

$$\begin{aligned} \left| \mu(a^\top \theta_t) - \mu(a^\top \hat{\theta}_t) \right| &\leq k_\mu \left| a^\top (\theta_t - \hat{\theta}_t) \right| \leq k_\mu \left| a^\top G_t^{-1} (g_t(\theta_t) - g_t(\hat{\theta}_t)) \right| \\ &\leq k_\mu \|a\|_{G_t^{-1}} \left\| g_t(\theta_t) - g_t(\hat{\theta}_t) \right\|_{G_t^{-1}} \\ &\leq \frac{k_\mu}{c_\mu} \|a\|_{\Lambda_t^{-1}} \left\| g_t(\theta_t) - g_t(\hat{\theta}_t) \right\|_{\Lambda_t^{-1}} \\ &\leq \frac{k_\mu}{c_\mu} \|a\|_{\Lambda_t^{-1}} \left\| g_t(\theta_t) - g_t(\theta'_t) \right\|_{\Lambda_t^{-1}} \\ &= \frac{k_\mu}{c_\mu} \|a\|_{\Lambda_t^{-1}} \left\| \sum_{\tau=1}^{t-1} \left(\mu(a_\tau^\top \theta_t) - \mu(a_\tau^\top \theta_\tau) \right) a_\tau + \sum_{\tau=1}^{t-1} \left(\mu(a_\tau^\top \theta_\tau) - R_\tau \right) a_\tau + \lambda c_\mu \theta_t \right\|_{\Lambda_t^{-1}} \\ &\leq \frac{k_\mu}{c_\mu} \|a\|_{\Lambda_t^{-1}} \left(k_\mu \max_{\tau \leq t} \|\theta_t - \theta_\tau\|_2 \sum_{\tau=1}^{t-1} \|a_\tau\|_{\Lambda_t^{-1}} + \left\| \sum_{\tau=1}^{t-1} \eta_\tau a_\tau \right\|_{\Lambda_t^{-1}} + \sqrt{\lambda} c_\mu \right) \\ &\hspace{15em} (\text{define } \eta_\tau = \mu(a_\tau^\top \theta_\tau) - R_\tau) \\ &\leq \frac{k_\mu}{c_\mu} \|a\|_{\Lambda_t^{-1}} \left(k_\mu \frac{c_\mu \Delta_{[1,t]}}{k_\mu^2 d \sqrt{\log(T/\delta)}} \sqrt{dt} + \sqrt{d \log(c_\mu T/\delta)} + \sqrt{\lambda} c_\mu \right) \\ &\leq \frac{k_\mu}{c_\mu} \|a\|_{\Lambda_t^{-1}} \left(\frac{c_\mu \rho(t)}{k_\mu} \sqrt{\frac{t}{d \log(T/\delta)}} + \sqrt{d \log(c_\mu T/\delta)} + \sqrt{\lambda} c_\mu \right) \\ &\hspace{15em} (\text{by the assumption } \Delta_{[1,t]} \leq \rho(t)) \\ &\leq 2\beta \|a\|_{\Lambda_t^{-1}}. \end{aligned} \tag{38}$$

Thus,

$$\begin{aligned} \sum_{\tau=1}^t (\tilde{f}_\tau - R_\tau) &= \sum_{\tau=1}^t (\tilde{f}_\tau - \mu(a_\tau^\top \theta_\tau)) + \sum_{\tau=1}^t (\mu(a_\tau^\top \theta_\tau) - R_\tau) \\ &\leq \sum_{\tau=1}^t (\mu(a_\tau^\top \hat{\theta}_\tau) - \mu(a_\tau^\top \theta_\tau)) + 2\beta \sum_{\tau=1}^t \|a_\tau\|_{\Lambda_t^{-1}} + \mathcal{O} \left(\sqrt{t \log(T/\delta)} \right) \end{aligned}$$

$$\begin{aligned}
 &\leq \mathcal{O} \left(\beta \sum_{\tau=1}^t \|a_\tau\|_{\Lambda_t^{-1}} \right) = \mathcal{O} \left(\beta \sqrt{dt \log(T/\delta)} \right) \\
 &= \mathcal{O}(t\rho(t)) = \mathcal{O}(t\rho(t) + t\Delta_{[1,t]}).
 \end{aligned}$$

This verifies Eq. (32). Furthermore, by Eq. (38),

$$\tilde{f}_t = \max_{a \in \mathcal{A}} \left(\mu(a^\top \hat{\theta}_t) + 2\beta \|a\|_{\Lambda_t^{-1}} \right) \geq \max_{a \in \mathcal{A}} \mu(a^\top \theta_t) = f_t^* \geq \min_{\tau \in [1,t]} f_\tau^* - \Delta_{[1,t]}.$$

This verifies Eq. (31).

I.4. Q-UCB for Finite-horizon Tabular MDPs

Algorithm 10: Q-UCB for finite-horizon tabular MDPs

input: S (number of states), A (number of actions), H, T, δ .

$Q_h(s, a) \leftarrow H$, $N_h(s, a) \leftarrow 0$ for all h, s, a .

for $t = 1, \dots, T$ **do**

for $h = 1, \dots, H$ **do**

 Choose $a_h^t \leftarrow \operatorname{argmax}_a Q_h(s_h^t, a)$.

$\tau = N_h(s_h^t, a_h^t) \leftarrow N_h(s_h^t, a_h^t) + 1$, $b_\tau \leftarrow c\sqrt{H^3 \log(SAT/\delta)/\tau}$.

$\triangleleft c$ is a universal constant

$Q_h(s_h^t, a_h^t) \leftarrow (1 - \alpha_\tau)Q_h(s_h^t, a_h^t) + \alpha_\tau [r_h^t(s_h^t, a_h^t) + V_{h+1}(s_{h+1}^t) + b_\tau]$. $\triangleleft \alpha_\tau \triangleq \frac{H+1}{H+\tau}$

$V_h(s_h^t) \leftarrow \min \{H, \max_a Q_h(s_h^t, a)\}$.

end

end

The finite-horizon tabular MDP problem fits in our framework with Π being the set of deterministic policies on the MDP, and $f_t(\pi)$ being the expected reward of policy π in episode t . Q-UCB (Hoeffding-style) is a model-free algorithm for finite-horizon tabular MDPs proposed by Jin et al. (2018), whose pseudocode is in Algorithm 10. Let H denote the horizon length, s_h^t, a_h^t denote the state and actions visited at step h of episode t , and r_h^t, p_h^t denote the reward and transition functions at step h of episode t . Without loss of generality, we assume that $s_1^t = s_1$ for all t (i.e., the initial state is fixed).

It has been shown in the proof of (Mao et al., 2021, Theorem 1) that Q-UCB satisfy Assumption 1' with the following choices:

$$\Delta(t) = H \sum_{h=1}^H \max_{s,a} |r_h^t(s, a) - r_h^{t+1}(s, a)| + H^2 \sum_{h=1}^H \max_{s,a} \|p_h^t(\cdot|s, a) - p_h^{t+1}(\cdot|s, a)\|_1,$$

$$\tilde{f}_t = V_h^t(s_1), \quad (V_h^t \text{ is the } V_h \text{ in Algorithm 10 at the beginning of episode } t)$$

$$\rho(t) = \tilde{\mathcal{O}} \left(\sqrt{\frac{H^5 SA}{t}} + \frac{H^3 SA}{t} \right).$$

The proof details are omitted here.

Algorithm 11: LSVI-UCB for finite-horizon linear MDP

input: \mathcal{S} (state space), \mathcal{A} (action space), $\phi(\cdot, \cdot) : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$, H, T, δ .

for $t = 1, \dots, T$ **do**

for $h = H, \dots, 1$ **do**

$\Lambda_h \leftarrow \sum_{\tau=1}^{t-1} \phi(s_h^\tau, a_h^\tau) \phi(s_h^\tau, a_h^\tau)^\top + I$.

$w_h \leftarrow \Lambda_h^{-1} \sum_{\tau=1}^{t-1} \phi(s_h^\tau, a_h^\tau) [r_h^\tau(x_h^\tau, a_h^\tau) + \max_{a \in \mathcal{A}} Q_{h+1}(x_{h+1}^\tau, a)]$.

$Q_h(\cdot, \cdot) \leftarrow \min \left\{ w_h^\top \phi(\cdot, \cdot) + 2\beta \left(\phi(\cdot, \cdot) \Lambda_h^{-1} \phi(\cdot, \cdot) \right)^{1/2}, H \right\}$.

$\triangleleft \beta = cdH\sqrt{\log(T/\delta)}$ for some universal constant c

end

for $h = 1, \dots, H$ **do**

 Take action $a_h^t \leftarrow \operatorname{argmax}_{a \in \mathcal{A}} Q_h(s_h^t, a)$.

end

end

I.5. LSVI-UCB for Finite-horizon Linear MDPs

See (Zhou et al., 2020; Touati and Vincent, 2020) for the non-stationary finite-horizon linear MDP setting. We assume that the reward function and the transition function at step h of episode t are $r_h^t(s, a) = \phi(s, a)^\top \theta_h^t$ and $p_h^t(s'|s, a) = \phi(s, a)^\top \mu_h^t(s')$ where $\phi(\cdot, \cdot)$ is the feature function that maps a state-action pair to a d -dimensional feature vector. The problem fits in our framework with Π being the set of deterministic policies, and $f_t(\pi)$ being the expected reward of policy π in episode t . The LSVI-UCB algorithm is an optimism-based algorithm proposed by Jin et al. (2020b), whose pseudocode is shown in Algorithm 11. We define $Q_h^t, w_h^t, \Lambda_h^t$ to be the Q_h, w_h, Λ_h at Line 11 of round t . Furthermore, define $V_h^t(s) = \max_{a \in \mathcal{A}} Q_h^t(s, a)$. Again, without loss of generality, we assume $s_1^t = s_1$ (the initial state is fixed).

We define

$$\Delta(t) = dH\sqrt{\log(T/\delta)} \left(\sum_{h=1}^H \|\theta_h^t - \theta_h^{t+1}\|_2 + H \sum_{h=1}^H \|\mu_h^t - \mu_h^{t+1}\|_F \right),$$

$$\tilde{f}_t = V_1^t(s_1),$$

$$\rho(t) = c\sqrt{\frac{d^3 H^4}{t}} \log(T/\delta) = \beta H \sqrt{\frac{d \log(T/\delta)}{t}}. \quad (c \text{ and } \beta \text{ defined in Algorithm 11})$$

Below, we verify that LSVI-UCB satisfies Assumption 1' with the Δ, \tilde{f} , and ρ defined above. Assume that $\Delta_{[1,t]} \leq \rho(t)$. By similar arguments as in the proof of (Zhou et al., 2020, Lemma 3), we have

$$\left| \phi(s, a)^\top w_h^t - Q_h^*(s, a) - \mathbb{P}_h^t(V_h^t - V_h^*)(s, a) \right| \leq \left(\beta + \sqrt{dt} B_{\theta, [1,t]} + H\sqrt{dt} B_{\mu, [1,t]} \right) \|\phi(s, a)\|_{(\Lambda_h^t)^{-1}} \quad (39)$$

where $B_{\theta, [1,t]} = \sum_{\tau=1}^{t-1} \sum_{h=1}^H \|\theta_h^\tau - \theta_h^{\tau+1}\|_2$ and $B_{\mu, [1,t]} = \sum_{\tau=1}^{t-1} \sum_{h=1}^H \|\mu_h^\tau - \mu_h^{\tau+1}\|_F$. By the definition of $\Delta(t)$, the right-hand side of Eq. (39) can be further upper bound by

$$\left(\beta + \frac{1}{H} \sqrt{\frac{t}{d \log(T/\delta)}} \Delta_{[1,t]} \right) \|\phi(s, a)\|_{(\Lambda_h^t)^{-1}} \leq 2\beta \|\phi(s, a)\|_{(\Lambda_h^t)^{-1}}, \quad (40)$$

where the inequality is by the assumption that $\Delta_{[1,t]} \leq \rho(t)$. Similar to the proof of (Zhou et al., 2020, Lemma 4), we can then show that for any t, h ,

$$\begin{aligned} Q_h^t(s, a) - Q_h^*(s, a) &= \phi(s, a)^\top w_h^t - Q_h^*(s, a) + 2\beta \|\phi(s, a)\|_{(\Lambda_h^t)^{-1}} \\ &\geq \max_{s'} (V_{h+1}^t(s') - V_{h+1}^*(s')) \quad (\text{by Eq. (39) and Eq. (40)}) \end{aligned}$$

and further using induction to show that $V_1^t(s) \geq V_1^*(s)$. Thus, $\tilde{f}_t = V_1^t(s_1) \geq V_1^*(s_1)$, which verifies Eq. (31). One can also show that $\sum_{\tau=1}^t (\tilde{f}_\tau - R_\tau) = \mathcal{O}(t\rho(t))$ using the standard analysis of LSVI-UCB (e.g., (Jin et al., 2020b, Theorem 3.1), (Zhou et al., 2020, Theorem 5)). This verifies Eq. (32).

I.6. ILOVETOCONBANDITS for Contextual Bandits

Algorithm 12: ILOVETOCONBANDITS for contextual bandits

input: Π (policy set), \mathcal{A} (action set), T, δ .

for $t = 1, \dots, T$ **do**

 Calculate $Q_t \in \Delta_\Pi$ that satisfies the following constraints with some universal constant $c' > 0$:

$$\begin{aligned} \sum_{\pi} Q(\pi) \widehat{\text{Reg}}_{[1,t-1]}(\pi) &\leq 2c' A \mu_t \\ \forall \pi \in \Pi, \quad \frac{1}{t-1} \sum_{\tau=1}^{t-1} \frac{1}{Q^{\mu_t}(\pi(x_\tau)|x_\tau)} &\leq 2A + \frac{\widehat{\text{Reg}}_{[1,t-1]}(\pi)}{c' \mu_t} \end{aligned}$$

where $\mu_t \triangleq \sqrt{\frac{\log(|\Pi|T/\delta)}{At}}$, $Q^\mu(a|x) \triangleq (1 - A\mu) \sum_{\pi \in \Pi} Q(\pi) \mathbb{1}[\pi(x) = a] + \mu$, and

$$\widehat{\text{Reg}}_{\mathcal{I}}(\pi) \triangleq \frac{1}{|\mathcal{I}|} \max_{\pi'} \sum_{\tau \in \mathcal{I}} (\hat{r}_\tau(\pi'(x_\tau)) - \hat{r}_\tau(\pi(x_\tau))), \quad \hat{r}_\tau(a) \triangleq \frac{R_\tau \mathbb{1}[a_\tau = a]}{p_\tau(a)}.$$

 Let $p_t(\cdot) = Q^{\mu_t}(\cdot|x_t)$ and sample $a_t \sim p_t$.

end

In the contextual bandit problem, in each round, the learner first sees a context $x_t \in \mathcal{X}$, and then chooses an action $a_t \in [A]$ based on it. The learner then receives the reward $r_t(a_t) \in \mathbb{R}$. We assume that (x_t, r_t) is sampled from the distribution \mathcal{D}_t . The goal of the learner is to be comparable to the best mapping $\pi : \mathcal{X} \rightarrow [A]$ within a given set of mappings Π (which are called *policies*), i.e., the learner wants to minimize $\sum_t (r_t(\pi_t^*(x_t)) - r_t(a_t))$ where $\pi_t^* \triangleq \max_{\pi' \in \Pi} \mathbb{E}_{(x,r) \sim \mathcal{D}_t} [r(\pi'(x))]$. See (Agarwal et al., 2014) for more detailed description of the problem. This problem fits in our framework with the same Π and $f_t(\pi) = \mathbb{E}_{(x,r) \sim \mathcal{D}_t} [r(\pi(x))]$.

The algorithm ILOVETOCONBANDITS (Algorithm 12) by Agarwal et al. (2014) achieves the optimal regret bound in the i.i.d. case. The analysis for ILOVETOCONBANDITS is more involved. Fortunately, Chen et al. (2019) already has helpful lemmas for ILOVETOCONBANDITS in the non-stationary case, and we can simply reuse them. We show a more general result that Assumption 1 is satisfied no matter how large $\Delta_{[1,t]}$ is.

Let $\mathcal{R}_{\mathcal{I}}(\pi) = \frac{1}{|\mathcal{I}|} \sum_{\tau \in \mathcal{I}} \mathbb{E}_{(x,r) \sim \mathcal{D}_{\tau}} [r(\pi(x))]$ be the expected of policy π in the interval \mathcal{I} , $\widehat{\mathcal{R}}_{\mathcal{I}}(\pi) = \frac{1}{|\mathcal{I}|} \sum_{\tau \in \mathcal{I}} \widehat{r}_{\tau}(\pi(x_{\tau}))$ be an unbiased estimator of $\mathcal{R}_{\mathcal{I}}(\pi)$, with \widehat{r}_{τ} an unbiased estimator for the action reward constructed with inverse propensity weighting at time τ . Let $\text{Reg}_{\mathcal{I}}(\pi) = \max_{\pi'} \mathcal{R}_{\mathcal{I}}(\pi') - \mathcal{R}_{\mathcal{I}}(\pi)$ and $\widehat{\text{Reg}}_{\mathcal{I}}(\pi) = \max_{\pi'} \widehat{\mathcal{R}}_{\mathcal{I}}(\pi') - \widehat{\mathcal{R}}_{\mathcal{I}}(\pi)$. Below, we will show that ILOVE-TOCONBANDITS satisfies Assumption 1' with the following definitions:

$$\begin{aligned} \Delta(t) &\triangleq \|\mathcal{D}_t - \mathcal{D}_{t+1}\|_{\text{TV}} = \int_{\mathcal{r}} \int_{\mathcal{x}} |\mathcal{D}_t(x, r) - \mathcal{D}_{t+1}(x, r)| dx dr, \\ \widetilde{f}_t &\triangleq \max_{\pi} \widehat{\mathcal{R}}_{[1, t-1]}(\pi) + c_2 A \mu_{t-1} \quad (\text{for some universal constant } c_2 > 0) \\ \rho(t) &\triangleq \sqrt{\frac{A \log(|\Pi|T/\delta)}{t}}. \end{aligned}$$

Note that $\Delta(t)$ upper bounds $|\mathbb{E}_{(x,r) \sim \mathcal{D}_t} [r(\pi(x))] - \mathbb{E}_{(x,r) \sim \mathcal{D}_{t+1}} [r(\pi(x))]|$.

Combining the proofs of Lemma 14 and Lemma 16 in (Chen et al., 2019), we get the following guarantee with probability at least $1 - \delta$ for any policy π :

$$\left| \widehat{\mathcal{R}}_{[1, t]}(\pi) - \mathcal{R}_{[1, t]}(\pi) \right| \leq c_1 \widehat{\text{Reg}}_{[1, t]}(\pi) + c_2 A \mu_t + c_3 \Delta_{[1, t]}, \quad (41)$$

where $\mu_t = \sqrt{\frac{\log(|\Pi|T/\delta)}{At}}$ and c_1, c_2, c_3 are universal constants. To see how to get Eq. (41), notice that Lemma 14 of (Chen et al., 2019) gives $\left| \widehat{\mathcal{R}}_{[1, t]}(\pi) - \mathcal{R}_{[1, t]}(\pi) \right| \leq \mathcal{O}\left(\frac{\mu_t}{t} \sum_{\tau=1}^t U_{\tau} + \frac{\log(|\Pi|T/\delta)}{t\mu_t}\right)$, and they further upper bound U_{τ} by $\mathcal{O}\left(\frac{\text{Reg}_{[1, t]}}{\mu_t} + A + \frac{\Delta_{[1, t]}}{\mu_t}\right)$ in the second-to-last line in their proof of Lemma 16. Combining them yields Eq. (41). Notice that they have an additional $\log T$ factor which we do not suffer.

Below, let $\bar{\pi}_t = \arg\max_{\pi} \mathcal{R}_{[1, t]}(\pi)$. Then we have

$$\begin{aligned} \max_{\pi} \widehat{\mathcal{R}}_{[1, t]}(\pi) &\geq \widehat{\mathcal{R}}_{[1, t]}(\bar{\pi}_t) \geq \mathcal{R}_{[1, t]}(\bar{\pi}_t) - c_3 \Delta_{[1, t]} - c_2 A \mu_t \\ &= \max_{\pi} \mathcal{R}_{[1, t]}(\pi) - c_3 \Delta_{[1, t]} - c_2 A \mu_t, \end{aligned} \quad (42)$$

where in the second inequality we use Eq. (41) with the fact that $\text{Reg}_{[1, t]}(\bar{\pi}_t) = 0$. Therefore, if we choose $\widetilde{f}_t = \max_{\pi} \widehat{\mathcal{R}}_{[1, t-1]}(\pi) + c_2 A \mu_{t-1}$, then

$$\begin{aligned} \widetilde{f}_t &\geq \max_{\pi} \mathcal{R}_{[1, t-1]}(\pi) - c_3 \Delta_{[1, t-1]} \quad (\text{using Eq. (42) and the definition of } \widetilde{f}_t) \\ &\geq \max_{\pi} \max_{\tau \in [1, t]} \mathcal{R}_{\tau}(\pi) - (c_3 + 1) \Delta_{[1, t]} \end{aligned} \quad (43)$$

which proves Eq. (31). Next, we show Eq. (32):

$$\begin{aligned} \widetilde{f}_t - \mathbb{E}_t[R_t] &\leq \sum_{\pi} Q_t(\pi) \left(\widetilde{f}_t - \mathcal{R}_t(\pi) \right) + \mathcal{O}(A \mu_t) \\ &\quad (\text{by the algorithm, which uses } \mathcal{O}(A \mu_t) \text{ probability to explore actions}) \\ &= \sum_{\pi} Q_t(\pi) \left(\max_{\pi'} \widehat{\mathcal{R}}_{[1, t-1]}(\pi') - \mathcal{R}_t(\pi) \right) + \mathcal{O}(A \mu_t) \end{aligned}$$

$$\begin{aligned}
 &\leq \sum_{\pi} Q_t(\pi) \left(\max_{\pi'} \widehat{\mathcal{R}}_{[1,t-1]}(\pi') - \mathcal{R}_{[1,t-1]}(\pi) \right) + \mathcal{O}(A\mu_t + \Delta_{[1,t]}) \\
 &= \sum_{\pi} Q_t(\pi) \left(\widehat{\text{Reg}}_{[1,t-1]}(\pi) + \widehat{\mathcal{R}}_{[1,t-1]}(\pi) - \mathcal{R}_{[1,t-1]}(\pi) \right) + \mathcal{O}(A\mu_t + \Delta_{[1,t]}) \\
 &\leq \sum_{\pi} Q_t(\pi) \left(\widehat{\text{Reg}}_{[1,t-1]}(\pi) + c_1 \text{Reg}_{[1,t-1]}(\pi) \right) + \mathcal{O}(A\mu_t + \Delta_{[1,t]}) \quad (\text{by Eq. (41)}) \\
 &\leq (1 + 2c_1) \sum_{\pi} Q_t(\pi) \widehat{\text{Reg}}_{[1,t-1]}(\pi) + \mathcal{O}(A\mu_t + \Delta_{[1,t]}) \tag{44}
 \end{aligned}$$

where the last inequality is by Lemma 16 of (Chen et al., 2019), which bounds $\text{Reg}_{[1,t-1]}(\pi)$ by $2\widehat{\text{Reg}}_{[1,t-1]}(\pi) + \mathcal{O}(A\mu_t + \Delta_{[1,t]})$. By the algorithm, $\sum_{\pi} Q_t(\pi) \widehat{\text{Reg}}_{[1,t-1]}(\pi)$ is of order $\mathcal{O}(A\mu_t)$. Therefore, the last expression can further be upper bounded by $\mathcal{O}(A\mu_t + \Delta_{[1,t]})$. Finally, with the above calculation and Azuma's inequality, we get

$$\sum_{\tau=1}^t (\tilde{f}_{\tau} - R_{\tau}) \leq \sum_{\tau=1}^t (\tilde{f}_{\tau} - \mathbb{E}_{\tau}[R_{\tau}]) + \sum_{\tau=1}^t (\mathbb{E}_{\tau}[R_{\tau}] - R_{\tau}) \leq \mathcal{O}\left(\sqrt{At \log(|\Pi|T/\delta)} + \Delta_{[1,t]}\right)$$

Since we choose $\rho(t) = \sqrt{\frac{A \log(|\Pi|T/\delta)}{t}}$, Eq. (32) is also satisfied.

I.7. FALCON for Contextual Bandits

Algorithm 13: FALCON for realizable contextual bandits

input: Φ (reward function class), \mathcal{A} (action sets), T, δ .

for $t = 1, \dots, T$ **do**

Let $\gamma_t = \sqrt{At / \log(|\Phi|T/\delta)}$.

Compute $\hat{\phi}_t = \operatorname{argmin}_{\phi \in \Phi} \sum_{\tau=1}^{t-1} (\phi(x_{\tau}, a_{\tau}) - r_{\tau}(a_{\tau}))^2$

Observe context x_t .

Let $\hat{a}_t = \operatorname{argmax}_{a \in \mathcal{A}} \hat{\phi}_t(x_t, a)$. Define

$$p_t(a) \triangleq \begin{cases} \frac{1}{A + \gamma_t (\hat{\phi}_t(x_t, \hat{a}_t) - \hat{\phi}_t(x_t, a))}, & \text{for } a \neq \hat{a}_t, \\ 1 - \sum_{a' \neq \hat{a}_t} p_t(a'), & \text{for } a = \hat{a}_t. \end{cases}$$

Sample $a_t \sim p_t$ and observe reward R_t .

end

FALCON is an algorithm for stationary contextual bandits. It relies on the assumption that the expected reward of action a under context x is given by an unknown function $\phi^*(x, a) : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]$. The learner is given the function class Φ that contains ϕ^* . For each $\phi \in \Phi$, one can derive a policy $\pi_{\phi} : \mathcal{X} \rightarrow \mathcal{A}$ such that $\pi_{\phi}(x) = \operatorname{argmax}_{a \in \mathcal{A}} \phi(x, a)$. It is straightforward to see that the optimal policy is π_{ϕ^*} , and the learner's goal is to be competitive with it. This problem falls into our framework with $\Pi = \{\pi_{\phi} : \phi \in \Phi\}$ and $f_t(\pi) = \mathbb{E}_{x \in \mathcal{D}_t} [\phi^*(x, \pi(x))]$ where \mathcal{D}_t is the distribution of context at time t . The algorithm FALCON is shown in Algorithm 13.

Below, we show that it also satisfies Assumption 1'.

At time t , the context x_t is sampled from \mathcal{D}_t , and the reward is generated by $\mathbb{E}[r_t(x_t, a_t)] = \phi_t^*(x_t, a_t)$. We slightly modify their algorithm so that at every round t , the algorithm call the regression oracle once and obtain $\hat{\phi}_t = \operatorname{argmin}_{\phi \in \Phi} \sum_{\tau=1}^{t-1} (\phi(x_\tau, a_\tau) - R_\tau)^2$ (the original algorithm does this only when the time index doubles), and then construct a mapping from context to action distribution $p_t(\cdot|\cdot)$ as specified in their algorithm. Analogous to their definitions, we define

$$\begin{aligned} \mathcal{R}_{[1,t-1]}(\pi) &= \frac{1}{t-1} \sum_{\tau=1}^{t-1} \mathbb{E}_{x \sim \mathcal{D}_\tau} [\phi_\tau^*(x, \pi(x))], \\ \widehat{\mathcal{R}}_{[1,t-1]}(\pi) &= \frac{1}{t-1} \sum_{\tau=1}^{t-1} \mathbb{E}_{x \sim \mathcal{D}_\tau} [\widehat{\phi}_t(x, \pi(x))], \\ \widehat{\operatorname{Reg}}_{[1,t-1]}(\pi) &= \widehat{\mathcal{R}}_{[1,t-1]}(\pi_{\widehat{\phi}_t}) - \widehat{\mathcal{R}}_{[1,t-1]}(\pi), \\ \operatorname{Reg}_{[1,t-1]}(\pi) &= \max_{\phi \in \Phi} \mathcal{R}_{[1,t-1]}(\pi_\phi) - \mathcal{R}_{[1,t-1]}(\pi), \\ V_t(p, \pi) &= \mathbb{E}_{x \sim \mathcal{D}_t} \left[\frac{1}{p(\pi(x)|x)} \right], \\ \mathcal{V}_t(\pi) &= \max_{\tau \in [1,t]} V_\tau(p_\tau, \pi). \end{aligned}$$

We will show that FALCON satisfies Assumption 1' with the following definitions:

$$\begin{aligned} \Delta(t) &= \sqrt{A} \max_{x,a} |\phi_t^*(x, a) - \phi_{t+1}^*(x, a)| + \int_x |\mathcal{D}_t(x) - \mathcal{D}_{t+1}(x)| dx, \\ \rho(t) &= \sqrt{\frac{A \log(|\Phi|T/\delta)}{t}} \end{aligned}$$

By the same calculation as in Lemma 7 of (Simchi-Levi and Xu, 2020), for any π ,

$$\begin{aligned} &(t-1) \left| \widehat{\mathcal{R}}_{[1,t-1]}(\pi) - \mathcal{R}_{[1,t-1]}(\pi) \right|^2 \\ &= \frac{1}{t-1} \left(\sum_{\tau=1}^{t-1} \mathbb{E}_{x \sim \mathcal{D}_\tau} [\widehat{\phi}_t(x, \pi(x)) - \phi_\tau^*(x, \pi(x))] \right)^2 \\ &\leq \sum_{\tau=1}^{t-1} \left(\mathbb{E}_{x \sim \mathcal{D}_\tau} [\widehat{\phi}_t(x, \pi(x)) - \phi_\tau^*(x, \pi(x))] \right)^2 \\ &\leq \sum_{\tau=1}^{t-1} \left(\mathbb{E}_{x \sim \mathcal{D}_\tau} \left[\sqrt{\frac{1}{p_\tau(\pi(x)|x)} p_\tau(\pi(x)|x)} \left(\widehat{\phi}_t(x, \pi(x)) - \phi_\tau^*(x, \pi(x)) \right)^2 \right] \right)^2 \\ &\leq \sum_{\tau=1}^{t-1} \left(\mathbb{E}_{x \sim \mathcal{D}_\tau} \left[\sqrt{\frac{1}{p_\tau(\pi(x)|x)} \mathbb{E}_{a \sim p_\tau(\cdot|x)} \left(\widehat{\phi}_t(x, a) - \phi_\tau^*(x, a) \right)^2} \right] \right)^2 \\ &\leq \sum_{\tau=1}^{t-1} \mathbb{E}_{x \sim \mathcal{D}_\tau} \left[\frac{1}{p_\tau(\pi(x)|x)} \right] \mathbb{E}_{x \sim \mathcal{D}_\tau} \mathbb{E}_{a \sim p_\tau(\cdot|x)} \left[\left(\widehat{\phi}_t(x, a) - \phi_\tau^*(x, a) \right)^2 \right] \\ &\leq \sum_{\tau=1}^{t-1} V_\tau(p_\tau, \pi) \mathbb{E}_{a \sim p_\tau(\cdot|x)} \left[\left(\widehat{\phi}_t(x, a) - \phi_\tau^*(x, a) \right)^2 \right] \end{aligned}$$

$$\leq \mathcal{V}_{t-1}(\pi) \sum_{\tau=1}^{t-1} \mathbb{E}_{a \sim p_\tau(\cdot|x)} \left[\left(\widehat{\phi}_t(x, a) - \phi_\tau^*(x, a) \right)^2 \right].$$

Using [Lemma 30](#) and [Lemma 31](#) below, when $\widehat{\Delta}_{[1,t]} \leq \mathcal{O}(\rho(t))$, we have

$$\sum_{\tau=1}^{t-1} \mathbb{E}_{x \sim \mathcal{D}_\tau, a \sim p_\tau(\cdot|x)} \left[\left(\widehat{\phi}_t(x, a) - \phi_\tau^*(x, a) \right)^2 \right] \leq \mathcal{O}(\log(T|\Phi|/\delta))$$

and $\mathcal{V}_{t-1}(\pi) \leq \mathcal{O}(A) + \max_{\tau \in [1, t-2]} \gamma_\tau \widehat{\text{Reg}}_{[1,\tau]}(\pi)$, where $\gamma_t = \Theta\left(\sqrt{\frac{At}{\log(T|\Phi|/\delta)}}\right)$. Note that they are actually of the same order as in the [Lemma 7](#) of ([Simchi-Levi and Xu, 2020](#)) since the additional terms contributed by $\Delta_{[1,t]}$ are dominated by other terms. Thus, the bound we get for $|\mathcal{R}_{[1,t-1]} - \widehat{\mathcal{R}}_{[1,t-1]}|$ is of the same order as their [Lemma 7](#), which is

$$\begin{aligned} \left| \mathcal{R}_{[1,t-1]}(\pi) - \widehat{\mathcal{R}}_{[1,t-1]}(\pi) \right| &\leq \mathcal{O} \left(\sqrt{\frac{\log(T|\Phi|/\delta)}{t}} \left(A + \max_{\tau \in [1, t-2]} \gamma_\tau \widehat{\text{Reg}}_{[1,\tau]}(\pi) \right) \right) \quad (45) \\ &\leq \frac{1}{16} \max_{\tau \in [1, t-2]} \widehat{\text{Reg}}_{[1,\tau]}(\pi) + \mathcal{O}(\rho(t)). \quad (\text{by AM-GM}) \end{aligned}$$

(46)

Then one can follow the derivation in their [Lemma 8](#) using [Eq. \(46\)](#), and get

$$\begin{aligned} \text{Reg}_{[1,t]}(\pi) - \widehat{\text{Reg}}_{[1,t]}(\pi) &\leq \frac{1}{8} \max_{\tau \in [1, t-1]} \widehat{\text{Reg}}_{[1,\tau]}(\pi) + \mathcal{O}(\rho(t)), \\ \widehat{\text{Reg}}_{[1,t]}(\pi) - \text{Reg}_{[1,t]}(\pi) &\leq \frac{1}{8} \max_{\tau \in [1, t-1]} \widehat{\text{Reg}}_{[1,\tau]}(\pi) + \mathcal{O}(\rho(t)). \end{aligned}$$

Using these two inequalities, together with $|\text{Reg}_{[1,\tau]}(\pi) - \text{Reg}_{[1,t]}(\pi)| \leq \mathcal{O}(\Delta_{[1,t]}) = \mathcal{O}(\rho(t))$, we can also prove

$$\text{Reg}_{[1,t]}(\pi) \leq 2\widehat{\text{Reg}}_{[1,t]}(\pi) + \mathcal{O}(\rho(t)), \quad \widehat{\text{Reg}}_{[1,t]}(\pi) \leq 2\text{Reg}_{[1,t]}(\pi) + \mathcal{O}(\rho(t)) \quad (47)$$

by induction as their [Lemma 8](#). One can see that all bounds we obtain are of the same order as in the stationary case shown in ([Simchi-Levi and Xu, 2020](#)), thanks to the condition $\Delta_{[1,t]} = \mathcal{O}(\rho(t))$.

Then following their [Lemmas 9](#) and [10](#), we obtain regret bound $\max_\phi t\mathcal{R}_{[1,t]}(\pi_\phi) - \sum_{\tau=1}^t R_\tau = \mathcal{O}\left(\sqrt{At \log(T|\Phi|/\delta)}\right)$.

Similar to the calculation in [Eq. \(43\)](#), by picking $\widetilde{f}_t = \widehat{\mathcal{R}}_{[1,t-1]}(\pi_{\widehat{\phi}_t}) + c\sqrt{\frac{A \log(T|\Phi|/\delta)}{t}}$ with large enough c , we have

$$\begin{aligned} \widetilde{f}_t &\geq \max_\phi \mathcal{R}_{[1,t-1]}(\pi_\phi) + c\sqrt{\frac{A \log(T|\Phi|/\delta)}{t}} - \mathcal{O}(\rho(t)) \quad (\text{by Eq. (46)}) \\ &\geq \mathcal{R}_{[1,t-1]}(\pi_{\phi_1^*}) \geq \mathcal{R}_1(\pi_{\phi_1^*}) - \mathcal{O}(\Delta_{[1,t]}) \geq \min_{\tau \in [1,t]} \max_\phi \mathcal{R}_\tau(\pi_\phi) - \mathcal{O}(\Delta_{[1,t]}), \end{aligned}$$

which verifies [Eq. \(31\)](#). To upper bound $\sum_{\tau=1}^t (\widetilde{f}_\tau - R_\tau)$, we follow a similar calculation as [Eq. \(44\)](#), and use the condition $\Delta_{[1,t]} = \mathcal{O}(\rho(t))$. This verifies [Eq. \(32\)](#).

Lemma 30 *If $\Delta_{[1,t]} \leq \mathcal{O}(\rho(t))$, then*

$$\sum_{\tau=1}^{t-1} \mathbb{E}_{x \sim \mathcal{D}_\tau, a \sim p_\tau(\cdot|x)} \left[\left(\widehat{\phi}_t(x, a) - \phi_\tau^*(x, a) \right)^2 \right] = \mathcal{O}(\log(T|\Phi|/\delta)).$$

Proof First, we consider a specific ϕ . Define $Y_\tau = (\phi(x_\tau, a_\tau) - R_\tau)^2 - (\phi_\tau^*(x_\tau, a_\tau) - R_\tau)^2$. Then we have $\mathbb{E}[Y_\tau] = \mathbb{E}[(\phi(x_\tau, a_\tau) - \phi_\tau^*(x_\tau, a_\tau))^2]$ and $\mathbb{E}[Y_\tau^2] \leq 4\mathbb{E}[(\phi(x_\tau, a_\tau) - \phi_\tau^*(x_\tau, a_\tau))^2] = 4\mathbb{E}[Y_\tau]$. By Freedman's inequality,

$$\begin{aligned} \sum_{\tau=1}^{t-1} Y_\tau &\geq \sum_{\tau=1}^{t-1} \mathbb{E}[Y_\tau] - c_1 \sqrt{\sum_{\tau=1}^{t-1} \mathbb{E}[Y_\tau^2] \log(T/\delta)} - c_2 \log(T|\Phi|/\delta) \\ &\geq \sum_{\tau=1}^{t-1} \mathbb{E}[Y_\tau] - 2c_1 \sqrt{\sum_{\tau=1}^{t-1} \mathbb{E}[Y_\tau] \log(T/\delta)} - c_2 \log(T|\Phi|/\delta). \end{aligned}$$

The above implies (by solving for $\sum_{\tau=1}^{t-1} \mathbb{E}[Y_\tau]$)

$$\sum_{\tau=1}^{t-1} \mathbb{E}[Y_\tau] \leq 2 \sum_{\tau=1}^{t-1} Y_\tau + 4(c_1^2 + c_2) \log(T|\Phi|/\delta). \quad (48)$$

For the other direction, we also have

$$\sum_{\tau=1}^{t-1} Y_\tau \leq 2 \sum_{\tau=1}^{t-1} \mathbb{E}[Y_\tau] + \left(\frac{c_1^2}{4} + c_2 \right) \log(T|\Phi|/\delta). \quad (49)$$

Then we can bound

$$\begin{aligned} &\sum_{\tau=1}^{t-1} \mathbb{E}_{x \sim \mathcal{D}_\tau, a \sim p_\tau(\cdot|x)} \left[\left(\widehat{\phi}_t(x, a) - \phi_\tau^*(x, a) \right)^2 \right] \quad (\text{using Eq. (48)}) \\ &\leq 2 \sum_{\tau=1}^{t-1} \left(\widehat{\phi}_t(x_\tau, a_\tau) - R_\tau \right)^2 - 2 \sum_{\tau=1}^{t-1} \left(\phi_\tau^*(x_\tau, a_\tau) - R_\tau \right)^2 + 4(c_1^2 + c_2) \log(T|\Phi|/\delta) \\ &\leq 2 \sum_{\tau=1}^{t-1} \left(\phi_1^*(x_\tau, a_\tau) - R_\tau \right)^2 - 2 \sum_{\tau=1}^{t-1} \left(\phi_\tau^*(x_\tau, a_\tau) - R_\tau \right)^2 + 4(c_1^2 + c_2) \log(T|\Phi|/\delta) \\ &\quad (\text{by the optimality of } \widehat{\phi}_t) \\ &\leq 4 \sum_{\tau=1}^{t-1} \mathbb{E} \left[\left(\phi_1^*(x_\tau, a_\tau) - \phi_\tau^*(x_\tau, a_\tau) \right)^2 \right] + c_3 \log(T|\Phi|/\delta) \quad (\text{using Eq. (49)}) \\ &\leq \frac{4(t-1)}{A} \Delta_{[1,t]}^2 + c_3 \log(T|\Phi|/\delta). \quad (\text{by the definition of } \Delta_{[\cdot, \cdot]}) \end{aligned}$$

By the condition on $\Delta_{[1,t]}$, we have $\frac{4(t-1)}{A} \Delta_{[1,t]}^2 = \mathcal{O}(\log(T|\Phi|/\delta))$, which proves the lemma. \blacksquare

Lemma 31 *If $\Delta_{[1,t]} \leq \mathcal{O}(\rho(t))$, then*

$$\mathcal{V}_t(\pi) \leq \mathcal{O}(A) + \max_{\tau \in [1, t-1]} \gamma_\tau \widehat{\text{Reg}}_{[1, \tau]}(\pi)$$

where $\gamma_t = \Theta\left(\sqrt{\frac{At}{\log(|\Phi|T/\delta)}}\right)$.

Proof Similar to Lemma 6 of FALCON, for $\tau \in [1, t]$,

$$\begin{aligned} V_\tau(p_\tau, \pi) &\leq A + \gamma_{\tau-1} \mathbb{E}_{x \sim \mathcal{D}_\tau} \left[\widehat{\phi}_\tau(x, \pi_{\widehat{\phi}_\tau}(x)) - \widehat{\phi}_\tau(x, \pi(x)) \right] \\ &\leq A + \frac{\gamma_{\tau-1}}{\tau-1} \sum_{s=1}^{\tau-1} \mathbb{E}_{x \sim \mathcal{D}_s} \left[\widehat{\phi}_\tau(x, \pi_{\widehat{\phi}_\tau}(x)) - \widehat{\phi}_\tau(x, \pi(x)) \right] + \gamma_{\tau-1} \Delta_{[1, t]} \\ &\leq A + \gamma_{\tau-1} \widehat{\text{Reg}}_{[1, \tau-1]}(\pi) + \gamma_{\tau-1} \Delta_{[1, t]}. \end{aligned}$$

By the condition $\widehat{\Delta}_{[1, t]} \leq \mathcal{O}(\rho(t))$, that last term $\gamma_{\tau-1} \Delta_{[1, t]}$ is of order $\mathcal{O}(A)$. By the definition of $\mathcal{V}_t(\pi)$, this finishes the proof. \blacksquare