

---

# Provably Safe Reinforcement Learning with Step-wise Violation Constraints

---

Nuoya Xiong<sup>1,\*</sup>, Yihan Du<sup>2,†</sup>, Longbo Huang<sup>1,‡,\*</sup>

<sup>1</sup>Institute for Interdisciplinary Information Sciences, Tsinghua University

<sup>2</sup>University of Illinois at Urbana-Champaign

\*nuoyaxiong@gmail.com

†yihandu@illinois.edu

‡longbohuang@tsinghua.edu.cn

## Abstract

We investigate a novel safe reinforcement learning problem with step-wise violation constraints. Our problem differs from existing works in that we focus on stricter step-wise violation constraints and do not assume the existence of safe actions, making our formulation more suitable for safety-critical applications that need to ensure safety in all decision steps but may not always possess safe actions, e.g., robot control and autonomous driving. We propose an efficient algorithm SUCBVI, which guarantees an  $\tilde{O}(\sqrt{ST})$  or gap-dependent  $\tilde{O}(S/C_{\text{gap}} + S^2AH^2)$  step-wise violation and an  $\tilde{O}(\sqrt{H^3SAT})$  regret. Lower bounds are provided to validate the optimality in both violation and regret performance with respect to the number of states  $S$  and the total number of steps  $T$ . Moreover, we further study an innovative safe reward-free exploration problem with step-wise violation constraints. For this problem, we design the algorithm SRF-UCRL to find a near-optimal safe policy, which achieves a nearly state-of-the-art sample complexity  $\tilde{O}((\frac{S^2AH^2}{\epsilon} + \frac{H^4SA}{\epsilon^2})(\log(\frac{1}{\delta}) + S))$ , and guarantees an  $\tilde{O}(\sqrt{ST})$  violation during exploration. Experimental results demonstrate the superiority of our algorithms in safety performance and corroborate our theoretical results.

## 1 Introduction

In recent years, reinforcement learning (RL) (Sutton & Barto, 2018) has become a powerful framework for decision-making and learning in unknown environments. Despite the ground-breaking success of RL in games (Lanctot et al., 2019), recommendation systems (Afsar et al., 2022) and complex tasks in simulation environments (Zhao et al., 2020), most existing RL algorithms focus on optimizing the cumulative reward and do not take into consideration the risk aspect, e.g., the agent runs into catastrophic situations during control. The lack of strong safety guarantees hinders the application of RL to broader safety-critical scenarios such as autonomous driving, robotics and healthcare. For example, for robotic control in complex environments, it is crucial to prevent the robot from getting into dangerous situations, e.g., hitting walls or falling into water pools, at all times.

To handle the safety requirement, a common approach is to formulate safety as a long-term expected violation constraint in each episode. This approach focuses on seeking a policy whose cumulative expected violation in each episode is below a certain threshold. However, for applications where an agent needs to avoid disastrous situations throughout the decision process, e.g., a robot needs to avoid hitting obstacles at each step, merely reducing the long-term expected violation is not sufficient to guarantee safety.

---

\*Corresponding author.

Motivated by this fact, we investigate safe reinforcement learning with a more fine-grained constraint, called *step-wise* violation constraint, which aggregates all nonnegative violations at each step (no offset between positive and negative violations permitted). We name this problem Safe-RL-SW. Our step-wise violation constraint differs from prior expected violation constraint (Wachi & Sui, 2020; Efroni et al., 2020b; Kalagarla et al., 2021) in two aspects: (i) Minimizing the step-wise violation enables the agent to learn an optimal policy that avoids unsafe regions deterministically, while reducing the expected violation only guarantees to find a policy with low expected violation, instead of a per-step zero-violation policy. (ii) Reducing the aggregated nonnegative violation allows us to have a risk control for each step, while a small cumulative expected violation can still result in a large cost at some individual step and cause danger, if other steps with smaller costs offset the huge cost.

Our problem faces two unique challenges. First, the step-wise violation requires us to guarantee a small violation at each step, which demands very different algorithm design and analysis from that for the expected violation (Wachi & Sui, 2020; Efroni et al., 2020b; Kalagarla et al., 2021). Second, in safety-critical scenarios, the agent needs to identify not only unsafe states but also potentially unsafe states, which are states that may appear to be safe but will ultimately lead to unsafe regions with a non-zero probability. For example, a self-driving car needs to learn to slow down or change directions early, foreseeing the potential danger in advance in order to ensure safe driving (Thomas et al., 2021). Existing safe RL works focus mainly on the expected violation (Wachi & Sui, 2020; Liu et al., 2021b; Wei et al., 2022), or requiring some other assumptions such as imposing the prior knowledge of a safe action for each state (Amani et al., 2021). Moreover, many previous works also require strong assumptions that exclude discrete tabular MDP (Amani et al., 2021; Wachi et al., 2021; Wang et al., 2023) which is considered in our paper. Hence, techniques in previous works cannot be applied to handle step-wise violations. More detailed comparisons are provided in Section 2.

To systematically handle these two challenges, we formulate safety as an unknown cost function for each state without assuming safe actions, and consider minimizing the step-wise violation instead of the expected violation. We propose a general algorithmic framework called **Safe UCBVI** (SUCBVI). Specifically, in each episode, we first estimate the transition kernel and cost function in an optimistic manner, tending to regard a state as safe at the beginning of learning. After that, we introduce novel dynamic programming to identify potentially unsafe states and determine safe actions, based on our estimated transition and costs. Finally, we employ the identified safe actions to conduct value iteration. This mechanism can adaptively update dangerous regions, and help the agent plan for the future, which keeps her away from all states that may lead to unsafe states. As our estimation becomes more accurate over time, the safety violation becomes smaller and eventually converges to zero. Note that without the assumption of safe actions, the agent knows nothing about the environment at the beginning. Thus, it is impossible for her to achieve an absolute zero violation. Nevertheless, we show that SUCBVI can achieve a sub-linear  $\tilde{O}(\sqrt{ST})$  cumulative violation or an  $\tilde{O}(S/C_{\text{gap}} + S^2AH^2)$  gap-dependent violation that is independent of  $T$ . This violation implies that as the RL game proceeds, the agent eventually learns how to avoid unsafe states. We also provide a matching lower bound to demonstrate the optimality of SUCBVI in both violation and regret.

Furthermore, we apply our step-wise safe RL framework to the reward-free exploration (Jin et al., 2020) setting. In this novel safe reward-free exploration, the agent needs to guarantee small step-wise violations during exploration, and also output a near-optimal safe policy. Our algorithm achieves  $\varepsilon$  cumulative step-wise violation during exploration, and also identifies a  $\varepsilon$ -optimal and safe policy. See the definition of  $\varepsilon$ -optimal and  $\varepsilon$ -safe policy in Section 6.1. Another interesting application of our framework is safe zero-sum Markov games, which we discuss in Appendix B.

The main contributions of our paper are as follows.

- We formulate the safe RL with step-wise violation constraint problem (Safe-RL-SW), which models safety as a cost function for states and aims to minimize the cumulative step-wise violation. Our formulation is particularly useful for safety-critical applications where avoiding disastrous situations at each decision step is desirable, e.g., autonomous driving and robotics.
- We provide a general algorithmic framework SUCBVI, which is equipped with an innovative dynamic programming to identify potentially unsafe states and distinguish safe actions. We establish an  $\tilde{O}(\sqrt{H^3SAT})$  regret and an  $\tilde{O}(\sqrt{ST})$  or  $\tilde{O}(S/C_{\text{gap}} + S^2AH^2)$  gap-dependent violation guarantees, which exhibits the capability of SUCBVI in attaining high rewards while maintaining small violation.

- We further establish an  $\Omega(\sqrt{HST})$  regret and an  $\Omega(\sqrt{ST})$  violation lower bounds for Safe-RL-SW. The lower bounds demonstrate the optimality of algorithm SUCBVI in both regret minimization and safety guarantee, with respect to factors  $S$  and  $T$ .
- We consider step-wise safety constraints in the reward-free exploration setting, which is called the Safe-RFE-SW problem. In this setting, we design an efficient algorithm SRF-UCRL, which ensures  $\varepsilon$  step-wise violations during exploration and plans a  $\varepsilon$ -optimal and  $\varepsilon$ -safe policy for any reward functions with probability at least  $1 - \delta$ . We obtain an  $\tilde{O}((\frac{S^2AH^2}{\varepsilon} + \frac{H^4SA}{\varepsilon^2})(\log(\frac{1}{\delta}) + S))$  sample complexity and an  $\tilde{O}(\sqrt{ST})$  violation guarantee for SRF-UCRL, which shows the efficiency of SRF-UCRL in sampling and danger avoidance even without reward signals. To the best of our knowledge, this work is the first to study the step-wise violation constraint in the RFE setting.

## 2 Related Work

**Safe RL.** Safety is an important topic in RL, which has been extensively studied. The constrained Markov decision process (CMDP)-based approaches handle safety via cost functions, and aim to minimize the expected episode-wise violation, e.g., (Yu et al., 2019; Wachi & Sui, 2020; Qiu et al., 2020; Efroni et al., 2020b; Turchetta et al., 2020; Ding et al., 2020; Singh et al., 2020; Kalagarla et al., 2021; Simão et al., 2021; Ding et al., 2021), or achieve zero episode-wise violation, e.g., (Liu et al., 2021b; Bura et al., 2021; Wei et al., 2022; Sootla et al., 2022). Apart from CMDP-based approaches, there are also other works that tackle safe RL by control-based approaches (Berkenkamp et al., 2017; Chow et al., 2018; Dalal et al., 2018; Wang et al., 2022), policy optimization (Uchibe & Doya, 2007; Achiam et al., 2017; Tessler et al., 2018; Liu et al., 2020; Stooke et al., 2020) and safety shields (Alshiekh et al., 2018).

In recent years, there are also some works studying step-wise violations with additional assumptions. Now we provide detailed comparisons between existing papers with instantaneous constraints. Turchetta et al. (2016); Wachi et al. (2018) propose a GP-based algorithm, which assumes that the transition is deterministic and known, while modeling the reward and cost functions using Gaussian Processes. By using this particular structure, they can infer the safety cost by estimating the parameters. Wachi et al. (2021) assumes the reward and cost functions have a generalized linear structure. Their algorithm explores in a safe space until a time  $t^*$  when the agent explores sufficiently. However, the upper bound of the exploring time  $t^*$  is not given in their paper. In fact, under the tabular MDP setting  $t^*$  can be infinite.

Amani et al. (2021) further considers the reward and cost functions to have a linear structure. It makes two assumptions: (a) There exists a safe action in each state, which prevents the agent from going to a potentially unsafe state. (b) The feature set is a star convex set, which helps them change actions continuously. However, this makes their works infeasible in tabular MDPs: The feature set in tabular MDPs consists of one-hot vectors and is not a star convex set. Hence, the work Amani et al. (2021) cannot solve our problem. Shi et al. (2023) considers safe RL in linear mixture MDPs. Their work also contains assumption (b), making it infeasible in tabular MDPs. Moreover, although they do not have assumption (a), it assumes the transition set  $\Delta(s, a)$  is known, which is not needed in our paper. Thus our paper is more challenging since we need to estimate  $\Delta(s, a)$  in our algorithm adaptively.

There are some other papers investigating the safety of RL problems. Alshiekh et al. (2018) represents the safe state by the reactive system, and uses shielding to calculate and restrict the agent within a safe trajectory completely. The main difference between their work and our work is that we need to dynamically update the estimated safe state, while they require to know the mechanism and state to calculate the shield. Dalal et al. (2018) considers restricting the safe action by projecting the action into the closest safe actions. They achieve this goal by solving a convex optimization problem on the continuous action set. However, in their paper, they do not consider the situation where a state can have no safe actions. To be more specific, they do not consider the situation when the convex optimization problem has no solutions. Le et al. (2019) considers the decision-making problem with a pre-collected dataset. Turchetta et al. (2020) and Sootla et al. (2022) both consider cumulative cost constraints rather than step-wise constraints. The former uses a teacher for intervention to keep the agent away from the unsafe region, while the latter encourages safe exploration by augmenting a safety state to measure safety during training.

**Reward-free Exploration with Safety.** Motivated by sparse reward signals in realistic applications, reward-free exploration (RFE) has been proposed in Jin et al. (2020), and further developed

in Kaufmann et al. (2021); Ménard et al. (2021). In RFE, the agent explores the environment without reward signals. After enough exploration, the reward function is given, and the agent needs to plan a near-optimal policy based on his knowledge collected in exploration. Safety is also important in RFE: We need to not only guarantee that our outputted policy is safe, but also ensure small violations during exploration.

Recently, Miryoosefi & Jin (2022); Huang et al. (2022) also study RFE with safety constraints. Compared to our paper, Miryoosefi & Jin (2022) only considers the safety constraint after the exploration phase. Huang et al. (2022) differs from our work in the following aspects: (i) They allow different reward functions during exploration and after exploration, and the agent can directly know the true costs during training. In our work, the cost function is the same during and after exploration, but the agent can only observe the noisy costs of a state when she arrives at that state. (ii) They require prior knowledge of safe baseline policies, while we do not need such an assumption. (iii) They consider the zero-expected violation during exploration, while we focus on keeping small step-wise violations.

### 3 The Safe MDP Model

**Episodic MDP.** In this paper, we consider the finite-horizon episodic Markov Decision Process (MDP), represented by a tuple  $(\mathcal{S}, \mathcal{A}, H, \mathbb{P}, r)$ . Here  $\mathcal{S}$  is the state space,  $\mathcal{A}$  is the action space, and  $H$  is the length of each episode.  $\mathbb{P} = \{\mathbb{P}_h : \mathcal{S} \times \mathcal{A} \mapsto \Delta_{\mathcal{S}}\}_{h \in [H]}$  is the transition kernel, and  $\mathbb{P}_h(s' | s, a)$  gives the transition probability from  $(s, a)$  to  $s'$  at step  $h$ .  $r = \{r_h : \mathcal{S} \times \mathcal{A} \mapsto [0, 1]\}_{h \in [H]}$  is the reward function, and  $r_h(s, a)$  gives the reward of taking action  $a$  in state  $s$  at step  $h$ . A policy  $\pi = \{\pi_h : \mathcal{S} \mapsto \mathcal{A}\}_{h \in [H]}$  consists of  $H$  mappings from the state space to action space. In each episode  $k$ , the agent first chooses a policy  $\pi^k$ . At each step  $h \in [H]$ , the agent observes a state  $s_h^k$ , takes an action  $a_h^k$ , and then goes to a next state  $s_{h+1}^k$  with probability  $\mathbb{P}_h(s_{h+1}^k | s_h^k, a_h^k)$ . The algorithm executes  $T = HK$  steps. Moreover, the state value function  $V_h^\pi(s, a)$  and state-action value function  $Q_h^\pi(s, a)$  for a policy  $\pi$  can be defined as

$$V_h^\pi(s) := \mathbb{E}_\pi \left[ \sum_{h'=h}^H r_{h'}(s_{h'}, \pi_{h'}(s_{h'})) \middle| s_h = s \right],$$

$$Q_h^\pi(s, a) := \mathbb{E}_\pi \left[ \sum_{h'=h}^H r_{h'}(s_{h'}, \pi_{h'}(s_{h'})) \middle| s_h = s, a_h = a \right].$$

**Safety Constraint.** To model unsafe regions in the environment, similar to Wachi & Sui (2020); Yu et al. (2022), we define a safety cost function  $c : \mathcal{S} \mapsto [0, 1]$ . Let  $\tau \in [0, 1]$  denote the safety threshold. A state is called *safe* if  $c(s) \leq \tau$ , and called *unsafe* if  $c(s) > \tau$ . Similar to Efroni et al. (2020b); Amani et al. (2021), when the agent arrives in a state  $s$ , she will receive a cost signal  $z(s) = c(s) + \zeta$ , where  $\zeta$  is an independent, zero-mean and 1-sub-Gaussian noise. Denote  $(x)_+ = \max\{x, 0\}$ . The violation in state  $s$  is defined as  $(c(s) - \tau)_+$ , and the cumulative step-wise violation till episode  $K$  is

$$C(K) = \sum_{k=1}^K \sum_{h=1}^H (c(s_h^k) - \tau)_+. \quad (1)$$

Eq. (1) represents the accumulated step-wise violation during training. When the agent arrives in state  $s_h^k$  at step  $h$  in episode  $k$ , she will suffer violation  $(c(s_h^k) - \tau)_+$ . This violation setting is significantly different from the previous CMDP setting (Qiu et al., 2020; Ding et al., 2020; Efroni et al., 2020b; Ding et al., 2021; Wachi & Sui, 2020; Liu et al., 2021a; Kalagarla et al., 2021). They study the *episode-wise expected violation*  $C'(K) = \sum_{k=1}^K \left( \mathbb{E} \left[ \sum_{h=1}^H c(s_h^k, a_h^k) \right] - \mu \right)$ . There are also some papers (Efroni et al., 2020a; Simão et al., 2021) considering a stricter constraint named *episode-wise clipped expected violation*:  $C''(K) = \sum_{k=1}^K \left( \mathbb{E} \left[ \sum_{h=1}^H c(s_h^k, a_h^k) \right] - \mu \right)_+$ . Compared to the episode-wise violation (including expected violation and clipped expected violation), our step-wise violation has two main differences: (i) First, the episode-wise violation constraints allow the agent to get into unsafe states occasionally. Instead, the step-wise violation constraint forces the agent to stay in safe regions at all times. (ii) Second, in the episode-wise constraints, the average violation

$\mathbb{E}[c(s_h, a_h)] - \mu/H$  at step  $h$  is allowed to be positive or negative, and they can cancel out in one episode to achieve  $\mathbb{E}[\sum_{h=1}^H c(s_h^k, a_h^k)] \leq \mu$ . Instead, we consider a *nonnegative* function  $(c(s) - \tau)_+$  at each step in our step-wise violation, which imposes a stricter constraint.

Define  $\mathcal{U} := \{s \in \mathcal{S} \mid c(s) > \tau\}$  as the set of all *unsafe states*. Let  $\iota = \{s_1, a_1, \dots, s_H, a_H\}$  denote a trajectory. Since a feasible policy needs to satisfy the constraint at each step, we define the set of feasible policies as  $\Pi = \{\pi \mid \Pr\{\exists h \in [H], s_h \in \mathcal{U} \mid \iota \sim \pi\} = 0\}$ . The feasible policy set  $\Pi$  consists of all policies under which one never reaches any unsafe state in an episode.

**Learning Objective.** In this paper, we consider the regret minimization objective. Specifically, define  $\pi^* = \operatorname{argmax}_{\pi \in \Pi} V_1^\pi$ ,  $V^* = V^{\pi^*}$  and  $Q^* = Q^{\pi^*}$ . The regret till  $K$  episodes is then defined as

$$R(K) = \sum_{t=1}^K (V_1^*(s_1) - V_1^{\pi^k}(s_1)),$$

where  $\pi^k$  is the policy taken in episode  $k$ . Our objective is to minimize  $R(K)$  to achieve a good performance, and minimize the violation  $C(K)$  to guarantee the safety at the same time.

## 4 Safe RL with Step-wise Violation Constraints

### 4.1 Assumptions and Problem Features

Before introducing our algorithms, we first state the important assumptions and problem features for Safe-RL-SW.

Suppose  $\pi$  is a feasible policy. Then, if we arrive at  $s_{H-1}$  at step  $H-1$ ,  $\pi$  needs to select an action that guarantees  $s_H \notin \mathcal{U}$ . Define the *transition set*  $\Delta_h(s, a) = \{s' \mid \mathbb{P}_h(s' \mid s, a) > 0\}$  for any  $\forall (s, a, h) \in \mathcal{S} \times \mathcal{A} \times [H]$ , which represents the set of possible next states after taking action  $a$  in state  $s$  at step  $h$ . Then, at the former step  $H-1$ , the state  $s$  is potentially unsafe if it satisfies that  $\Delta_{H-1}(s, a) \cap \mathcal{U} \neq \emptyset$  for all  $a \in \mathcal{A}$ . (i.e., no matter taking what action, there is a positive probability of transitioning to an unsafe next state). Therefore, we can recursively define the set of *potentially unsafe states at step  $h$*  as

$$\mathcal{U}_h = \mathcal{U}_{h+1} \cup \{s \mid \forall a \in \mathcal{A}, \Delta_h(s, a) \cap \mathcal{U}_{h+1} \neq \emptyset\}, \quad (2)$$

where  $\mathcal{U}_H = \mathcal{U}$ . Intuitively, if we are in a state  $s_h \in \mathcal{U}_h$  at step  $h$ , no action can be taken to completely avoid reaching potentially unsafe states  $s_{h+1} \in \mathcal{U}_{h+1}$  at step  $h+1$ . Thus, in order to completely prevent from getting into unsafe states  $\mathcal{U}$  throughout all steps, one needs to avoid potentially unsafe states in  $\mathcal{U}_h$  at step  $h$ . From the above argument, we have that  $s_1 \notin \mathcal{U}_1$  is equivalent to the existence of feasible policies. The detailed proof is provided in Appendix D. Thus, we make the following necessary assumption.

**Assumption 4.1** (Existence of feasible policies). The initial state  $s_1$  satisfies  $s_1 \notin \mathcal{U}_1$ .

For any  $s \in \mathcal{S}$  and  $h \in [H-1]$ , we define the set of safe actions for state  $s$  at step  $h$  as

$$A_h^{safe}(s) = \{a \in \mathcal{A} \mid \Delta_h(s, a) \cap \mathcal{U}_{h+1} = \emptyset\}, \quad (3)$$

and let  $A_H^{safe}(s) = \mathcal{A}$ .  $A_h^{safe}(s)$  stands for the set of all actions at step  $h$  which will not lead to potentially unsafe states in  $\mathcal{U}_{h+1}$ . Here,  $\{\mathcal{U}_h\}_{h \in [H]}$  and  $\{A_h^{safe}(s)\}_{h \in [H]}$  are defined by dynamic programming: If we know sets of all possible next state  $\{\Delta_h(s, a)\}_{h \in [H]}$  and unsafe state set  $\mathcal{U} = \mathcal{U}_H$ , we can calculate all potentially unsafe state sets  $\{\mathcal{U}_h\}_{h \in [H]}$  and safe action sets  $\{A_h^{safe}(s)\}_{h \in [H]}$ , and choose feasible policies to completely avoid unsafe states.

### 4.2 Algorithm SUCBVI

Now we present our main algorithm **Safe UCBVI (SUCBVI)**, which is based on previous classic RL algorithm UCBVI (Azar et al., 2017), and equipped with a novel dynamic programming to identify potentially unsafe states and safe actions. The pseudo-code is shown in Algorithm 1. First, we provide some intuitions about how SUCBVI works. At each episode, we first estimate all the unsafe states based on the historical data. Then, we perform a dynamic programming procedure introduced

---

**Algorithm 1** SUCBVI

---

```
1: Initialize:  $\Delta_h^1(s, a) = \emptyset$ ,  $N_h^1(s, a) = N_h^1(s, a, s') = 0$  for all  $s \in \mathcal{S}$ ,  $a \in \mathcal{A}$ ,  $s' \in \mathcal{S}$  and  $h \in [H]$ .
2: for  $k = 1, 2, \dots, K$  do
3:   Update the optimistic estimates of cost.
4:   Update the empirical cost  $\hat{c}(s)$  and calculate  $\bar{c}(s) = \hat{c}(s) - \beta(N^k(s), \delta)$  for all  $s \in \mathcal{S}$ .  $\triangleright$ 
5:   Define  $\mathcal{U}_H^k = \{s \mid \bar{c}(s) > \tau\}$  and  $\mathcal{U}_h^k = \mathcal{U}_{h+1}^k \cup \{s \mid \forall a \in \mathcal{A}, \Delta_h^k(s, a) \cap \mathcal{U}_{h+1}^k \neq \emptyset\}$  for all
    $h \in [H]$  recursively. Calculate  $\{A_h^{k, safe}(s)\}_{h \in [H]}$  by Eq. (3) with  $\{\mathcal{U}_h^k\}_{h \in [H]}$ .
    $\triangleright$  Perform value iteration with previous estimates.
6:   for  $h = H, H-1, \dots, 1$  do
7:     for  $s \in \mathcal{S}$  do
8:       for  $a \in \mathcal{A}$  do
9:         Compute  $\hat{\mathbb{P}}_h^k(s' \mid s, a) = \frac{N_h^k(s, a, s')}{N_h^k(s, a)}$ .
10:         $Q_h^k(s, a) = \min\{H, r_h(s, a) + \sum_{s'} \hat{\mathbb{P}}_h^k(s' \mid s, a) V_{h+1}^k(s') + \alpha(N_h^k(s, a))\}$ .
11:      end for
12:      if  $s \notin \mathcal{U}_h^k$  then
13:         $V_h^k(s) = \max_{a \in A_h^{k, safe}(s)} Q_h^k(s, a)$ ,  $\pi_h^k(s) = \arg \max_{a \in A_h^{k, safe}(s)} Q_h^k(s, a)$ .
14:      else
15:         $V_h^k(s) = \max_{a \in \mathcal{A}} Q_h^k(s, a)$ ,  $\pi_h^k(s) = \arg \max_{a \in \mathcal{A}} Q_h^k(s, a)$ .
16:      end if
17:    end for
18:  end for
19:  for  $h = 1, 2, \dots, H$  do
20:    Take action  $a_h^k = \pi_h^k(s_h^k)$  and observe state  $s_{h+1}^k$ .
     $\triangleright$  Update the estimates of  $\Delta(s, a)$ .
21:     $\Delta_h^{k+1}(s_h^k, a_h^k) = \Delta_h^k(s_h^k, a_h^k) \cup \{s_{h+1}^k\}$ . Increase  $N_h^{k+1}(s_h^k, a_h^k)$ ,  $N_h^{k+1}(s_h^k, a_h^k, s_{h+1}^k)$  by 1.
22:  end for
23: end for
```

---

in Section 4.1 and calculate the safe action set  $A_h^{safe}(s)$  for each state  $s$ . Then, we perform value iteration in the estimated safe action set. As the estimation becomes more accurate, SUCBVI will eventually avoid potentially unsafe states and achieve both sublinear regrets and violations.

Now we begin to introduce our algorithm. In the beginning, we initialize  $\Delta_h(s, a) = \emptyset$  for all  $(h, s, a) \in [H] \times \mathcal{S} \times \mathcal{A}$ . It implies that the agent considers all actions to be safe at first, because no action will lead to unsafe states from the agent's perspective. In each episode, we first estimate the empirical cost  $\hat{c}(s)$  based on historical cost feedback  $z(s)$ , and regard state  $s$  as safe if  $\bar{c}(s) = \hat{c}(s) - \beta > \tau$  for some bonus term  $\beta$ , which aims to guarantee  $\bar{c}(s) \leq c(s)$  (Line 4). Then, we calculate the estimated unsafe state set  $\mathcal{U}_H^k$ , which is a subset of the true unsafe state set  $\mathcal{U} = \mathcal{U}_H$  with a high probability by optimism. With  $\mathcal{U}_H^k$  and  $\Delta_h^k(s, a)$ , we can estimate potentially unsafe state sets  $\mathcal{U}_h^k$  for all  $h \in [H]$  by Eq. (2) recursively.

Then, we perform value iteration to compute the optimistically estimated optimal policy. Specifically, for any hypothesized safe state  $s \notin \mathcal{U}_h^k$ , we update the greedy policy on the estimated safe actions, i.e.,  $\pi_h^k(s) = \max_{a \in A_h^{k, safe}(s)} Q_h^k(s, a)$  (Line 13). On the other hand, for any hypothesized unsafe state  $s_h \in \mathcal{U}_h$ , since there is no action that can completely avoid unsafe states, we ignore safety costs and simply update the policy by  $\pi_h^k(s) = \max_{a \in \mathcal{A}} Q(s, a)$  (Line 15). After that, we calculate the estimated optimal policy  $\pi^k$  for episode  $k$ , and the agent follows  $\pi^k$  and collects a trajectory. Then, we update  $\Delta_h^{k+1}(s, a)$  by incorporating the observed state  $s_{h+1}^k$  into the set  $\Delta_h^k(s_h^k, a_h^k)$ . Under this updating rule, it holds that  $\Delta_h^k(s, a) \subseteq \Delta_h^{k+1}(s, a)$  for all  $(s, a) \in \mathcal{S} \times \mathcal{A}$ .

The performance of Algorithm 1 is summarized below in Theorem 4.2.

**Theorem 4.2.** *Let  $\alpha(n, \delta) = 7H \sqrt{\frac{\ln(5SAHK/\delta)}{n}}$  and  $\beta(n, \delta) = \sqrt{\frac{2}{n} \log(SK/\delta)}$ . With probability at least  $1 - \delta$ , the regret and step-wise violation of Algorithm 1 are bounded by*

$$R(K) = \tilde{O}(\sqrt{H^2SAT}), \quad C(K) = \tilde{O}(\sqrt{ST} + S^2AH^2).$$

Moreover, if  $\mathcal{C}_{\text{gap}} \triangleq \min_{s \in \mathcal{U}} (c(s) - \tau)_+ > 0$ , we have  $C(K) = \tilde{O}(S/\mathcal{C}_{\text{gap}} + S^2AH^2)$ .

Theorem 4.2 shows that SUCBVI achieves both sublinear regret and violation. Moreover, when all the unsafe states have a large cost compared to the safety threshold  $\tau$ , i.e.,  $\mathcal{C}_{\text{gap}} = \min_{s \in \mathcal{U}} (c(s) - \tau)_+ > 0$  is a constant, we can distinguish the unsafe states easily and get a constant violation. In particular, when  $\tau = 1$ , Safe-RL-SW degenerates to the unconstrained MDP and Theorem 4.2 maintains the same regret  $\tilde{O}(\sqrt{H^2SAT})$  as UCBVI-CH (Azar et al., 2017), while CMDP algorithms (Efroni et al., 2020b; Liu et al., 2021a) suffer a larger  $\tilde{O}(\sqrt{H^3S^3AT})$  regret.

We provide the analysis idea of Theorem 4.2 here. First, by the updating rule of  $\Delta_h^k(s, a)$ , we show that  $\mathcal{U}_h^k$  and  $A_h^{k, \text{safe}}(s)$  have the following crucial properties:  $\mathcal{U}_h^k \subseteq \mathcal{U}_h$ ,  $A_h^{\text{safe}}(s) \subseteq A_h^{k, \text{safe}}(s)$ . Based on this property, we can prove that  $Q_h^*(s, a) \leq Q_h^k(s, a)$  and  $V_h^*(s) \leq V_h^k(s)$  for all  $(s, a)$ , and then apply the regret decomposition techniques to derive the regret bound.

Recall that  $\pi^k$  is a feasible policy with respect to the estimated unsafe state set  $\mathcal{U}_H^k$  and transition set  $\{\Delta_h^k(s, a)\}_{h \in [H]}$ . If the agent takes policy  $\pi^k$  and the transition follows  $\Delta_h^k(s, a)$ , i.e.,  $s_{h+1}^k \in \Delta_h^k(s, a)$ , the agent never arrive any estimated unsafe state  $s \in \mathcal{U}_H^k$  in this episode. Hence the agent will suffer at most an  $\tilde{O}(\sqrt{ST})$  step-wise violation or an  $\tilde{O}(S/\mathcal{C}_{\text{gap}} + S^2AH^2)$  gap-dependent bounded violation. Yet, the situation  $s_{h+1} \in \Delta_h^k(s, a)$  does not always hold for all  $h \in [H]$ . In the case when  $s_{h+1}^k \notin \Delta_h^k(s, a)$ , we add the newly observed state  $s_{h+1}^k$  into  $\Delta_h^{k+1}(s_h^k, a_h^k)$  (Line 21). We can show that this case appears at most  $S^2AH$  times, and thus incurs  $O(S^2AH^2)$  additional violation. Combining the above two cases, the total violation can be upper bounded.

## 5 Lower Bounds for Safe-RL-SW

In this section, we provide a matching lower bound for Safe-RL-SW in Section 4. The lower bound shows that if an algorithm always achieves a sublinear regret in Safe-RL-SW, it must incur an  $\Omega(\sqrt{ST})$  violation. This result matches our upper bound in Theorem 4.2, showing that SUCBVI achieves the optimal violation performance.

**Theorem 5.1.** *If an algorithm has an expected regret  $\mathbb{E}_\pi[R(K)] \leq \frac{HK}{24}$  for all MDP instances, there exists an MDP instance in which the algorithm suffers expected violation  $\mathbb{E}_\pi[C(K)] = \Omega(\sqrt{ST})$ .*

Now we validate the optimality in terms of regret. Note that if we do not consider safety constraints, the lower bound for classic RL (Osband & Van Roy, 2016) can be applied to our setting. Thus, we also have an  $\Omega(\sqrt{T})$  regret lower bound. To understand the essential hardness brought by safety constraints, we further investigate whether safety constraints will lead to an  $\Omega(\sqrt{T})$  regret, given that we can achieve an  $o(\sqrt{T})$  regret on some good instances without the safety constraints.

**Theorem 5.2.** *For any  $\alpha \in (0, 1)$ , there exists a parameter  $n$  and  $n$  MDPs  $M_1, \dots, M_n$  satisfying that:*

1. *If we do not consider any constraint, there is an algorithm that achieves an  $\tilde{O}(T^{(1-\alpha)/2})$  regret compared to the unconstrained optimal policy on all MDPs.*
2. *If we consider the safety constraint, any algorithm with a  $O(T^{1-\alpha})$  expected violation will achieve an  $\Omega(\sqrt{HST})$  regret compared to the constrained optimal policy on one of MDPs.*

Intuitively, Theorem 5.2 shows that if one achieves sublinear violation, she must suffer at least an  $\Omega(\sqrt{T})$  regret even if she can achieve an  $o(\sqrt{T})$  regret without considering constraints. This theorem demonstrates the hardness particularly brought by the step-wise constraint, and corroborates the optimality of our results. Combining with Theorem 5.1, the two lower bounds show an essential trade-off between the violation and performance.

## 6 Safe Reward-Free Exploration with Step-wise Violation Constraints

### 6.1 Formulation of Safe-RFE-SW

In this section, we consider Safe RL in the reward-free exploration (RFE) setting (Jin et al., 2020; Kaufmann et al., 2021; Ménard et al., 2021) called Safe-RFE-SW, to show the generality of our proposed framework. In the RFE setting, the agent does not have access to reward signals and only receives random safety cost feedback  $z(s) = c(s) + \zeta$ . To impose safety requirements, Safe-RFE-SW requests the agent to keep small safety violations during exploration, and outputs a near-optimal safe policy after receiving the reward function.

**Definition 6.1** ( $(\varepsilon, \delta)$ -optimal safe algorithm for Safe-RFE-SW). An algorithm is  $(\varepsilon, \delta)$ -optimal safe for Safe-RFE-SW if it outputs the triple  $(\hat{\Pi}, \hat{\Delta}(s, a), \hat{\mathcal{U}}_H)$  such that for any reward function  $r$ , with probability at least  $1 - \delta$ ,

$$V_1^*(s_1; r) - V_1^{\hat{\pi}^*}(s_1; r) \leq \varepsilon, \quad \mathbb{E}_\pi \left[ \sum_{h=1}^H (c(s_h) - \tau)_+ \right] \leq \varepsilon, \quad \forall \pi \in \Pi \quad (4)$$

where  $\hat{\pi}^*$  is the optimal feasible policy with respect to  $(\hat{\Pi}, \hat{\Delta}(s, a), \hat{\mathcal{U}}_H, r)$ ,  $\Pi$  is the set of feasible policies with respect to  $(\Delta(s, a), \mathcal{U}_H)$ , and  $V(s; r)$  is the value function under reward function  $r$ . We say that a policy is  $\varepsilon$ -optimal if it satisfies the left inequality in Eq. (4) and  $\varepsilon$ -safe if it satisfies the right inequality in Eq. (4). We measure the performance by the number of episodes used before the algorithm terminates, i.e., *sample complexity*. Moreover, the cumulative step-wise violation till episode  $K$  is defined as  $C(K) = \sum_{k=1}^K \sum_{h=1}^H (c(s_h^k) - \tau)_+$ . In Safe-RFE-SW, our goal is to design an  $(\varepsilon, \delta)$ -optimal safe algorithm, and minimize both sample complexity and violation.

### 6.2 Algorithm SRF-UCRL

The Safe-RFE-SW problem requires us to consider extra safety constraints for both the exploration phase and final output policy, which needs new algorithm design and techniques compared to previous RFE algorithms. Also, the techniques for Safe-RL-SW in Section 4.2 are not sufficient for guaranteeing the safety of output policy, because SUCBVI only guarantees a step-wise violation during exploration.

We design an efficient algorithm **Safe RF-UCRL** (SRF-UCRL), which builds upon previous RFE algorithm RF-UCRL (Kaufmann et al., 2021). SRF-UCRL distinguishes potentially unsafe states and safe actions by backward iteration, and establishes a new uncertainty function to guarantee the safety of output policy. Algorithm 2 illustrates the procedure of SRF-UCRL. Specifically, in each episode  $k$ , we first execute a policy  $\pi_h^k$  computed from previous episodes, and then update the estimated next state set  $\{\Delta_h^k(s, a)\}_{h \in [H]}$  and unsafe state set  $\mathcal{U}_H^k$  by optimistic estimation. Then, we use Eq. (2) to calculate the unsafe state set  $\mathcal{U}_h$  for all steps  $h \in [H]$ . After that, we update the uncertainty function  $\overline{W}^k$  defined in Eq. (5) below and compute the policy  $\pi^{k+1}$  that maximizes the uncertainty to encourage more exploration in the next episode.

Now we provide the definition of the *uncertainty function*, which measures the estimation error between the empirical MDP and true MDP. For any safe state-action pair  $s \notin \mathcal{U}_h, a \in A_h^{k, safe}(s)$ , we define

$$\overline{W}_h^k(s, a) = \min \left\{ H, M(N_h^k(s, a), \delta) + \sum_{s'} \hat{\mathbb{P}}_h^k(s' | s, a) \max_{b \in A_{h+1}^{k, safe}(s')} \overline{W}_{h+1}^k(s', b) \right\}. \quad (5)$$

where  $M(N_h^k(s, a), \delta) = 2H \sqrt{\frac{2\gamma(N_h^k(s, a), \delta)}{N_h^k(s, a)}} + \frac{SH\gamma(N_h^k(s, a), \delta)}{N_h^k(s, a)}$ , and  $\gamma(n, \delta)$  is a logarithmic term that is formally defined in Theorem 6.2. For other state-action pairs  $(s, a)$ , we relax the restriction of  $b \in A_{h+1}^{k, safe}(s')$  to  $b \in \mathcal{A}$  in the max function. Our algorithm stops when  $\overline{W}_1^K(s_1, \pi^K(s_1))$  shrinks to within  $\varepsilon/2$ . Compared to previous RFE works (Kaufmann et al. (2021); Ménard et al. (2021)), our uncertainty function has two distinctions. First, for any safe state-action pair  $(s, a) \in (\mathcal{S} \setminus \mathcal{U}_h, A_h^{k, safe}(s))$ , Eq. (5) considers only safe actions  $A_{h+1}^{k, safe}(s)$ , which guarantees that the agent focuses on safe policies. Second, Eq. (5) incorporates another term  $(SH\gamma(N_h^k(s, a), \delta)/N_h^k(s, a))$  to control the expected violation for feasible policies. Now we present our result for Safe-RFE-SW.

---

**Algorithm 2** SRF-UCRL

---

```
1: Initialize:  $k = 1$ ,  $\bar{W}_h^0(s, a) = H$ ,  $\pi_h^1(s) = a_1$  for all  $s \in \mathcal{S}$ ,  $a \in \mathcal{A}$  and  $h \in [H]$ .
2: while  $\bar{W}_h^{k-1}(s_1, \pi_1^k(s_1)) \leq \varepsilon/2$  do
3:   for  $h = 1, \dots, H$  do
4:     Observe state  $s_h^k$ . Take action  $a_h^k = \pi_h^k(s_h^k)$  and observe  $s_{h+1}^k$ .
      $\triangleright$  Update the optimistic estimate of cost.
5:     Calculate empirical cost  $\hat{c}(s)$  and  $\bar{c}(s) = \hat{c}(s) - \beta(N^k(s), \delta)$  for all  $s \in \mathcal{S}$ .
      $\triangleright$  Update the estimates of  $\Delta_h(s, a)$  and  $\mathcal{U}_h$ .
6:     Update  $\Delta_h^k(s_h^k, a_h^k) = \Delta_h^{k-1}(s_h^k, a_h^k) \cup \{s_{h+1}^k\}$  and  $\{\mathcal{U}_h^k\}_{h \in [H]}$  by Eq. (2).
7:     Calculate  $A_h^{k, safe}(s)$  for all  $s \in \mathcal{S}$  by Eq. (3).
8:   end for
9:   Update  $N_h^k(s, a, s')$ ,  $N_h^k(s, a)$  and  $\hat{\mathbb{P}}_h^k$  for all  $s \in \mathcal{S}$ ,  $a \in \mathcal{A}$ ,  $s' \in \mathcal{S}$  and  $h \in [H]$ .
10:  Compute  $\bar{W}^k$  according to Eq. (5).
   $\triangleright$  Calculate the greedy policy.
11:  for  $h \in [H]$ ,  $s \in \mathcal{S}$  do
12:    if  $A_h^{k, safe}(s) \neq \emptyset$  then
13:       $\pi_h^{k+1}(s) = \arg \max_{a \in A_h^{k, safe}(s)} \bar{W}_h^k(s, a)$ .
14:    else
15:       $\pi_h^{k+1}(s) = \arg \max_{a \in \mathcal{A}} \bar{W}_h^k(s, a)$ .
16:    end if
17:  end for
18:  Set  $k = k + 1$ .
19: end while
20: return the tuple  $(\hat{\mathbb{P}}_h^{k-1}, \Delta_h^{k-1}(s, a), \mathcal{U}_H^{k-1})$ .
```

---

**Theorem 6.2.** Let  $\gamma(n, \delta) = 2(\log(2SAH/\delta) + (S-1)\log(e(1+n/(S-1))))$ , Algorithm 2 is a  $(\varepsilon, \delta)$ -PAC algorithm with sample complexity at most<sup>2</sup>

$$K = \tilde{\mathcal{O}} \left( \left( \frac{S^2AH^2}{\varepsilon} + \frac{H^4SA}{\varepsilon^2} \right) \left( \log \left( \frac{1}{\delta} \right) + S \right) \right),$$

The step-wise violation of Algorithm 2 during exploration is  $C(K) = \tilde{\mathcal{O}}(S^2AH^2 + \sqrt{ST})$ .

Compared to previous work (Kaufmann et al., 2021) with an  $\tilde{\mathcal{O}}((H^4SA/\varepsilon^2)(\log(1/\delta) + S))$  sample complexity, our result has an additional term  $\tilde{\mathcal{O}}((S^2AH^2/\varepsilon)(\log(1/\delta) + S))$ . This term is due to the additional safety requirement for the final output policy, which was not considered in previous RFE algorithms (Kaufmann et al., 2021; Ménard et al., 2021). When  $\varepsilon$  and  $\delta$  are sufficiently small, the leading term is  $\tilde{\mathcal{O}}((H^4SA/\varepsilon^2)\log(1/\delta))$ , which implies that our algorithm satisfies the safety constraint without suffering additional regret.<sup>3</sup>

### 6.3 Analysis for Algorithm SRF-UCRL

For the analysis of step-wise violation (Eq. (1)), similar to algorithm SUCBVI, algorithm SRF-UCRL estimates the next state set  $\Delta_h(s, a)$  and potentially unsafe state set  $\mathcal{U}_h$ , which guarantees a  $\tilde{\mathcal{O}}(\sqrt{ST})$  step-wise violation. Now we give a proof sketch for the  $\varepsilon$ -safe property of output policy (Eq. (4)).

First, if  $\pi$  is a feasible policy for  $(\hat{\mathbb{P}}^k, \Delta^k(s, a), \{\mathcal{U}_h^k\}_{h \in [H]})$  and  $s_{h+1} \in \Delta_h^k(s_h, a_h)$  for all  $h \in [H]$ , the agent who follows policy  $\pi$  will only visit the estimated safe states. Since each estimated safe state

---

<sup>2</sup>Here  $\tilde{\mathcal{O}}(\cdot)$  ignores all  $\log S$ ,  $\log A$ ,  $\log H$ ,  $\log(1/\varepsilon)$  and  $\log(\log(1/\delta))$  terms.

<sup>3</sup>Ménard et al. (2021) improve the result by a factor  $H$  and replace  $\log(1/\delta) + S$  by  $\log(1/\delta)$  via the Bernstein-type inequality. While they do not consider the safety constraints, we believe that similar improvement can also be applied to our framework, without significant changes in our analysis. However, since our paper mainly focuses on tackling the safety constraint for reward-free exploration, we use the Hoeffding-type inequality to keep the succinctness of our statements.

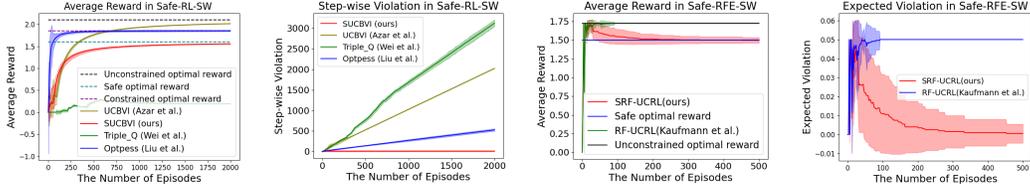


Figure 1: Experimental results for Safe-RL-SW and Safe-RFE-SW. The left two figures show the average rewards and step-wise violations of algorithms SUCBVI, UCBVI (Azar et al., 2017), OptCMDP-bonus (Efroni et al., 2020b), Triple-Q (Wei et al., 2022) and Optpess (Liu et al., 2021a). The right two figures show the reward and expected violation of the policies outputted by algorithms SRF-UCRL and RF-UCRL (Kaufmann et al., 2021).

only suffers a  $\mathcal{O}(1/\sqrt{N_t(s_h, \pi_h(s_h))})$  violation, the violation led by this situation is bounded by  $\overline{W}_1^k(s, \pi_1(s_1))$ . Next, we bound the probability that  $s_{h+1} \notin \Delta_h^K(s_h, a_h)$  for some step  $h \in [H]$ . For any state-action pair  $(s, a)$ , if there is a probability  $\mathbb{P}(s' | s, a) \geq \tilde{\mathcal{O}}(\log(S/\delta)/N_h^K(s, a))$  that the agent transitions to next state  $s'$  from  $(s, a)$  at step  $h$ , the state  $s'$  is put into  $\Delta_h^K(s, a)$  with probability at least  $1 - \delta/S$ . Then, we can expect that all such states are put into our estimated next state set  $\Delta_h^K(s, a)$ . Thus, the probability that  $s_{h+1} \notin \Delta_h^K(s_h, a_h)$  is no larger than  $\tilde{\mathcal{O}}(S \log(S/\delta)/N_h^K(s, a))$  by a union bound over all possible next states  $s'$ . Based on this argument,  $\overline{W}_h^k(s_h, a_h)$  is an upper bound for the total probability that  $s_{h+1} \notin \Delta_h^K(s_h, a_h)$  for some step  $h$ . This will lead to additional  $\overline{W}_1^K(s_1, \pi_1(s_1))$  expected violation. Hence the expected violation of output policy is upper bounded by  $2\overline{W}_1^K(s_1, \pi_1^K(s_1)) \leq \varepsilon$ . The complete proof is provided in Appendix A.

## 7 Experiments

In this section, we provide experiments for Safe-RL-SW and Safe-RFE-SW to validate our theoretical results. For Safe-RL-SW, we compare our algorithm SUCBVI with a classical RL algorithm UCBVI (Azar et al., 2017) and three state-of-the-art CMDP algorithms OptCMDP-bonus (Efroni et al., 2020b), Optpess (Liu et al., 2021a) and Triple-Q (Wei et al., 2022). For Safe-RFE-SW, we report the average reward in each episode and cumulative step-wise violation. For Safe-RFE-SW, we compare our algorithm SRF-UCRL with a state-of-the-art RFE algorithm RF-UCRL (Kaufmann et al., 2021). We do not plot the regret because unconstrained MDP or CMDP algorithms do not guarantee step-wise violation. Applying them to the step-wise constrained setting can lead to negative or large regret and large violations, making the results meaningless. Detailed experiment setup is in Appendix E.

As shown in Figure 1, the rewards of SUCBVI and SRF-UCRL converge to the optimal rewards under safety constraints (denoted by "safe optimal reward"). In contrast, the rewards of UCBVI and UCRL converge to the optimal rewards without safety constraints (denoted by "unconstrained optimal reward"), and those of CMDP algorithms converge to a policy with low expected violation (denoted by "constrained optimal reward"). For Safe-RFE-SW, the expected violation of output policy of SRF-UCRL converges to zero while that of RF-UCRL does not. This corroborates the ability of SRF-UCRL in finding policies that simultaneously achieve safety and high rewards.

## 8 Conclusion

In this paper, we investigate a novel safe reinforcement learning problem with step-wise safety constraints. We first provide an algorithmic framework SUCBVI to achieve both an  $\tilde{\mathcal{O}}(\sqrt{H^3 SAT})$  regret and an  $\tilde{\mathcal{O}}(\sqrt{ST})$  step-wise or an  $\tilde{\mathcal{O}}(S/C_{\text{gap}} + S^2 AH^2)$  gap-dependent bounded violation that is independent of  $T$ . Then, we provide two lower bounds to validate the optimality of SUCBVI in both violation and regret in terms of  $S$  and  $T$ . Further, we extend our framework to the safe RFE with a step-wise violation and provide an algorithm SRF-UCRL that identifies a near-optimal safe policy given any reward function  $r$  and guarantees an  $\tilde{\mathcal{O}}(\sqrt{ST})$  violation during exploration.

## Acknowledgements

This work is supported by the Technology and Innovation Major Project of the Ministry of Science and Technology of China under Grant 2020AAA0108400 and 2020AAA0108403 and the Tsinghua Precision Medicine Foundation 10001020109.

## References

- Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *International conference on machine learning*, pp. 22–31. PMLR, 2017.
- M Mehdi Afsar, Trafford Crump, and Behrouz Far. Reinforcement learning based recommender systems: A survey. *ACM Computing Surveys*, 55(7):1–38, 2022.
- Mohammed Alshiekh, Roderick Bloem, Rüdiger Ehlers, Bettina Könighofer, Scott Niekum, and Ufuk Topcu. Safe reinforcement learning via shielding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Sanae Amani, Christos Thrampoulidis, and Lin Yang. Safe reinforcement learning with linear function approximation. In *International Conference on Machine Learning*, pp. 243–253. PMLR, 2021.
- Peter Auer, Thomas Jaksch, and Ronald Ortner. Near-optimal regret bounds for reinforcement learning. *Advances in neural information processing systems*, 21, 2008.
- Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, pp. 263–272. PMLR, 2017.
- Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. Safe model-based reinforcement learning with stability guarantees. *Advances in neural information processing systems*, 30, 2017.
- Archana Bura, Aria HasanzadeZonuzi, Dileep Kalathil, Srinivas Shakkottai, and Jean-Francois Chamberland. Safe exploration for constrained reinforcement learning with provable guarantees. *arXiv preprint arXiv:2112.00885*, 2021.
- Yinlam Chow, Ofir Nachum, Edgar Duenez-Guzman, and Mohammad Ghavamzadeh. A lyapunov-based approach to safe reinforcement learning. *Advances in neural information processing systems*, 31, 2018.
- Gal Dalal, Krishnamurthy Dvijotham, Matej Vecerik, Todd Hester, Cosmin Paduraru, and Yuval Tassa. Safe exploration in continuous action spaces. *arXiv preprint arXiv:1801.08757*, 2018.
- Dongsheng Ding, Kaiqing Zhang, Tamer Basar, and Mihailo Jovanovic. Natural policy gradient primal-dual method for constrained markov decision processes. *Advances in Neural Information Processing Systems*, 33:8378–8390, 2020.
- Dongsheng Ding, Xiaohan Wei, Zhuoran Yang, Zhaoran Wang, and Mihailo Jovanovic. Provably efficient safe exploration via primal-dual policy optimization. In *International Conference on Artificial Intelligence and Statistics*, pp. 3304–3312. PMLR, 2021.
- Yonathan Efroni, Shie Mannor, and Matteo Pirota. Exploration-exploitation in constrained mdps. *arXiv preprint arXiv:2003.02189*, 2020a.
- Yonathan Efroni, Shie Mannor, and Matteo Pirota. Exploration-exploitation in constrained mdps. *arXiv preprint arXiv:2003.02189*, 2020b.
- Wassily Hoeffding. Probability inequalities for sums of bounded random variables. In *The collected works of Wassily Hoeffding*, pp. 409–426. Springer, 1994.
- Ruiquan Huang, Jing Yang, and Yingbin Liang. Safe exploration incurs nearly no additional sample complexity for reward-free rl. *arXiv preprint arXiv:2206.14057*, 2022.

- Chi Jin, Akshay Krishnamurthy, Max Simchowitz, and Tiancheng Yu. Reward-free exploration for reinforcement learning. In *International Conference on Machine Learning*, pp. 4870–4879. PMLR, 2020.
- Krishna C Kalagarla, Rahul Jain, and Pierluigi Nuzzo. A sample-efficient algorithm for episodic finite-horizon mdp with constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 8030–8037, 2021.
- Emilie Kaufmann, Pierre Ménard, Omar Darwiche Domingues, Anders Jonsson, Edouard Leurent, and Michal Valko. Adaptive reward-free exploration. In *Algorithmic Learning Theory*, pp. 865–891. PMLR, 2021.
- Marc Lanctot, Edward Lockhart, Jean-Baptiste Lespiau, Vinicius Zambaldi, Satyaki Upadhyay, Julien Pérolat, Sriram Srinivasan, Finbarr Timbers, Karl Tuyls, Shayegan Omidshafiei, et al. Openspiel: A framework for reinforcement learning in games. *arXiv preprint arXiv:1908.09453*, 2019.
- Hoang Le, Cameron Voloshin, and Yisong Yue. Batch policy learning under constraints. In *International Conference on Machine Learning*, pp. 3703–3712. PMLR, 2019.
- Tao Liu, Ruida Zhou, Dileep Kalathil, Panganamala Kumar, and Chao Tian. Learning policies with zero or bounded constraint violation for constrained mdps. *Advances in Neural Information Processing Systems*, 34:17183–17193, 2021a.
- Tao Liu, Ruida Zhou, Dileep Kalathil, Panganamala Kumar, and Chao Tian. Learning policies with zero or bounded constraint violation for constrained mdps. *Advances in Neural Information Processing Systems*, 34:17183–17193, 2021b.
- Yongshuai Liu, Jiaxin Ding, and Xin Liu. Ipo: Interior-point policy optimization under constraints. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 4940–4947, 2020.
- Pierre Ménard, Omar Darwiche Domingues, Anders Jonsson, Emilie Kaufmann, Edouard Leurent, and Michal Valko. Fast active learning for pure exploration in reinforcement learning. In *International Conference on Machine Learning*, pp. 7599–7608. PMLR, 2021.
- Sobhan Miryoosefi and Chi Jin. A simple reward-free approach to constrained reinforcement learning. In *International Conference on Machine Learning*, pp. 15666–15698. PMLR, 2022.
- Ian Osband and Benjamin Van Roy. On lower bounds for regret in reinforcement learning. *arXiv preprint arXiv:1608.02732*, 2016.
- Shuang Qiu, Xiaohan Wei, Zhuoran Yang, Jieping Ye, and Zhaoran Wang. Upper confidence primal-dual reinforcement learning for cmdp with adversarial loss. *Advances in Neural Information Processing Systems*, 33:15277–15287, 2020.
- Ming Shi, Yingbin Liang, and Ness Shroff. A near-optimal algorithm for safe reinforcement learning under instantaneous hard constraints. *arXiv preprint arXiv:2302.04375*, 2023.
- Thiago D Simão, Nils Jansen, and Matthijs TJ Spaan. Always safe: Reinforcement learning without safety constraint violations during training. 2021.
- Max Simchowitz and Kevin G Jamieson. Non-asymptotic gap-dependent regret bounds for tabular mdps. *Advances in Neural Information Processing Systems*, 32, 2019.
- Rahul Singh, Abhishek Gupta, and Ness B Shroff. Learning in markov decision processes under constraints. *arXiv preprint arXiv:2002.12435*, 2020.
- Aivar Sootla, Alexander Cowen-Rivers, Jun Wang, and Haitham Bou Ammar. Enhancing safe exploration using safety state augmentation. *Advances in Neural Information Processing Systems*, 35:34464–34477, 2022.
- Adam Stooke, Joshua Achiam, and Pieter Abbeel. Responsive safety in reinforcement learning by pid lagrangian methods. In *International Conference on Machine Learning*, pp. 9133–9143. PMLR, 2020.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

- Chen Tessler, Daniel J Mankowitz, and Shie Mannor. Reward constrained policy optimization. *arXiv preprint arXiv:1805.11074*, 2018.
- Garrett Thomas, Yuping Luo, and Tengyu Ma. Safe reinforcement learning by imagining the near future. *Advances in Neural Information Processing Systems*, 34:13859–13869, 2021.
- Matteo Turchetta, Felix Berkenkamp, and Andreas Krause. Safe exploration in finite markov decision processes with gaussian processes. *Advances in neural information processing systems*, 29, 2016.
- Matteo Turchetta, Andrey Kolobov, Shital Shah, Andreas Krause, and Alekh Agarwal. Safe reinforcement learning via curriculum induction. *Advances in Neural Information Processing Systems*, 33:12151–12162, 2020.
- Eiji Uchibe and Kenji Doya. Constrained reinforcement learning from intrinsic and extrinsic rewards. In *2007 IEEE 6th International Conference on Development and Learning*, pp. 163–168. IEEE, 2007.
- Akifumi Wachi and Yanan Sui. Safe reinforcement learning in constrained markov decision processes. In *International Conference on Machine Learning*, pp. 9797–9806. PMLR, 2020.
- Akifumi Wachi, Yanan Sui, Yisong Yue, and Masahiro Ono. Safe exploration and optimization of constrained mdps using gaussian processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Akifumi Wachi, Yunyue Wei, and Yanan Sui. Safe policy optimization with local generalized linear function approximations. *Advances in Neural Information Processing Systems*, 34:20759–20771, 2021.
- Yixuan Wang, Simon Sinong Zhan, Ruochen Jiao, Zhilu Wang, Wanxin Jin, Zhuoran Yang, Zhaoran Wang, Chao Huang, and Qi Zhu. Enforcing hard constraints with soft barriers: Safe reinforcement learning in unknown stochastic environments. *arXiv preprint arXiv:2209.15090*, 2022.
- Yixuan Wang, Simon Sinong Zhan, Ruochen Jiao, Zhilu Wang, Wanxin Jin, Zhuoran Yang, Zhaoran Wang, Chao Huang, and Qi Zhu. Enforcing hard constraints with soft barriers: Safe reinforcement learning in unknown stochastic environments. In *International Conference on Machine Learning*, pp. 36593–36604. PMLR, 2023.
- Honghao Wei, Xin Liu, and Lei Ying. Triple-q: A model-free algorithm for constrained reinforcement learning with sublinear regret and zero constraint violation. In *International Conference on Artificial Intelligence and Statistics*, pp. 3274–3307. PMLR, 2022.
- Dongjie Yu, Haitong Ma, Shengbo Li, and Jianyu Chen. Reachability constrained reinforcement learning. In *International Conference on Machine Learning*, pp. 25636–25655. PMLR, 2022.
- Ming Yu, Zhuoran Yang, Mladen Kolar, and Zhaoran Wang. Convergent policy optimization for safe reinforcement learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- Tiancheng Yu, Yi Tian, Jingzhao Zhang, and Suvrit Sra. Provably efficient algorithms for multi-objective competitive rl. In *International Conference on Machine Learning*, pp. 12167–12176. PMLR, 2021.
- Wenshuai Zhao, Jorge Peña Queralta, and Tomi Westerlund. Sim-to-real transfer in deep reinforcement learning for robotics: a survey. In *2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 737–744. IEEE, 2020.

# Appendix

<b>A</b>	<b>Proofs of Main Results</b>	<b>14</b>
A.1	Proof of Theorem 4.2 . . . . .	14
A.2	Proof of Theorem 6.2 . . . . .	18
A.3	Proof of Theorem 5.1 . . . . .	26
A.4	Proof of Theorem 5.2 . . . . .	27
<b>B</b>	<b>Safe Markov Games</b>	<b>29</b>
B.1	Definitions . . . . .	29
B.2	Algorithms and Results . . . . .	30
B.3	Proofs . . . . .	32
<b>C</b>	<b>Technical Lemmas</b>	<b>34</b>
C.1	Lemma C.1 . . . . .	34
C.2	Lemma C.2 . . . . .	35
<b>D</b>	<b>Detailed Analysis of Assumption 4.1</b>	<b>35</b>
<b>E</b>	<b>Experiment Setup</b>	<b>35</b>

## A Proofs of Main Results

### A.1 Proof of Theorem 4.2

We start the proof with the confidence bound of optimistic estimation:

**Lemma A.1.** *With probability at least  $1 - \delta$ , the empirical optimistic estimation of cost function  $\bar{c}^k(s) = \hat{c}^k(s) - \sqrt{\frac{2}{N^k(s)} \log(SK/\delta)}$  are always less than the true cost  $c(s)$  for any episode  $k \in [K]$ .*

*In other words, let the event be*

$$\mathcal{E}_1 = \{\forall k \in [K] \text{ and } s \in \mathcal{S}, \bar{c}^k(s) \leq c(s)\}, \quad (6)$$

*then  $\Pr(\mathcal{E}_1^c) \leq \delta$ .*

*Proof.* Since we observe  $c(s) + \varepsilon$  with 1-subgaussian noise  $\varepsilon$ , by Hoeffding's inequality [Hoeffding \(1994\)](#), for a fixed episode  $k$  and state  $s$ , the statement

$$|\hat{c}^k(s) - c(s)| \leq \sqrt{\frac{2}{N^k(s)} \log\left(\frac{SK}{\delta}\right)}$$

holds with probability at least  $1 - \frac{\delta}{SK}$ . Taking a union bound on all  $s \in \mathcal{S}$  and  $k \in [K]$ , we have

$$c(s) \geq \hat{c}^k(s) - \sqrt{\frac{2}{N^k(s)} \log\left(\frac{SK}{\delta}\right)} = \bar{c}^k(s),$$

and the Lemma A.1 has been proved. □

**Lemma A.2.** Let the event  $\mathcal{E}_2$  and  $\mathcal{E}_3$  are

$$\begin{aligned}\mathcal{E}_2 &= \left\{ \forall s, a, s', k, h, |\hat{\mathbb{P}}_h(s' | s, a) - \mathbb{P}_h(s' | s, a)| \right. \\ &\quad \left. \leq \sqrt{\frac{2\mathbb{P}_h(s' | s, a)}{N_h^k(s, a)} \log\left(\frac{2S^2 AHK}{\delta}\right)} + \frac{2}{3N_h^k(s, a)} \log\left(\frac{2S^2 AHK}{\delta}\right) \right\} \\ \mathcal{E}_3 &= \left\{ \sum_{k=1}^K \sum_{h=1}^H \left( P_h(V_{h+1}^k - V_{h+1}^{\pi^k})(s_h^k, a_h^k) - (V_{h+1}^k - V_{h+1}^{\pi^k})(s_{h+1}^k) \right) \leq \sqrt{\frac{H^3 K}{2} \log(1/\delta)} \right\}\end{aligned}$$

Then  $\Pr(\mathcal{E}_2^c) \leq \delta$ ,  $\Pr(\mathcal{E}_3^c) \leq \delta$ .

*Proof.* This lemma is very standard in analysis of UCBVI algorithm. The  $\Pr(\mathcal{E}_2^c)$  can be derived by Bernstein's inequality, and  $\Pr(\mathcal{E}_3^c) \leq \delta$  can be derived by Azuma-hoeffding inequality.  $\square$

The following lemma shows the optimistic property we discussed in the paper.

**Lemma A.3.** Under the event  $\mathcal{E}_1$ ,  $\mathcal{U}_h^k \subseteq \mathcal{U}_h$  and  $A_h^{safe}(s) \subseteq A_h^{k,safe}(s)$  for all  $s \in \mathcal{U}_h^k$  and  $h \in [H]$ .

*Proof.* Under the event  $\mathcal{E}_1$ , for all  $s \in \mathcal{U}_H^k$ ,  $c(s) \geq \bar{c}(s) > \tau$ . So  $s \in \mathcal{U}_H^*$ . Now from our definition of  $\Delta_h^K(s, a)$ , we can easily have  $\Delta_h^k(s, a) \subseteq \Delta_h^K(s, a)$  for all step  $h \in [H]$  because we add the possible next state one by one into estimated set  $\Delta_h^K(s, a)$ .

Now we prove the lemma by induction. When  $h = H$ ,  $\mathcal{U}_H^k \subseteq \mathcal{U}_H^*$ . Assume that  $\mathcal{U}_{h+1}^k \subseteq \mathcal{U}_{h+1}^*$ , then

$$\begin{aligned}A_h^{safe}(s) &= \{a \in \mathcal{A} \mid \Delta_h(s, a) \cap \mathcal{U}_{h+1}^* = \emptyset\} \\ &\subseteq \{a \in \mathcal{A} \mid \Delta_h^K(s, a) \cap \mathcal{U}_{h+1}^* = \emptyset\} \\ &\subseteq \{a \in \mathcal{A} \mid \Delta_h^K(s, a) \cap \mathcal{U}_{h+1}^k = \emptyset\} \\ &= A_h^{k,safe}(s),\end{aligned}$$

and

$$\begin{aligned}\mathcal{U}_h^k &= \mathcal{U}_{h+1}^k \cup \{s \mid \exists a \in \mathcal{A}, \Delta_h^K(s, a) \cap \mathcal{U}_{h+1}^k = \emptyset\} \\ &\subseteq \mathcal{U}_{h+1}^* \cup \{s \mid \exists a \in \mathcal{A}, \Delta_h(s, a) \cap \mathcal{U}_{h+1}^* = \emptyset\} \\ &\subseteq \mathcal{U}_h.\end{aligned}$$

Hence we complete the proof by induction.  $\square$

Similar to UCBVI algorithm, we prove that  $Q_h^k(s, a)$  and  $V_h^k(s)$  are always greater than  $Q_h^*(s)$  and  $V_h^*(s)$ .

**Lemma A.4.** With probability at least  $1 - \delta$ , choosing  $b(n, \delta) = 7H\sqrt{\frac{\ln(5SAT/\delta)}{n}}$ , we will have

$$Q_h^k(s, a) \geq Q_h^*(s, a), \quad V_h^k(s) \geq V_h^*(s),$$

for any state  $s \in \mathcal{S}$ , action  $a \in \mathcal{A}$ , episode  $k$  and the step  $h \in [H]$ .

*Proof.* We proof this lemma by induction. First, consider  $h = H + 1$ , then  $Q_{H+1}^*(s, a) = Q_{H+1}^k(s, a) = 0$  and  $V_{H+1}^*(s) = V_{H+1}^k(s) = 0$  for all  $k \in [K]$ . Then our statements hold for  $h = H + 1$ . Now fixed an episode  $k$  and assume this statement holds for  $h + 1$ . If  $Q_h^k(s, a) = H$ , our proof is completed. Then assume  $Q_h^k(s, a) \leq H$  and

$$\begin{aligned}Q_h^k(s, a) - Q_h^*(s, a) &= (\hat{\mathbb{P}}_h^k)(V_{h+1} - V^*)(s, a) + (\hat{\mathbb{P}}_h^k - \mathbb{P}_h)(V^*)(s, a) + b(N_h^k(s, a)) \\ &\geq (\hat{\mathbb{P}}_h^k - \mathbb{P}_h)(V^*)(s, a) + b(N_h^k(s, a)) \\ &\geq 0\end{aligned}$$

where the first inequality is the induction hypothesis, and the second inequality follows naturally by Chernoff-Hoeffding's inequality. Then for  $s \in \mathcal{U}_h^k$

$$V_h^k(s) = \max_{a \in \mathcal{A}} Q_h^k(s, a) \geq \max_{a \in \mathcal{A}} Q_h^*(s, a) \geq V_h^*(s).$$

For  $s \notin \mathcal{U}_h^k$ , since we know  $\mathcal{U}_h \subseteq \mathcal{U}_h^k$ ,  $s \notin \mathcal{U}_h$ . Then

$$\begin{aligned} V_h^k(s) &= \max_{a \in A_h^{k, safe}(s)} Q_h^k(s, a) \geq \max_{a \in A_h^{k, safe}(s)} Q_h^*(s, a) \\ &\geq \max_{a \in A_h^{safe}(s)} Q_h^*(s, a) \\ &= V_h^*(s). \end{aligned} \quad (7)$$

where Eq. (7) is because  $A_h^{safe}(s) \subseteq A_h^{k, safe}(s)$  by our optimistic exploration mechanism.  $\square$

**Regret** Now we can start to prove our main theorem.

*Proof.* Having the previous analysis, our proof is similar to UCBVI algorithm. Note that the regret can be bounded by UCB:

$$R(K) = \sum_{k=1}^K (V_1^*(s_1) - V_1^{\pi^k}(s_1)) \leq \sum_{k=1}^K (V_1^k(s_1) - V_1^{\pi^k}(s_1)).$$

Because the action taking by  $\pi^k$  is equal to action taking at episode  $k$ , we will have

$$V_h^k(s_h^k) - V_h^{\pi^k}(s_h^k) = (Q_h^k - Q_h^{\pi^k})(s_h^k, a_h^k),$$

where  $a_h^k = \pi_h^k(s_h^k)$ . Then

$$\begin{aligned} V_h^k(s_h^k) - V_h^{\pi^k}(s_h^k) &= (Q_h^k - Q_h^{\pi^k})(s_h^k, a_h^k) \\ &= (\hat{\mathbb{P}}_h^k(V_{h+1}^k) - \mathbb{P}_h(V_{h+1}^{\pi^k}))(s_h^k, a_h^k) + b(N_h^k(s_h^k, a_h^k), \delta) \\ &= (\hat{\mathbb{P}}_h^k - \mathbb{P}_h)(V_{h+1}^*)(s_h^k, a_h^k) + (\hat{\mathbb{P}}_h^k - \mathbb{P}_h)(V_{h+1}^k - V_{h+1}^*)(s_h^k, a_h^k) \\ &\quad + P_h(V_{h+1}^k - V_{h+1}^*)(s_h^k, a_h^k) + b(N_h^k(s_h^k, a_h^k), \delta). \end{aligned} \quad (8)$$

First, by Chernoff-Hoeffding's bound,  $(\hat{\mathbb{P}}_h^k - \mathbb{P}_h)(V_{h+1}^*)(s_h^k, a_h^k) \leq b(N_h^k(s_h^k, a_h^k), \delta)$ . Second, by the standard analysis of UCBVI (Theorem 1 in UCBVI), under the event  $\mathcal{E}_2$ , we will have

$$\begin{aligned} &(\hat{\mathbb{P}}_h^k - \mathbb{P}_h)(V_{h+1}^k - V_{h+1}^*)(s_h^k, a_h^k) \\ &\leq \sum_{s' \in \mathcal{S}} \left( \sqrt{\frac{2\mathbb{P}_h(s' | s_h^k, a_h^k)\iota}{N_h^k(s_h^k, a_h^k)}} + \frac{2\iota}{3N_h^k(s_h^k, a_h^k)} \right) (V_{h+1}^k - V_{h+1}^*)(s') \\ &\leq \sum_{s' \in \mathcal{S}} \left( \frac{\mathbb{P}_h(s' | s_h^k, a_h^k)}{H} + \frac{H\iota}{2N_h^k(s_h^k, a_h^k)} + \frac{2\iota}{3N_h^k(s_h^k, a_h^k)} \right) (V_{h+1}^k - V_{h+1}^*)(s') \\ &\leq \frac{1}{H} \mathbb{P}_h(V_{h+1}^k - V_{h+1}^*)(s_h^k, a_h^k) + \frac{2H^2S\iota}{N_h^k(s_h^k, a_h^k)} \\ &\leq \frac{1}{H} \mathbb{P}_h(V_{h+1}^k - V_{h+1}^{\pi^k})(s_h^k, a_h^k) + \frac{2H^2S\iota}{N_h^k(s_h^k, a_h^k)}, \end{aligned}$$

where the first inequality holds under the event  $\mathcal{E}_2$ , and the second inequality is derived by  $\sqrt{ab} \leq \frac{a+b}{2}$ , and the last inequality is because  $V_{h+1}^{\pi^k} \leq V_{h+1}^*$ . Thus combining with Eq. (8), we can get

$$V_h^k(s_h^k) - V_h^{\pi^k}(s_h^k) \leq \left(1 + \frac{1}{H}\right) \mathbb{P}_h(V_{h+1}^k - V_{h+1}^{\pi^k})(s_h^k, a_h^k) + 2b(N_h^k(s_h^k, a_h^k), \delta) + \frac{2H^2S\iota}{N_h^k(s_h^k, a_h^k)}.$$

Denote

$$\begin{aligned} \alpha_h^k &= \mathbb{P}_h(V_{h+1}^k - V_{h+1}^{\pi^k})(s_h^k, a_h^k) - (V_{h+1}^k - V_{h+1}^{\pi^k})(s_h^k, a_h^k). \\ \beta_h^k &= \frac{H^2S\iota}{N_h^k(s_h^k, a_h^k)}. \\ \gamma_h^k &= b(N_h^k(s_h^k, a_h^k), \delta). \end{aligned}$$

we have

$$\begin{aligned} V_h^k(s_h^k) - V_h^{\pi^k}(s_h^k) &\leq \left(1 + \frac{1}{H}\right) (V_{h+1}^k - V_{h+1}^{\pi^k})(s_{h+1}^k) + \left(1 + \frac{1}{H}\right) \alpha_h^k + 2\beta_h^k + 2\gamma_h^k \\ &\leq \left(1 + \frac{1}{H}\right) (V_{h+1}^k - V_{h+1}^{\pi^k})(s_{h+1}^k) + 2\alpha_h^k + 2\beta_h^k + 2\gamma_h^k, \end{aligned}$$

then by recursion,

$$\begin{aligned} V_1^k(s_1) - V_1^{\pi^k}(s_1) &\leq \sum_{h=1}^H \left(1 + \frac{1}{H}\right)^H (2\alpha_h^k + 2\beta_h^k + 2\gamma_h^k) \\ &\leq e \cdot \sum_{h=1}^H (2\alpha_h^k + 2\beta_h^k + 2\gamma_h^k). \end{aligned}$$

Under the event  $\mathcal{E}_3$  with Bernstein's inequality,  $\sum_{k=1}^K \sum_{h=1}^H \alpha_h^k \leq \sqrt{\frac{H^3 K}{2} \log(1/\delta)}$ , and

$$\sum_{k=1}^K \sum_{h=1}^H \beta_h^k = \sum_{k=1}^K \sum_{h=1}^H \frac{H^2 S \iota}{N_h^k(s_h^k, a_h^k)} \leq H^3 S^2 A \iota \log(HK). \quad (9)$$

Also, for  $\gamma_h^k$ , we can get

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \gamma_h^k &= \sum_{k=1}^K \sum_{h=1}^H 7H \sqrt{\frac{\iota}{N_h^k(s_h^k, a_h^k)}} \\ &\leq 7H \sqrt{\iota} \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \sum_{i=1}^{N^k(s,a)} \sqrt{1/i} \\ &\leq 14H \sqrt{\iota} \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} \sqrt{N^k(s,a)} \\ &\leq 14H \sqrt{\iota S A T}. \end{aligned}$$

Then the regret will be bounded by  $O(\sqrt{H^2 S A T})$ .

**Violation** Now we turn to consider the violation during the process.

First, if the agent follows the  $\pi^k$  at episode  $k$ , and  $s_{h+1}^k \in \Delta_h^K(s_h^k, a_h^k)$  for all step  $h \in [H-1]$ , from the definition of  $\{\mathcal{U}_h\}_{1 \leq h \leq H+1}$  and  $\Delta_h^K(s_h^k, a_h^k)$ , all states  $s_1^k, s_2^k, \dots, s_H^k \notin \mathcal{U}_H^k$  are safe. Since  $\mathcal{U}_H^k = \{s \mid \bar{c}(s) > \tau\}$ , for each state  $s \in \{s_1^k, s_2^k, \dots, s_H^k\}$ , we can have

$$c(s) - \tau \leq \bar{c}^k(s) + 2\sqrt{\frac{2}{N^k(s)} \log(SK/\delta)} - \tau \leq 2\sqrt{\frac{2}{N^k(s)} \log(SK/\delta)},$$

where the first inequality from the event  $\mathcal{E}_1$  and the second inequality is because  $s \notin \mathcal{U}_H^k$ . Then

$$\sum_{h=1}^H (c(s_h^k) - \tau)_+ \leq \sum_{h=1}^H 2\sqrt{\frac{2}{N^k(s_h^k)} \log(SK/\delta)}.$$

Then we consider the situation that  $s_{h+1}^k \notin \Delta_h^K(s_h^k, a_h^k)$  for some step  $h \in [H]$  at episode  $k \in [K]$ . In this particular case, we can only bound the violation on this episode by  $H$  because the future state are not in control. Fortunately, in this situation we will add at least one new state  $s_{h+1}^k$  to  $\Delta_h^K(s_h^k, a_h^k)$ , then  $|\Delta_h^{k+1}(s_h^k, a_h^k)| \geq |\Delta_h^k(s_h^k, a_h^k)| + 1$ , and the summation

$$\sum_{(h,s,a) \in [H] \times \mathcal{S} \times \mathcal{A}} |\Delta_h^{k+1}(s,a)| \geq \sum_{(h,s,a) \in [H] \times \mathcal{S} \times \mathcal{A}} |\Delta_h^k(s,a)| + 1.$$

Note that for any episode  $k$ , the summation has a trivial upper bound

$$\sum_{(h,s,a) \in [H] \times \mathcal{S} \times \mathcal{A}} |\Delta_h^k(s, a)| \leq S^2 AH,$$

this situation will also appears for at most  $S^2 AH$  episode and leads at most  $S^2 AH^2$  extra violation.

Hence the total violation can be upper bounded by

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H (c(s_h^k) - \tau)_+ &\leq S^2 AH^2 + \sum_{k=1}^K \sum_{h=1}^H \left( 2\sqrt{\frac{2}{N^k(s_h^k)} \log(SK/\delta)} \wedge 1 \right) \\ &\leq S^2 AH^2 + 2\sqrt{2 \log(SK/\delta)} \sum_{k=1}^K \sum_{h=1}^H \left( \sqrt{\frac{1}{N^k(s_h^k) \vee 1}} \right) \\ &\leq S^2 AH^2 + 2\sqrt{2 \log(SK/\delta)} \sum_{s \in \mathcal{S}} \sum_{i=1}^K \left( \frac{N^{i+1}(s) - N^i(s)}{\sqrt{N^i(s) \vee 1}} \right) \quad (10) \\ &\leq S^2 AH^2 + 2\sqrt{8 \log(SK/\delta)} \sum_{s \in \mathcal{S}} \sqrt{N^K(s)} \\ &\leq S^2 AH^2 + 2\sqrt{8 \log(SK/\delta)} \sqrt{ST} \\ &= \tilde{\mathcal{O}}(S^2 AH^2 + \sqrt{ST}) \end{aligned}$$

with probability at least  $1 - 3\delta$ . The inequality Eq. (10) is derived by Lemma 19 in [Auer et al. \(2008\)](#). By replacing  $\delta$  to  $\delta/3$  we can complete our proof.

Moreover, for the unsafe state set  $\mathcal{U}$ , recall that  $d = \min_{s \in \mathcal{U}} (c(s) - \tau) > 0$ . For the state  $s_h^k$ , if  $s_h^k \notin \mathcal{U}_H^k$ , we have

$$c(s_h^k) - \tau \leq \bar{c}^k(s_h^k) + 2\sqrt{\frac{2}{N^k(s_h^k)} \log(SK/\delta)} - \tau \leq 2\sqrt{\frac{2}{N^k(s_h^k)} \log(SK/\delta)}.$$

Thus we can get

$$N^k(s_h^k) \leq \frac{8}{(c(s) - \tau)^2} \log(SK/\delta).$$

Now denote  $\mathcal{K}$  be the set of episodes  $k$  that  $s_h^{k+1} \in \Delta_h^k(s_h^k, a_h^k)$ . Now for each  $s \notin \mathcal{U}$ , denote  $k_s = \max\{k : \exists h \in [H], s = s_h^k\}$ , then our total violation can be bounded by

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H (c(s_h^k) - \tau)_+ &\leq S^2 AH^2 + \sum_{k \in \mathcal{K}} \sum_{h=1}^H (c(s_h^k) - \tau)_+ \\ &\leq S^2 AH^2 + \sum_{s \notin \mathcal{U}} N^{k_s}(s) (c(s) - \tau)_+ \\ &\leq S^2 AH^2 + \sum_{s \notin \mathcal{U}} \frac{8}{(c(s) - \tau)_+^2} (c(s_h^k) - \tau)_+ \log(SK/\delta) \\ &\leq S^2 AH^2 + \frac{8S}{d} \log(SK/\delta). \end{aligned}$$

□

## A.2 Proof of Theorem 6.2

First, apply the same argument in the Appendix A.1, the violation during the exploration phase can be bounded by

$$\begin{aligned}
\sum_{k=1}^K \sum_{h=1}^H (c(s_h^k) - \tau)_+ &\leq S^2 AH^2 + \sum_{k=1}^K \sum_{h=1}^H \sqrt{\frac{2}{N^k(s_h^k)} \log(SK/\delta)} \\
&\leq S^2 AH^2 + \sqrt{2 \log(SK/\delta)} \sum_{s \in \mathcal{S}} \sum_{i=1}^{N^K(s)} \sqrt{\frac{1}{i}} \\
&\leq S^2 AH^2 + \sqrt{8 \log(SK/\delta)} \sum_{s \in \mathcal{S}} \sqrt{N^K(s)} \\
&= \tilde{\mathcal{O}}(S^2 AH^2 + \sqrt{ST}).
\end{aligned}$$

Now we prove the Algorithm 2 will return  $\varepsilon$ -safe and  $\varepsilon$ -optimal policy. Lemma A.5 shows that  $\overline{W}_h^k(s_1, \pi_1^{k+1}(s_1))$  bounds the estimation error for any reward function  $r$ .

**Lemma A.5.** *For any feasible policy  $\pi \in \Pi^k$ , step  $h$ , episode  $k$  and reward function  $r$ , defining the estimation error as  $\hat{e}_h^{k,\pi}(s, a; r) = |\hat{Q}_h^{k,\pi}(s, a; r) - Q_h^{k,\pi}(s, a; r)|$ . Then, with probability at least  $1 - \delta$ , we have*

$$\hat{e}_h^{k,\pi}(s, a; r) \leq \overline{W}_h^k(s, a).$$

First, we prove the Lemma A.5. We provide an lemma to show a high probability event for concentration.

**Lemma A.6** (Kaufmann et al. (2021)). *Define the event as*

$$\mathcal{E}_4 = \left\{ \forall k, h, s, a, KL(\hat{\mathbb{P}}_h^k(\cdot | (s, a)), \mathbb{P}_h^k(\cdot | (s, a))) \leq \frac{\gamma(N_h^k(s, a), \delta)}{N_h^k(s, a)} \right\},$$

where  $\beta(n, \delta) = 2 \log(SAHK/\delta) + (S-1) \log(e(1+n/(S-1)))$ . Then  $\Pr(\mathcal{E}_4^c) \leq \delta$ .

By Bellman equations,

$$\begin{aligned}
\hat{Q}_h^\pi(s, a; r) - Q_h^\pi(s, a; r) &= \sum_{s'} (\hat{\mathbb{P}}_h(s' | s, a) - \mathbb{P}_h(s' | s, a)) Q_{h+1}^\pi(s', \pi(s'); r) \\
&\quad + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) (\hat{Q}_{h+1}^\pi(s', \pi(s'); r) - Q_{h+1}^\pi(s', \pi(s'); r)).
\end{aligned}$$

By the Pinsker's inequality and event  $\mathcal{E}_4$ ,

$$\begin{aligned}
\hat{e}_h^\pi(s, a; r) &\leq \sum_{s'} |\hat{\mathbb{P}}_h(s' | s, a) - \mathbb{P}_h(s' | s, a)| Q_{h+1}^\pi(s', \pi(s'); r) \\
&\quad + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) |\hat{Q}_{h+1}^\pi(s', \pi(s'); r) - Q_{h+1}^\pi(s', \pi(s'); r)| \\
&\leq H \|\hat{\mathbb{P}}_h(s' | s, a) - \mathbb{P}_h(s' | s, a)\|_1 + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) \hat{e}_{h+1}^\pi(s', \pi_{h+1}(s'); r) \\
&\leq H \sqrt{\frac{\gamma(N_h(s, a), \delta)}{N_h(s, a)}} + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) \hat{e}_{h+1}^\pi(s, \pi_{h+1}(s'); r).
\end{aligned}$$

By induction, if  $\hat{e}_{h+1}^\pi(s, a; r) \leq \overline{W}_{h+1}(s, a)$ , for  $s \notin \mathcal{U}_h^k, a \in A_h^{k, safe}(s)$ , we can have

$$\begin{aligned}
\hat{e}_h^\pi(s, a; r) &\leq \min \left\{ H, H \sqrt{\frac{\gamma(N_h(s, a), \delta)}{N_h(s, a)}} + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) \hat{e}_{h+1}^\pi(s, \pi_{h+1}(s'); r) \right\} \\
&\leq \min \left\{ H, H \sqrt{\frac{\gamma(N_h(s, a), \delta)}{N_h(s, a)}} + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) \max_{b \in A_{h+1}^{k, safe}(s')} \hat{e}_{h+1}^\pi(s', b; r) \right\} \quad (11) \\
&\leq \min \left\{ H, H \sqrt{\frac{\gamma(N_h(s, a), \delta)}{N_h(s, a)}} + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) \max_{b \in A_{h+1}^{k, safe}(s')} \overline{W}_{h+1}(s', b) \right\} \\
&\leq \overline{W}_h(s, a).
\end{aligned}$$

The second inequality is because  $\hat{\mathbb{P}}_h(s' | s, a) = 0$  for all  $s' \in \mathcal{U}_{h+1}$ , and then for  $s' \notin \mathcal{U}_{h+1}$ ,  $\pi_{h+1}(s') \in A_{h+1}^{k, safe}(s')$ . The third inequality holds by induction hypothesis, and the last inequality holds by the definition of  $\overline{W}_h(s, a)$ .

Also, for other state-action pair  $(s, a)$ , we can get

$$\begin{aligned} \hat{e}_h^\pi(s, a; r) &\leq \min \left\{ H, H \sqrt{\frac{\gamma(N_h(s, a), \delta)}{N_h(s, a)}} + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) \hat{e}_{h+1}^\pi(s, \pi_{h+1}(s'); r) \right\} \\ &\leq \min \left\{ H, H \sqrt{\frac{\gamma(N_h(s, a), \delta)}{N_h(s, a)}} + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) \max_{b \in A} \hat{e}_{h+1}^\pi(s', b; r) \right\} \\ &\leq \min \left\{ H, H \sqrt{\frac{\gamma(N_h(s, a), \delta)}{N_h(s, a)}} + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) \max_{b \in A} \overline{W}_{h+1}(s', b) \right\} \\ &\leq \overline{W}_h(s, a). \end{aligned}$$

Hence the Lemma A.5 is true. Now we assume the algorithm terminates at episode  $K + 1$ , for any feasible policy  $\pi \in \Pi^K$ , if we define the value function in model  $\hat{M}$  as  $\hat{V}$ , we will get

$$\begin{aligned} |V_1^\pi(s_1) - \hat{V}_1^\pi(s_1)| &= |Q_1^\pi(s_1, \pi_1(s_1)) - \hat{Q}_1^\pi(s_1, \pi_1(s_1))| \\ &\leq \overline{W}_h^K(s_1, \pi_1(s_1)) \\ &\leq \overline{W}_h^K(s_1, \pi_1^{K+1}(s_1)) \\ &\leq \varepsilon/2. \end{aligned}$$

where the second inequality is because  $\pi^{K+1}$  is the greedy policy with respect to  $\overline{W}_h^K$  and  $\pi$  is in the feasible policy set  $\Pi^K$  at episode  $K$ . Then for any reward function  $r$ , denote the  $\hat{\pi}^* \in \Pi^K$  to be the optimal policy in estimated MDP  $\hat{M}$ . Recall that  $\pi^* \in \Pi^K \subseteq \Pi^*$ , we will have

$$\begin{aligned} V_1^{\pi^*}(s_1) - V_1^{\hat{\pi}^*}(s_1) &\leq \hat{V}_1^{\pi^*}(s_1) - V_1^{\hat{\pi}^*}(s_1) + \varepsilon/2 \\ &\leq \hat{V}_1^{\hat{\pi}^*}(s_1) - V_1^{\hat{\pi}^*}(s_1) + \varepsilon/2 \\ &\leq \varepsilon. \end{aligned}$$

Now we consider the expected violation. We will show that for all feasible policy  $\pi \in \Pi^K$ , the expected violation will be upper bounded by  $2\overline{W}_1(s_1, \pi_1(s_1))$ .

We first consider the event for feasible policy  $\pi$

$$\mathcal{P}_\pi = \{\forall h \in [H - 1], s_{h+1} \in \Delta_h^K(s_h, a_h) \mid s_h, a_h \sim \pi\}.$$

Assume that the event holds, then because  $\pi \in \Pi^K$  is the feasible policy with respect to  $\Delta_h^K(s, a)$ , we have  $s_h \notin \mathcal{U}_h$  for all  $1 \leq h \leq H + 1$ . Then

$$(c(s_h) - \tau)_+ \leq \sqrt{\frac{2}{N^K(s)} \log(SK/\delta)}.$$

and under the event  $\mathcal{P}_\pi$

$$\mathbb{E}_\pi \left[ \sum_{h=1}^H (c(s_h) - \tau)_+ \right] \leq \mathbb{E}_\pi \left[ \sum_{h=1}^H \sqrt{\frac{2}{N^K(s_h)} \log\left(\frac{SK}{\delta}\right)} \right]. \quad (12)$$

Now define

$$\begin{aligned} &\overline{F}_h^\pi(s, a) \\ &= \min \left\{ H, \sqrt{\frac{2}{N(s_h)} \log\left(\frac{SK}{\delta}\right)} + H \sqrt{\frac{\gamma(N_h(s, a), \delta)}{N_h(s, a)}} + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) F_{h+1}^\pi(s', \pi_{h+1}(s')) \right\}. \end{aligned}$$

Since  $\sqrt{\frac{2}{N(s_h)} \log\left(\frac{SK}{\delta}\right)} \leq H \sqrt{\frac{\gamma(N_h(s,a), \delta)}{N_h(s,a)}}$ , we have  $\overline{F}_h^\pi(s, a) \leq \overline{W}_h(s, a)$  by recursion. The proof is similar to A.5.

Also denote

$$F_h^\pi(s, a) = \mathbb{E}_\pi \left[ \sum_{h'=h}^H \sqrt{\frac{2}{N^K(s_h)} \log\left(\frac{SK}{\delta}\right)} \middle| s_h = s, a_h = a \right].$$

Next we prove  $\overline{F}_h^\pi(s, a) \geq F_h^\pi(s, a)$ . To show this fact, we use the induction. When  $h = H + 1$ ,  $F_{H+1}^\pi(s, a) = \overline{F}_{H+1}^\pi(s, a) = 0$ . Now assume that  $\overline{F}_{h+1}^\pi(s, a) \geq F_{h+1}^\pi(s, a)$ , then

$$\begin{aligned} & \overline{F}_h^\pi(s, a) \\ &= \min \left\{ H, \sqrt{\frac{2}{N(s)} \log\left(\frac{SK}{\delta}\right)} + H \sqrt{\frac{\gamma(N_h(s, a), \delta)}{N_h(s, a)}} + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) \overline{F}_{h+1}^\pi(s', \pi_{h+1}(s')) \right\} \\ &\geq \min \left\{ H, \sqrt{\frac{2}{N(s)} \log\left(\frac{SK}{\delta}\right)} + H \sqrt{\frac{\gamma(N_h(s, a), \delta)}{N_h(s, a)}} + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) F_{h+1}^\pi(s', \pi_{h+1}(s')) \right\} \\ &\geq \min \left\{ H, \sqrt{\frac{2}{N(s_h)} \log\left(\frac{SK}{\delta}\right)} + H \|\hat{\mathbb{P}}_h(s' | s, a) - \mathbb{P}_h(s' | s, a)\|_1 \right. \\ &\quad \left. + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) F_{h+1}^\pi(s', \pi_{h+1}(s')) \right\} \\ &\geq \min \left\{ H, \sqrt{\frac{2}{N(s_h)} \log\left(\frac{SK}{\delta}\right)} + \sum_{s'} \mathbb{P}_h(s' | s, a) F_{h+1}^\pi(s', \pi_{h+1}(s')) \right\} \\ &= F_h^\pi(s, a). \end{aligned}$$

Thus by induction, we know  $F_1^\pi(s_1, \pi_1(s_1)) \leq \overline{F}_1^\pi(s_1, \pi_1(s_1)) \leq \overline{W}_1^K(s, \pi_1(s_1))$  because  $\sqrt{\frac{2}{N(s)} \log\left(\frac{SK}{\delta}\right)} \leq H \sqrt{\frac{\gamma(N_h(s,a), \delta)}{N_h(s,a)}}$ .

From now, we have proved that under the condition  $\mathcal{P}_\pi$ , the expected violation can be upper bounded by  $\overline{W}_1(s, \pi_1(s_1))$ . Next we state that  $\Pr(\mathcal{P}_\pi^c)$  is small. Recall that  $\mathcal{P}_\pi = \{\forall h \in [H-1], s_{h+1} \in \Delta_h^K(s_h, a_h) | \pi\}$ , we define

$$G_h^\pi(s, a) = \Pr\{\exists h' \in [h, H-1], s_{h'+1} \notin \Delta_h^K(s_{h'}, a_{h'}) | s_h = s, a_h = a\}$$

and

$$\begin{aligned} & \overline{G}_h^\pi(s, a) \\ &= \min \left\{ H, \frac{SH \log(S^2 HA/\delta)}{N_h(s, a)} + H \sqrt{\frac{\gamma(N_h(s, a), \delta)}{N_h(s, a)}} + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) \overline{G}_{h+1}^\pi(s', \pi_{h+1}(s')) \right\}. \end{aligned}$$

Then by recursion, we can easily show that  $\overline{G}_h^\pi(s, a) \leq \overline{W}_h^K(s, a)$  by  $\log(S^2 HA/\delta) \leq \gamma(N_h(s, a), \delta)$ .

**Lemma A.7.** For all  $h \in [H]$  and feasible policy  $\pi$ , we have  $\overline{G}_h^\pi(s, a) \geq H G_h^\pi(s, a)$ .

*Proof.* Define  $\tilde{G}_h^\pi(s, a)$  as

$$\tilde{G}_h^\pi(s, a) = \min \left\{ H, \frac{SH \log(S^2 HA/\delta)}{N_h(s, a)} + \sum_{s'} \mathbb{P}_h(s' | s, a) \tilde{G}_{h+1}^\pi(s', \pi_{h+1}(s')) \right\}.$$

We will show  $\overline{G}_h^\pi(s, a) \geq \tilde{G}_h^\pi(s, a) \geq HG_h^\pi(s, a)$ . We prove these two inequalities by induction. First, for  $h = H + 1$ , these inequalities holds. Now assume that they hold for  $h + 1$ , then

$$\begin{aligned}
& \overline{G}_h^\pi(s, a) \\
&= \min \left\{ H, \frac{SH \log(S^2 HA/\delta)}{N_h(s, a)} + H \sqrt{\frac{\gamma(N_h(s, a), \delta)}{N_h(s, a)}} + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) \overline{G}_{h+1}^\pi(s', \pi_{h+1}(s')) \right\} \\
&\geq \min \left\{ H, \frac{SH \log(S^2 HA/\delta)}{N_h(s, a)} + H \sqrt{\frac{\gamma(N_h(s, a), \delta)}{N_h(s, a)}} + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) \tilde{G}_{h+1}^\pi(s', \pi_{h+1}(s')) \right\} \\
&\geq \min \left\{ H, \frac{SH \log(S^2 HA/\delta)}{N_h(s, a)} + H \|\hat{\mathbb{P}}_h(s' | s, a) - \mathbb{P}_h(s' | s, a)\|_1 \right. \\
&\quad \left. + \sum_{s'} \hat{\mathbb{P}}_h(s' | s, a) \tilde{G}_{h+1}^\pi(s', \pi_{h+1}(s')) \right\} \\
&\geq \min \left\{ H, \frac{SH \log(S^2 HA/\delta)}{N_h(s, a)} + \sum_{s'} \mathbb{P}_h(s' | s, a) \tilde{G}_{h+1}^\pi(s', \pi_{h+1}(s')) \right\} \\
&= \tilde{G}_h^\pi(s, a).
\end{aligned}$$

Now to prove that  $\tilde{G}_h^\pi(s, a) \geq HG_h^\pi(s, a)$ , we need the following lemma.

**Lemma A.8.** *Fixed a step  $h \in [H]$  and state-action pair  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , then with high probability at least  $1 - \delta$ , for any  $s' \in \Delta_h(s, a)$  with  $P(s' | s, a) \geq \frac{\log(S/\delta)}{N_h(s, a)}$ ,  $s' \in \Delta_h^K(s, a)$ .*

*Proof.* Since we sample  $s'$  from  $P_h(s' | s, a)$  for  $N_h(s, a)$  times, if  $s' \notin \Delta_h^K(s, a)$ , then  $s'$  are not sampled. The probability that  $s'$  are not sampled will be  $(1 - p)^{N_h(s, a)}$ , where  $p = \mathbb{P}_h(s' | s, a)$ . When  $p \geq \frac{\log(S/\delta)}{N_h(s, a)}$ , the probability will be at most

$$(1 - p)^{N_h(s, a)} \leq \left(1 - \frac{\log(S/\delta)}{N_h(s, a)}\right)^{N_h(s, a)} \leq e^{-\log(S/\delta)} = \frac{\delta}{S}.$$

Taking the union bound for all possible next state  $s'$ , the lemma has been proved.  $\square$

By the Lemma A.8, define the event as

$$\mathcal{E}_5 = \left\{ \forall h, s, a \text{ and } s' \in \Delta_h(s, a) \text{ with } P(s' | s, a) \geq \frac{\log(S^2 HA/\delta)}{N_h(s, a)} : s' \in \Delta_h^K(s, a) \right\}.$$

Then  $\Pr(\mathcal{E}_5^c) \leq \delta$ . Under the event  $\mathcal{E}_5$ , at step  $h \in [H]$  and state-action pair  $(s_h, a_h)$ , we have

$$\mathbb{P}(s_{h+1} \notin \Delta_h^K(s_h, a_h)) \leq \frac{S \log(S^2 HA/\delta)}{N_h(s_h, a_h)}. \quad (13)$$

This inequality is because  $s_{h+1} \notin \Delta_h^K(s_h, a_h)$  only happens when  $P(s' | s, a) \leq \frac{\log(S^2 HA/\delta)}{N_h(s_h, a_h)}$  under the event  $\mathcal{E}_5$ . Now we first decompose the  $G_h^\pi(s, a)$  as

$$\begin{aligned}
H \cdot G_h^\pi(s, a) &= \Pr\{\exists h' \in [h, H - 1], s_{h'+1} \notin \Delta_{h'}^K(s_{h'}, a_{h'}) \mid s_h = s, a_h = a, \pi\} \\
&\leq H \sum_{h'=h}^{H-1} \Pr\{s_{h'+1} \notin \Delta_{h'}^K(s_{h'}, a_{h'}) \mid s_h = s, a_h = a, \pi\} \\
&\leq \mathbb{E}_\pi \left[ \sum_{h'=h}^{H-1} \frac{SH \log(S^2 HA/\delta)}{N_{h'}(s_{h'}, a_{h'})} \right].
\end{aligned}$$

The first inequality is to decompose the probability into the summation over step  $h$ , and the second inequality is derived by Eq. (13). Now denote  $g_h^\pi(s, a) = \min \left\{ H, \mathbb{E}_\pi \left[ \sum_{h'=h}^{H-1} \frac{SH \log(S^2 HA/\delta)}{N_{h'}(s_{h'}, a_{h'})} \right] \right\}$ , then we can easily have  $g_h^\pi(s, a) \leq \tilde{G}_h^\pi(s, a)$  by recursion. In fact, if

$$\begin{aligned}
& g_h^\pi(s, a) \\
&= \min \left\{ H, \mathbb{E}_\pi \left[ \sum_{h'=h}^{H-1} \frac{SH \log(S^2 HA/\delta)}{N_{h'}(s_{h'}, a_{h'})} \middle| (s_h, a_h) = (s, a) \right] \right\} \\
&= \min \left\{ H, \frac{SH \log(S^2 HA/\delta)}{N_h(s, a)} + \mathbb{E}_\pi \left[ \sum_{h'=h+1}^{H-1} \frac{SH \log(S^2 HA/\delta)}{N_{h'}(s_{h'}, a_{h'})} \middle| (s_h, a_h) = (s, a) \right] \right\} \\
&= \min \left\{ H, \frac{SH \log(S^2 HA/\delta)}{N_h(s, a)} \right. \\
&\quad \left. + \min \left\{ H, \mathbb{E}_\pi \left[ \sum_{h'=h+1}^{H-1} \frac{SH \log(S^2 HA/\delta)}{N_{h'}(s_{h'}, a_{h'})} \middle| (s_h, a_h) = (s, a) \right] \right\} \right\} \\
&= \min \left\{ H, \frac{SH \log(S^2 HA/\delta)}{N_h(s, a)} + \sum_{s' \in \mathcal{S}} \mathbb{P}_h(s' | s, a) g_{h+1}^\pi(s', \pi_{h+1}(s')) \right\}
\end{aligned}$$

has the same recursive equation as  $\tilde{G}_h^\pi(s, a)$ . Since  $g_H^\pi(s, a) = 0 \leq \tilde{G}_H^\pi(s, a)$ , by induction we can get  $g_h^\pi(s, a) \leq \tilde{G}_h^\pi(s, a)$ , and then

$$H \cdot G_h^\pi(s, a) \leq g_h^\pi(s, a) \leq \tilde{G}_h^\pi(s, a)$$

holds directly. The Lemma A.7 has been proved.

Now we return to our main theorem. Consider the policy  $\pi \in \Pi^K$  is one of feasible policy at episode  $K$ . For each possible  $\mathcal{T}_H = \{s_1, a_1, \dots, s_H, a_H\}$ , denote  $\Pr_\pi(\mathcal{T}_H)$  be the probability for generating  $\mathcal{T}_H$  following policy  $\pi$  and true transition kernel. If the event  $s_{h+1} \notin \Delta_h^K(s_h, a_h)$  happens, we define the trajectory  $\mathcal{T}_H$  are "unsafe". Denote the set of all "unsafe" trajectories are  $\mathcal{U}$ , we have

$$\begin{aligned}
& \mathbb{E}_\pi \left[ \sum_{h=1}^H (c(s_h) - \tau)_+ \right] \\
&= \sum_{\mathcal{T}_H} \Pr_\pi(\mathcal{T}_H) \left[ \sum_{h=1}^H (c(s_h) - \tau)_+ \right] \\
&\leq \sum_{\mathcal{T}_H \in \mathcal{U}} \Pr_\pi(\mathcal{T}_H) \left[ \sum_{h=1}^H (c(s_h) - \tau)_+ \right] + \sum_{\mathcal{T}_H \notin \mathcal{U}} \Pr_\pi(\mathcal{T}_H) \left[ \sum_{h=1}^H (c(s_h) - \tau)_+ \right] \\
&\leq H \sum_{\mathcal{T}_H \in \mathcal{U}} \Pr_\pi(\mathcal{T}_H) + \sum_{\mathcal{T}_H \notin \mathcal{U}} \Pr_\pi(\mathcal{T}_H) \left[ \sum_{h=1}^H \sqrt{\frac{2}{N^K(s_h)} \log(SK/\delta)} \right] \tag{14}
\end{aligned}$$

$$\begin{aligned}
& \leq H \Pr_\pi\{\mathcal{P}_\pi\} + \sum_{\mathcal{T}_H} \Pr_\pi(\mathcal{T}_H) \left[ \sum_{h=1}^H \sqrt{\frac{2}{N^K(s_h)} \log(SK/\delta)} \right] \\
& \leq H \cdot G_1^\pi(s, a) + F_1^\pi(s, a) \tag{15} \\
& \leq 2\overline{W}_1^K(s_1, \pi_1(s_1)) \\
& \leq 2\overline{W}_1^K(s_1, \pi_1^{K+1}(s_1)) \\
& \leq \varepsilon.
\end{aligned}$$

The Eq. (14) is because  $\mathcal{T}_H \in \mathcal{U}$  implies that the event  $\mathcal{P}_\pi$  happens, and we can apply Eq. (12). The Eq. (15) is derived by the definition of  $F_1^\pi(s, a)$  and  $G_1^\pi(s, a)$ . Until now, we have proved that if the algorithm terminates, all the requirements are satisfied. Now we turn to bound the sample complexity. We first assume that  $K \geq SA$

Define the average  $q_h^k = \sum_{(s,a)} \mathbb{P}_h^{\pi^{k+1}}(s,a) \overline{W}_h^k(s,a)$ . Since the  $\pi^{k+1}$  are greedy policy with respect to  $\overline{W}_h^k(s,a)$ , we have

$$\begin{aligned} & \overline{W}_h^k(s,a) \\ &= \min \left\{ H, M(N_h^k(s,a), \delta) + \sum_{s'} \hat{\mathbb{P}}_h^k(s' | s, a) \overline{W}_{h+1}^k(s', \pi_{h+1}^{t+1}(s')) \right\} \\ &\leq \min \left\{ H, M(N_h^k(s,a), \delta) + \sum_{s'} \mathbb{P}_h^k(s' | s, a) \overline{W}_{h+1}^k(s', \pi_{h+1}^{t+1}(s')) \right. \\ &\quad \left. + H \|\hat{\mathbb{P}}_h^k(s' | s, a) - \mathbb{P}_h^k(s' | s, a)\|_1 \right\} \\ &\leq M'(N_h^k(s,a), \delta) + \sum_{s'} \mathbb{P}_h^k(s' | s, a) \overline{W}_{h+1}^k(s', \pi_{h+1}^{t+1}(s')) \end{aligned}$$

where  $M'(n, \delta) = 3H \left( \sqrt{\frac{2\beta(n,\delta)}{n}} \wedge 1 \right) + SH \left( \frac{\beta(n,\delta)}{n} \wedge 1 \right)$ . Now

$$\begin{aligned} q_h^k &= \sum_{(s,a)} \mathbb{P}_h^{\pi^{k+1}}(s,a) \overline{W}_h^k(s,a) \\ &\leq \sum_{(s,a)} \mathbb{P}_h^{\pi^{k+1}}(s,a) M'(N_h^k(s,a), \delta) + \sum_{(s',s,a)} \mathbb{P}_h^{\pi^{k+1}}(s,a) \mathbb{P}_h^k(s' | s, a) \overline{W}_{h+1}^k(s', \pi_{h+1}^{k+1}(s')) \\ &= \sum_{(s,a)} \mathbb{P}_h^{\pi^{k+1}}(s,a) M'(N_h^k(s,a), \delta) + q_{h+1}^k. \end{aligned}$$

Thus

$$\varepsilon/2 < q_1^k \leq \sum_{h=1}^H \sum_{(s,a)} \mathbb{P}_h^{\pi^{k+1}}(s,a) M'(N_h^k(s,a), \delta), \forall k \in [K-1],$$

where the first equality is because  $q_1^k = \overline{W}_h^k(s_1, \pi^{k+1}(s_1)) > \varepsilon/2$ . Sum over  $k \in [K-1]$ , we can get

$$K\varepsilon/2 \leq \sum_{k=0}^{K-1} \sum_{h=1}^H \sum_{(s,a)} \mathbb{P}_h^{\pi^{k+1}}(s,a) M'(N_h^k(s,a), \delta).$$

Define the pseudo-count as  $\overline{N}_h^k(s,a) = \sum_{i=1}^k \mathbb{P}_h^{\pi^i}(s,a)$  and the event about the pseudo-count like [Kaufmann et al. \(2021\)](#):

$$\mathcal{E}_{cnt} = \left\{ \forall k, h, s, a : N_h^k(s,a) \geq \frac{1}{2} \overline{N}_h^k(s,a) - \log(SAH/\delta) \right\}.$$

Then  $\Pr\{\mathcal{E}_{cnt}\} \leq \delta$ . The proof is provided in [Kaufmann et al. \(2021\)](#). Also, the Lemma 7 in [Kaufmann et al. \(2021\)](#) show that we can change  $N_h^k(s,a)$  to  $\overline{N}_h^k(s,a)$  up to a constant.

**Lemma A.9** ([Kaufmann et al. \(2021\)](#)). *Under the event  $\mathcal{E}_{cnt}$ , for any  $s \in \mathcal{S}, a \in \mathcal{A}$  and  $k$ , if  $\beta(x, \delta)/x$  is non-increasing for  $x \geq 1$  and  $\beta(x, \delta)$  is increasing, we have*

$$\frac{\gamma(N_h^k(s,a), \delta)}{N_h^k(s,a)} \wedge 1 \leq \frac{4\gamma(\overline{N}_h^k(s,a), \delta)}{\overline{N}_h^k(s,a) \vee 1}.$$

The proof of Lemma A.9 can be found in Kaufmann et al. (2021). Then we can have

$$\begin{aligned}
K\varepsilon/2 &\leq \sum_{k=0}^{K-1} \sum_{h=1}^H \sum_{(s,a)} \mathbb{P}_h^{T^{k+1}}(s,a) M'(N_h^k(s,a), \delta) \tag{16} \\
&\leq \underbrace{3H \sum_{h=1}^H \sum_{k=0}^{K-1} \sum_{(s,a)} (\overline{N}_h^{k+1}(s,a) - \overline{N}_h^k(s,a)) \sqrt{\frac{2\gamma(\overline{N}_h^k(s,a), \delta)}{\overline{N}_h^k(s,a)} \wedge 1}}_{(a)} \\
&\quad + \underbrace{SH \sum_{h=1}^H \sum_{k=0}^{K-1} \sum_{(s,a)} (\overline{N}_h^{k+1}(s,a) - \overline{N}_h^k(s,a)) \left( \frac{\gamma(\overline{N}_h^k(s,a), \delta)}{\overline{N}_h^k(s,a)} \wedge 1 \right)}_{(b)}. \tag{17}
\end{aligned}$$

For part (a), we upper bound it by

$$\begin{aligned}
(a) &\leq 3H \sum_{h=1}^H \sum_{k=0}^{K-1} \sum_{(s,a)} (\overline{N}_h^{k+1}(s,a) - \overline{N}_h^k(s,a)) \sqrt{\frac{8\gamma(\overline{N}_h^k(s,a), \delta)}{\overline{N}_h^k(s,a)} \vee 1} \\
&\leq 3\sqrt{2}H \sqrt{\gamma(K, \delta)} \sum_{h=1}^H \sum_{k=0}^{K-1} \sum_{(s,a)} \frac{(\overline{N}_h^{k+1}(s,a) - \overline{N}_h^k(s,a))}{\sqrt{\overline{N}_h^k(s,a)} \vee 1} \tag{18}
\end{aligned}$$

$$\begin{aligned}
&\leq 3(2 + \sqrt{2})H \sqrt{\gamma(K, \delta)} \sum_{h=1}^H \sum_{(s,a)} \sqrt{\overline{N}_h^K(s,a)} \vee 1 \\
&\leq 3(2 + \sqrt{2})H \sqrt{\gamma(K, \delta)} \sum_{h=1}^H \sum_{(s,a)} (1 + \sqrt{\overline{N}_h^K(s,a)}) \\
&\leq 3(2 + \sqrt{2})H^2 \sqrt{\gamma(K, \delta)} (SA + \sqrt{SAK}) \tag{19} \\
&\leq 6(2 + \sqrt{2})H^2 \sqrt{\gamma(K, \delta)} \sqrt{SAK},
\end{aligned}$$

where Eq. (18) is because Lemma 19 in Auer et al. (2008), Eq. (19) is derived by  $\sum_{(s,a)} \overline{N}_h^K(s,a) = K$  and Cauchy's inequality, and the last inequality is because we assume  $K \geq SA$ .

Also for part (b),

$$\begin{aligned}
(b) &\leq SH \sum_{h=1}^H \sum_{k=0}^{K-1} \sum_{(s,a)} (\overline{N}_h^{k+1}(s,a) - \overline{N}_h^k(s,a)) \frac{4\gamma(\overline{N}_h^k(s,a), \delta)}{\overline{N}_h^k(s,a) \vee 1} \\
&\leq SH\gamma(K, \delta) \sum_{h=1}^H \sum_{k=0}^{K-1} \sum_{(s,a)} \frac{(\overline{N}_h^{k+1}(s,a) - \overline{N}_h^k(s,a))}{\overline{N}_h^k(s,a) \vee 1} \\
&\leq SH\gamma(K, \delta) \sum_{h=1}^H \sum_{(s,a)} 4\log(\overline{N}_h^K(s,a) + 1) \tag{20} \\
&\leq 4S^2AH^2\gamma(K, \delta) \log(K + 1),
\end{aligned}$$

where the Eq. (20) is because the Lemma 9 in Ménard et al. (2021). Now from the upper bound of (a) and (b), we can get

$$K\varepsilon/2 \leq 6(2 + \sqrt{2})H^2 \sqrt{\gamma(K, \delta)} \sqrt{SAK} + 4S^2AH^2\gamma(K, \delta) \log(K + 1).$$

Recall that  $\gamma(n, \delta) = 2\log(SAHK/\delta) + (S - 1)\log(e(1 + n/(S - 1))) \geq 1$ , by some technical transformation, performing Lemma C.1 we can get

$$K = \tilde{\mathcal{O}} \left( \left( \frac{H^4SA}{\varepsilon^2} + \frac{S^2AH^2}{\varepsilon} \right) \left( \log \left( \frac{1}{\delta} \right) + S \right) \right),$$

under the event  $\mathcal{E}_{cnt}, \mathcal{E}_1, \mathcal{E}_4, \mathcal{E}_5$ . Hence the theorem holds with probability at least  $1 - 4\delta$ , and we can prove the theorem by replacing  $\delta$  to  $\delta/4$ . □

### A.3 Proof of Theorem 5.1

*Proof.* We construct  $n$  MDP models with the same transition model but different safety cost. For each model, the states consist of a tree with degree  $A$  and layer  $H/3$ . For each state inside the tree (not a leaf), the action  $a_i$  will lead to the  $i$ -th branch of the next state. For leaf nodes, they are absorbing states, and we call them  $s_1, s_2, \dots, s_n$ . Then  $n = A^{H/3} = \frac{S(A-1)+1}{A} \geq \frac{S}{2}$ .

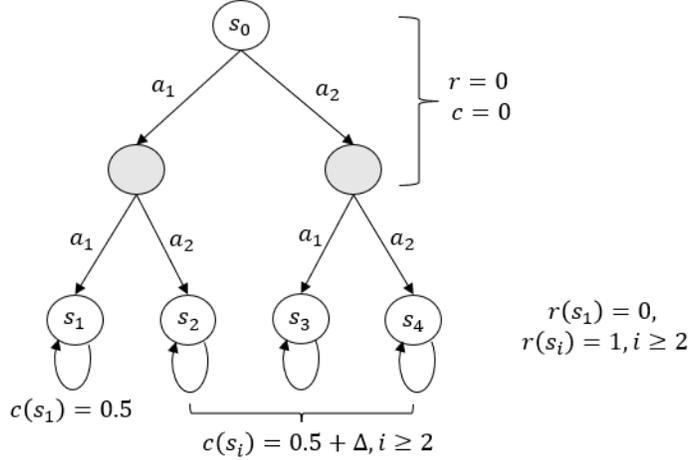


Figure 2: MDP instance  $\tau_1$

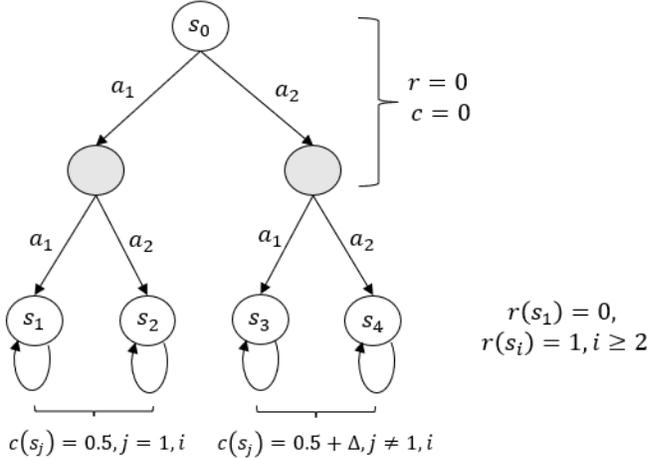


Figure 3: MDP instance  $\tau_i (i = 2)$

Now we first define the instance  $\tau_1$  as follows: For each state inside the tree, the safety cost and reward are all 0. For leaf node, the cost function for state  $s_i$  are  $c(s_i) = 1/2 + \Delta \mathbb{I}\{i \neq 1\}$ . The reward function are  $r(s_i) = \mathbb{I}\{i \neq 1\}$ , which implies that only state  $s_1$  has the reward 0 and other states  $s_2, \dots, s_n$  have reward 1. Now we define the instance  $\tau_j (2 \leq j \leq n)$ . For state inside the tree, the safety cost and reward are equal to the instance  $\tau_1$ . For leaf node, the cost function for state  $s_i$  are  $c(s_i) = 1/2 + \Delta \mathbb{I}\{i \neq 1, j\}$ ,  $r(s_i)$  unchanged. Then for  $\tau_j$ , the state  $s_1$  and  $s_j$  (and initial state  $s_0$ )

are safe. Then the safe optimal policy is to choose  $a_j$  at the first step. Thus the safe optimal reward is  $H$ . The MDP instances are shown in Figure 2 and Figure 3.

Denote  $H' = \frac{2}{3}H$  are the rest of length when we arrive at the leaf states  $s_1, \dots, s_n$ . Let  $R_K(\pi, \tau)$  represents the regret till step episode  $K$  for instance  $\tau$  and policy  $\pi$ , and  $C_K(\pi, \tau)$  represents the violation. For  $\tau_1$ , we can have  $C_K(\pi, \tau_1) \geq \mathbb{P}_{\tau_1}(t_1(K) \leq K/2) \cdot \frac{KH'\Delta}{2}$ , where  $t_i(K)$  means the times to choose the path to  $s_i$ . Also,  $R_K(\pi, \tau_i) \geq \mathbb{P}_{\tau_i}(t_1(K) > K/2) \frac{KH'}{2}$ .

Applying Bertagnolle-Huber inequality, we have

$$\mathbb{P}_{\tau_1}(t_1(K) \leq K/2) + \mathbb{P}_{\tau_i}(t_1(K) > K/2) \geq \frac{1}{2} \exp\{-D(\mathbb{P}_{\tau_1}, \mathbb{P}_{\tau_i})\}.$$

We choose  $i$  such that  $i = \arg \min_i \mathbb{E}_{\tau_i}[t_i(K)]$ , then by  $\sum_{i \neq 1} \mathbb{E}_{\tau_i}[t_i(K)] \leq K$ ,

$$\mathbb{E}_{\tau_1}[t_i(K)] \leq \frac{K}{n-1}.$$

By the definition of  $\mathbb{P}_{\tau_1}$  and  $\mathbb{P}_{\tau_i}$ ,

$$\begin{aligned} D(\mathbb{P}_{\tau_1}, \mathbb{P}_{\tau_i}) &= H' \mathbb{E}_{\tau_1}[t_i(K)] D(\mathcal{N}(1/2 + \Delta, 1), \mathcal{N}(1/2, 1)) \\ &\leq H' \mathbb{E}_{\tau_1}[t_i(K)] \cdot 4\Delta^2 \\ &\leq \frac{4KH'\Delta^2}{n-1}. \end{aligned}$$

Thus choose  $\Delta = \sqrt{n-1/4KH'} \leq \frac{1}{2}$ , (need  $n \leq T$ )

$$\mathbb{P}_{\tau_1}(t_1(K) \leq K/2) + \mathbb{P}_{\tau_i}(t_1(K) > K/2) \geq \frac{1}{2} \exp\{-D(\mathbb{P}_{\tau_1}, \mathbb{P}_{\tau_i})\} \geq \frac{1}{2e}.$$

If  $\mathbb{E}_{\pi}(R_K(\pi, \tau)) \leq \frac{T}{24} = \frac{KH}{24} \leq \frac{KH'}{16}$  for all  $\tau$ , then

$$\mathbb{P}_{\tau_i}(t_1(K) > K/2) \leq \frac{2R_K(\pi, \tau_i)}{KH'} \leq \frac{1}{8}.$$

Thus

$$\mathbb{P}_{\tau_1}(t_1(K) \leq K/2) \geq \frac{1}{2e} - \frac{1}{8} := C$$

and

$$C_K(\pi, \tau_1) \geq \frac{CKH'\Delta}{2} = \frac{C\sqrt{(n-1)KH'}}{4} = \Omega(\sqrt{SHK}) = \Omega(\sqrt{SHK}),$$

where  $n \geq \frac{S}{2}$ . □

#### A.4 Proof of Theorem 5.2

*Proof.* Construct the MDP  $\tau_1, \tau_2, \dots, \tau_n$  as follows: For each instance, the states consist of a tree with degree  $A-1$  and layer  $H/3$ . For each state inside the tree (not a leaf), the action  $a_i, 1 \leq i \leq A-1$  will lead to the  $i$ -th branch of the next state. For the leaf nodes, denote them as  $s_1, s_2, \dots, s_n$ , where  $n = (A-1)^{H/3} = \frac{(S-3)(A-2)+1}{A-1} = \Omega(S)$  and these states will arrive at the final absorbing state  $s_A$  and  $s_B$  with reward  $r(s_A) = 0$  and  $r(s_B) = 1$  respectively. Also, the action  $a_A$  will always lead the agent to the unique unsafe state  $s_U$  with  $r(s_U) = 0.5 + \Delta'$  and  $c(s_U) = 1$  for some parameter  $\Delta'$  in the states except two absorbing states. For all states except the unsafe state, the safety cost is 0.

Now for the instance  $\tau_1$ , the transition probability between the leaf states and absorbing states are

$$\begin{cases} p(s_A | s_1, a) = 0.5 - \Delta \\ p(s_B | s_1, a) = 0.5 + \Delta \\ p(s_A | s_i, a) = p(s_B | s_1, a) = 0.5 \quad i \geq 2 \end{cases}$$

For instance  $\tau_j$ , the transition probability are

$$\begin{cases} p(s_A | s_1, a) = 0.5 - \Delta \\ p(s_B | s_1, a) = 0.5 + \Delta \\ p(s_A | s_j, a) = 0.5 - 2\Delta \\ p(s_B | s_j, a) = 0.5 + 2\Delta \\ p(s_A | s_i, a) = p(s_B | s_1, a) = 0.5 \quad i \neq 1, j \end{cases}$$

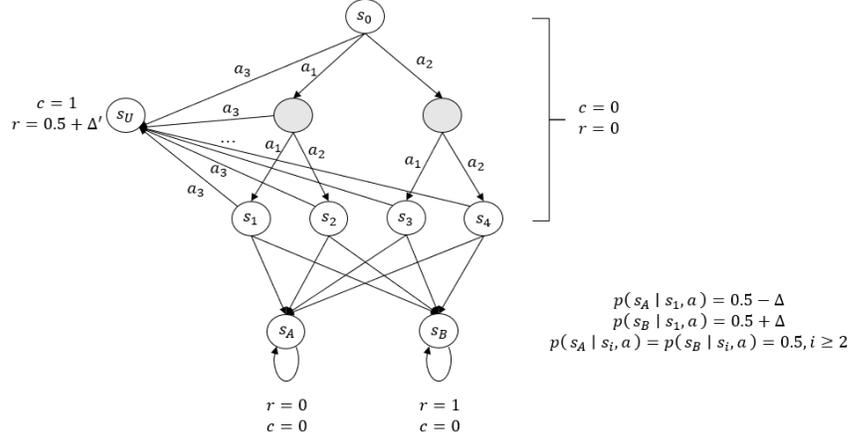


Figure 4: MDP instance  $\tau_1$

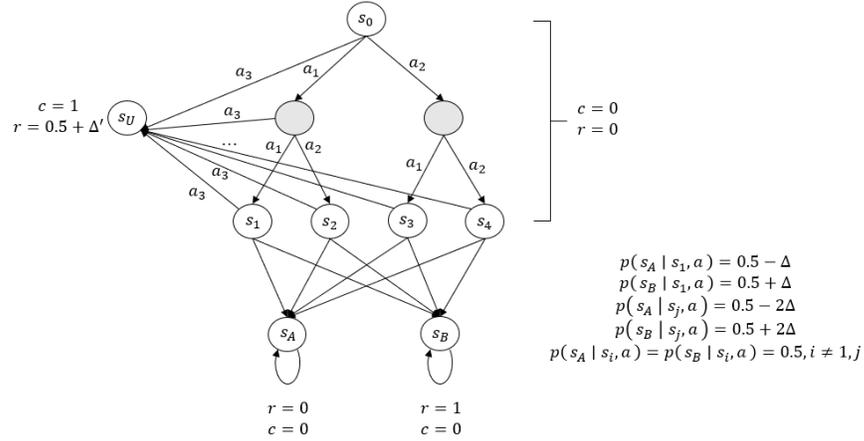


Figure 5: MDP instance  $\tau_i$

Then if we do not consider the constraint, choosing  $a_A$  and getting into the unsafe state  $s_U$  is always the optimal action, then define  $\text{gap}_{h,s,a} = V_h^*(s) - Q_h^*(s, a)$ , then the minimal gap [Simchowit & Jamieson \(2019\)](#) is

$$\text{gap}_{\min} = \min_{h,s,a} \{\text{gap}_{h,s,a} > 0 : \text{gap}_{h,s,a}\} \geq H'((0.5 + \Delta') - (0.5 + 2\Delta)) \geq H'(\Delta' - 2\Delta).$$

Thus the regret will be at most  $O\left(\frac{S^2 A \cdot \text{poly}(H) \log T}{\Delta' - \Delta}\right)$  by classical gap-dependent upper bound [Simchowit & Jamieson \(2019\)](#). As we will show later, we will choose  $\Delta = \Theta(T^{-1/2})$  and  $\Delta' = \Theta(T^{\frac{\alpha-1}{2}})$ , hence the regret will be

$$O\left(\frac{S^2 A \cdot \text{poly}(H) \log T}{\Delta' - \Delta}\right) = \tilde{O}\left(T^{\frac{1-\alpha}{2}}\right).$$

Now if we consider the safe threshold  $\tau = 0$ , the violation will be the number of time to arrive the unsafe state. Now consider any algorithm, if it achieves  $O(T^{1-\alpha})$  violation, they will arrive at unsafe state at most  $O(T^{1-\alpha})$  times. Since denote  $T_i(K), 1 \leq i \leq n+1$  as the number of times when agent arrive at state  $s_i$  till episode  $K$ . Then  $\mathbb{E}_\pi[T_{n+1}(K)] \leq CT^{1-\alpha}$ . Note that each episode we will arrive at least one of  $s_i, 1 \leq i \leq n+1$  once, we know  $\sum_{i=1}^{n+1} T_i(K) \geq K$ , and  $\sum_{i=1}^n T_i(K) \geq K - C \cdot T^{1-\alpha} \geq K/2$  when  $K \geq (2C)^{1/\alpha} H^{\frac{1-\alpha}{\alpha}}$ .

Now similar to proof of Theorem 5.1, let  $C_K(\pi, \tau)$  represents the violation with algorithm  $\pi = \{\pi_1, \dots, \pi_K\}$ . Denote  $H' = \frac{2}{3}H$ . For  $\tau_1$ , we have

$$\begin{aligned} R_K(\pi, \tau_1) &\geq \mathbb{P}_{\tau_1}(T_1(K) \leq K/4) \frac{KH'\Delta}{4} - (\Delta' - \Delta)\mathbb{E}_{\pi}[T_{n+1}(K)] \\ &\geq \mathbb{P}_{\tau_1}(T_1(K) \leq K/4) \frac{KH'\Delta}{4} - (\Delta' - \Delta)CT^{1-\alpha} \end{aligned}$$

$$\begin{aligned} R_K(\pi, \tau_i) &\geq \mathbb{P}_{\tau_i}(T_1(K) > K/4) \frac{KH'\Delta}{4} - (\Delta' - 2\Delta)\mathbb{E}_{\pi}[T_{n+1}(K)] \\ &\geq \mathbb{P}_{\tau_i}(T_1(K) > K/4) \frac{KH'\Delta}{4} - (\Delta' - 2\Delta)CT^{1-\alpha} \end{aligned}$$

Also, we have

$$\mathbb{P}_{\tau_1}(T_1(K) \leq K/4) + \mathbb{P}_{\tau_i}(T_1(K) > K/4) \geq \frac{1}{2} \exp\{-D(\mathbb{P}_{\tau_1}, \mathbb{P}_{\tau_i})\}$$

Now WLOG we have  $\mathbb{E}_{\tau_1}[T_i(K)] \leq \frac{K}{2(n-1)}$

$$D(\mathbb{P}_{\tau_1}, \mathbb{P}_{\tau_i}) \leq \mathbb{E}_{\tau_1}[T_i(K)]2\Delta^2 \leq \frac{K}{2(n-1)} \cdot 2\Delta^2 = \frac{K\Delta^2}{n-1}$$

Choose  $\Delta = \sqrt{\frac{n-1}{K}}$ , then

$$\mathbb{P}_{\tau_1}(T_1(K) \leq K/4) + \mathbb{P}_{\tau_i}(T_1(K) > K/4) \geq \frac{1}{2e}.$$

Then at least one of  $\tau_1$  and  $\tau_i$  (assume  $\tau_1$ ) have

$$R_K(\pi, \tau_1) \geq \frac{H'\sqrt{K(n-1)}}{8e} - C(\Delta' - 2\Delta)T^{1-\alpha}$$

Choose  $\Delta' = T^{\frac{\alpha-1}{2}}$ . Then assume  $T \geq 2^{\frac{2}{1-\alpha}}$  so that  $c \leq \frac{1}{2}$ .

$$R_K(\pi, \tau_1) \geq \frac{H'\sqrt{K(n-1)}}{8e} - CT^{(1-\alpha)/2} = \Omega(H\sqrt{SK}) = \Omega(\sqrt{HST}).$$

□

## B Safe Markov Games

### B.1 Definitions

In this section, we provide an extension for our general framework of Safe-RL-SW. We solve the safe zero-sum two-player Markov games [Yu et al. \(2021\)](#).

In zero-sum two-player Markov games, there is an agent and another adversary who wants to make the agent into the unsafe region and also minimize the reward. The goal of agent is to maximize the reward and avoid the unsafe region with the existence of adversary. To be more specific, define the action space for the agent and adversary are  $\mathcal{A}$  and  $\mathcal{B}$  respectively. In each episode the agent starts at an initial state  $s_1$ . At each step  $h \in [H]$ , the agent and adversary observe the state  $s_h$  and taking action  $a_h \in \mathcal{A}$  and  $b_h \in \mathcal{B}$  from the strategy  $\mu_h(\cdot | s)$  and  $\nu_h(\cdot | s)$  simultaneously. Then the agent receive a reward  $r_h(s, a, b)$  and the environment transitions to the next state  $s_{h+1} \sim P(s_{h+1} | s, a, b)$  by the transition kernel  $P$ .

One motivation of this setting is that, when a robot is in a dangerous environment, the robot should learn a completely safe policy even with the disturbance from the environment. Like the single-player

MDP, we define the value function and Q-value function:

$$V_h^{\mu,v}(s) = \mathbb{E}_\pi \left[ \sum_{h'=h}^H r_{h'}(s_{h'}, a_{h'}, b_{h'}) \mid s_h = s, \mu, v \right].$$

$$Q_h^{\mu,v}(s, a, b) = \mathbb{E}_\pi \left[ \sum_{h'=h}^H r_{h'}(s_{h'}, a_{h'}, b_{h'}) \mid s_h = s, a_h = a, b_h = b, \mu, v \right].$$

The cost function is a bounded function defined on the states  $c(s) \in [0, 1]$ ,  $s \in \mathcal{S}$ , and the unsafe states are defined  $\mathcal{U}_H = \{s \mid c(s) > \tau\}$  for some fixed threshold  $\tau$ . Similar to the Section 4, we define the possible next states for state  $s$ , step  $h$ , action pair  $(a, b)$  as

$$\Delta_h(s, a, b) = \{s' \mid \mathbb{P}_h(s' \mid s, a, b) > 0\}.$$

Hence we can define the unsafe states for step  $h$  recursively as:

$$\mathcal{U}_h = \mathcal{U}_{h+1} \cup \{s \mid \forall a \in \mathcal{A}, \exists b \in \mathcal{B}, \Delta_h(s, a, b) \cap \mathcal{U}_{h+1} \neq \emptyset\}.$$

Then to keep safety, the agent cannot arrive at the state  $s \in \mathcal{U}_h$  at the step  $1, 2, \dots, h$ . The safe action for the agent  $A_h^{safe}(s)$  should avoid the unsafe region no matter what action the adversary takes.

$$A_h^{safe}(s) = \{a \in \mathcal{A} \mid \forall b \in \mathcal{B}, \Delta_h(s, a, b) \cap \mathcal{U}_{h+1} = \emptyset\}.$$

Hence a feasible policy should keep the agent in safe state  $s \notin \mathcal{U}_h$  at step  $h$ , and always chooses  $a \in A_h^{safe}(s)$  to avoid the unsafe region regardless of the adversary. Analogous to the safe RL in Section 4, assume the the optimal Bellman equation can be written as

$$Q_h^*(s, a, b) = r(s, a) + \mathbb{P}_h(s' \mid s, a, b) V_{h+1}^*(s').$$

$$V_h^*(s) = \max_{a \in A_h^{safe}(s)} \min_{b \in \mathcal{B}} Q_h^*(s, a, b).$$

If we denote  $r(s, a, b) = -\infty$  for unsafe states  $s \in \mathcal{U}_H$ , fixed an strategy  $v$  for the adversary, the *best response* of the agent can be defined as  $\arg \max_\mu V_1^{\mu,v}(s_1)$ . Similarly, a *best response* of the adversary given agent's strategy  $\mu$  can be defined as  $\arg \min_v V_1^{\mu,v}(s_1)$ . By the minimax equality, for any step  $h \in [H]$  and  $s \in \mathcal{S}$ ,

$$\max_\mu \min_v V_h^{\mu,v}(s) = \min_v \max_\mu V_h^{\mu,v}(s).$$

A policy pair  $(\mu^*, v^*)$  satisfy that  $V_h^{\mu^*, v^*}(s) = \max_\mu \min_v V_h^{\mu,v}(s)$  for all step  $h \in [H]$  is called a *Nash Equilibrium*, and the value  $V_h^{\mu^*, v^*}(s)$  is called the *minimax value*. It is known that minimax value is unique for Markov games, and the agent's goal is make the reward closer to the minimax value.

Assume the agent and adversary executes the policy  $(\mu_1, v_1), (\mu_2, v_2), \dots, (\mu_K, v_K)$  for episode  $1, 2, \dots, K$ , the regret is defined as

$$R(K) = \sum_{k=1}^K V_1^{\mu^*, v^*}(s_1) - V_1^{\mu_k, v_k}(s_1).$$

Also, the actual state-wise violation is defined as

$$C(K) = \sum_{k=1}^K \sum_{h=1}^H (c(s_h^k) - \tau)_+.$$

## B.2 Algorithms and Results

The main algorithm is presented in Algorithm 3. The key idea is similar to the Algorithm 1. However, in each episode we need to calculate an optimal policy  $\pi^k$  by a sub-procedure "Planning". This sub-procedure calculates the Nash Equilibrium for the estimated MDP. Note that there is always a mixed Nash equilibrium strategy  $(\mu, v)$ , and the ZERO-SUM-NASH function in "Planning" procedure calculates the Nash equilibrium for a Markov games.

---

**Algorithm 3** Safe Markov Games

---

- 1: **Input:**  $\Delta_h^1(s, a, b) = \emptyset, N_h^1(s, a, b), N_h^1(s, a, b, s') = 0$ .
  - 2: **for**  $k = 1, 2, \dots, K$  **do**
  - 3:   Receive the initial state  $s_1$ .
  - 4:   Update the estimation  $\hat{c}(s)$  based on history and get optimistic estimation  $\bar{c}(s) = \hat{c}(s) - \beta(N^k(s), \delta)$  for each state  $s$ , where  $\beta$  is the bounds function.
  - 5:   Define  $\mathcal{U}_H^k = \{s \mid \bar{c}(s) > \tau\}$ .
  - 6:   Based on  $\Delta_h^K(s, a)$ , calculate  $\mathcal{U}_h^k = \mathcal{U}_{h+1}^k \cup \{s \mid \forall a \in A, \Delta_k(s, a, b) \cap \mathcal{U}_{h+1}^k \neq \emptyset\}$ .
  - 7:    $\hat{\mathbb{P}}_h(s' \mid s, a, b) = \frac{N_h^k(s, a, b, s')}{N_h^k(s, a, b)}$ .
  - 8:    $\pi_h^k = \text{Planning}(k, r, \{\mathcal{U}_h^k\}_{h \in [H]}, \Delta_h^K(s, a, b), N^k(s, a, b), \hat{\mathbb{P}})$ .
  - 9:   **for**  $h = 1, 2, \dots, H$  **do**
  - 10:     Take action  $a_h^k$  such that  $(a_h^k, b) \sim \pi_h^k(\cdot, \cdot \mid s)$  and arrive state  $s_{h+1}^k \sim P(s \mid s_h^k, a_h^k, b_h^k)$ , where  $b_h^k$  is the action of the adversary.
  - 11:      $\Delta_h^{k+1}(s_h^k, a_h^k, b_h^k) = \Delta_h^k(s_h^k, a_h^k, b_h^k) \cup \{s_{h+1}^k\}$ .
  - 12:   **end for**
  - 13:   Update  $N_h^{k+1}(s, a, b), N_h^{k+1}(s, a, b, s'), \Delta_h^{k+1}(s, a, b)$  for all  $(s, a, b)$ .
  - 14: **end for**
- 

---

**Algorithm 4**  $\text{Planning}(k, r, \{\mathcal{U}_h\}_{h \in [H]}, \Delta(s, a, b), N, \hat{\mathbb{P}})$ 

---

- 1: **Input**  $k, r, \{\mathcal{U}_h\}, \Delta(s, a, b), N, \hat{\mathbb{P}}$ .
  - 2: **for**  $h = H, H - 1, \dots, 1$  **do**
  - 3:   **for**  $(s, a, b) \in \mathcal{S} \times \mathcal{A} \times \mathcal{B}$  **do**
  - 4:      $\beta(N_h^k(s, a, b), \delta) \leftarrow 7H \sqrt{\log(5SABT/\delta)/N_h^k(s, a, b)}$ .
  - 5:      $Q_h^k(s, a, b) = \min\{(r_h + \sum_{s'} \hat{\mathbb{P}}_h^k(s' \mid s, a, b) V_{h+1}(s')) + \beta(N_h^k(s, a, b), \delta), H\}$ .
  - 6:   **end for**
  - 7:   **for**  $s \in \mathcal{S}$  **do**
  - 8:     **for**  $a \in \mathcal{A}$  **do**
  - 9:       **if**  $\exists b, \Delta_h(s, a, b) \cap \mathcal{U}_{h+1} \neq \emptyset$  **then**
  - 10:          $(Q_h^k)'(s, a, \cdot) = -\infty$ .
  - 11:       **else**
  - 12:          $(Q_h^k)'(s, a, \cdot) = Q_h^k(s, a, \cdot), a \in A_h^k(s)$ .
  - 13:       **end if**
  - 14:     **end for**
  - 15:      $\mu_h^k, v_h^k = \text{ZERO-SUM-NASH}((Q_h^k)'(s, \cdot, \cdot))$ .
  - 16:      $V_h^k(s) = \mathbb{E}_{a \sim \mu_h^k(\cdot \mid s), b \sim v_h^k(\cdot \mid s)}(a, b) Q_h(s, a, b)$ .
  - 17:   **end for**
  - 18: **end for**
- 

In the sub-procedure ‘‘Planning’’, we calculate the  $(Q_h^k)'(s, a, \cdot) = -\infty$  if there exists at least one adversary action  $b$  such that  $\Delta(s, a, b) \cap \mathcal{U}_{h+1} \neq \emptyset$ . In fact, the agent cannot take these actions at this time, hence by setting  $(Q_h^k)'(s, a, \cdot)$  for these actions  $a$ , the support of Nash Equilibrium policy  $\mu_h^k$  will not contains  $a$ . Otherwise, the minimax value is  $-\infty$ , and it means that  $a$  is now in  $\mathcal{U}_h^k$ , then at this time the agent can chooses the action arbitrarily

Our main result are presented in the following theorem, which shows that our algorithm achieves both  $O(\sqrt{T})$  regret and  $O(\sqrt{ST})$  state-wise violation.

**Theorem B.1.** *With probability at least  $1 - \delta$ , the Algorithm 3 have regret and state-wise violation*

$$\begin{aligned} R(T) &= O(H^3 \sqrt{SABT}) \\ C(T) &= O(\sqrt{ST} + S^2 ABH^2), \end{aligned}$$

where  $A = |\mathcal{A}|, B = |\mathcal{B}|$  and  $S = |\mathcal{S}|$ .

### B.3 Proofs

*Proof.* We first prove that, with probability at least  $1 - \delta$ ,  $Q_h^k(s, a, b) \geq Q_h^*(s, a, b)$  and  $V_h^k(s) \geq V_h^*(s)$  for any episode  $k \in [K]$ . We prove this fact by induction. Denote  $\mathbb{D}_{\mu \times v} Q(s) = \mathbb{E}_{a \sim \mu(\cdot|s), b \sim v(\cdot|s)} Q(s, a, b)$ . The statement holds for  $h = H + 1$ . Suppose the bounds hold for  $Q$ -values in the step  $h + 1$ , now we consider step  $h$ .

$$\begin{aligned}
V_{h+1}^k(s) &= \mathbb{D}_{\pi_h^k} Q_{h+1}^k(s) \\
&= \max_{\text{supp}(\mu) \subseteq A_h^{k, \text{safe}}(s)} \mathbb{D}_{\mu \times v_h^k} Q_{h+1}^k(s) \\
&\geq \max_{\mu \in \text{supp}(\mu) \subseteq A_h^k(s)} \mathbb{D}_{\mu \times v_h^k} Q_{h+1}^*(s) \\
&\geq \max_{\mu \in \text{supp}(\mu) \subseteq A_h^*(s)} \mathbb{D}_{\mu \times v_h^k} Q_{h+1}^*(s) \\
&\geq \min_v \max_{\mu \in \text{supp}(\mu) \subseteq A_h^*(s)} \mathbb{D}_{\mu \times v} Q_{h+1}^*(s) \\
&= V_{h+1}^*(s),
\end{aligned}$$

where  $A_h^{k, \text{safe}}(s)$  is all safe action for max-player at episode  $k$  and step  $h$  it estimated,  $A_h^*(s)$  is the true safe action, and  $\text{supp}(\mu) = \{a \in \mathcal{A} \mid \mu(a) > 0\}$ . From our algorithm, we know  $A_h(s) \subseteq A_h^k(s)$  for all time if the confidence event always holds.

Then

$$\begin{aligned}
Q_h^k(s, a, b) - Q_h^*(s, a, b) &= (\hat{\mathbb{P}}_h^k V_{h+1}^k - \mathbb{P}_h V_{h+1}^* + \beta_h^k)(s, a, b) \\
&\geq (\hat{\mathbb{P}}_h^k - \mathbb{P}_h)(V_{h+1}^*)(s, a, b) + \beta(N_h^k(s, a, b), \delta) \\
&\geq 0.
\end{aligned}$$

The last inequality is because of the Chernoff-Hoeffding's inequality. Thus define event

$$\mathcal{E}_6 = \{ \forall k, h, s, a, b, Q_h^k(s, a, b) \geq Q_h^*(s, a, b), V_h^k(s) \geq V_h^*(s) \}.$$

Then  $\Pr\{\mathcal{E}_6^c\} \leq \delta$ .

Now we start to bound the regret.

$$\begin{aligned}
R(K) &= \sum_{k=1}^K V_1^{\mu^*, v^*}(s_1) - V_1^{\mu_k, v_k}(s_1) \\
&\leq \sum_{k=1}^K V_1^k(s_1) - V_1^{\mu_k, v_k}(s_1).
\end{aligned}$$

Now we first assume  $v_k$  is the best response of  $\mu_k$ . Otherwise our regret can be larger. We calculate the regret by

$$\begin{aligned}
&V_h^k(s_h^k) - V_h^{\mu_k, v_k}(s_h^k) \\
&= \mathbb{D}_{\pi^k}(Q_h^k)(s_h^k) - \mathbb{D}_{\mu_k \times v_k} Q_h^{\mu_k \times v_k}(s_h^k) \\
&\leq \mathbb{D}_{\mu_k \times v_k}(Q_h^k)(s_h^k) - \mathbb{D}_{\mu_k \times v_k} Q_h^{\mu_k \times v_k}(s_h^k) \\
&= \mathbb{D}_{\mu_k \times v_k}(Q_h^k - Q_h^{\mu_k \times v_k})(s_h^k) \\
&= (Q_h^k - Q_h^{\mu_k \times v_k})(s_h^k, a_h^k, b_h^k) + \alpha_h^k \\
&= (\hat{\mathbb{P}}_h - \mathbb{P}_h)V_{h+1}^k(s_h^k, a_h^k, b_h^k) + \mathbb{P}_h(V_{h+1}^k - V_{h+1}^{\mu_k \times v_k})(s_h^k, a_h^k, b_h^k) + \beta(N_h^k(s_h^k, a_h^k, b_h^k), \delta) + \alpha_h^k \\
&\leq (V_{h+1}^k - V_{h+1}^{\mu_k \times v_k})(s_{h+1}^k) + \alpha_h^k + \gamma_h^k + 2\beta(N_h^k(s_h^k, a_h^k, b_h^k), \delta) \\
&\quad + (\hat{\mathbb{P}}_h - \mathbb{P}_h)(V_{h+1}^k - V_{h+1}^*)(s_h^k, a_h^k, b_h^k),
\end{aligned}$$

where

$$\alpha_h^k = \mathbb{P}_h(V_{h+1}^k - V_{h+1}^{\mu_k \times v_k})(s_h^k, a_h^k, b_h^k) - (V_{h+1}^k - V_{h+1}^{\mu_k \times v_k})(s_{h+1}^k),$$

$$\gamma_h^k = (\mathbb{D}_{\mu_k \times v_k}(Q_h^k - Q_h^{\mu_k \times v_k})(s_h^k) - (Q_h^k - Q_h^{\mu_k \times v_k})(s_h^k)).$$

The last inequality is because  $(\hat{\mathbb{P}}_h - \mathbb{P}_h)V_{h+1}^*(s_h^k, a_h^k, b_h^k) \leq \beta(N_h^k(s_h^k, a_h^k, b_h^k), \delta)$  by Chernoff-Hoeffding's inequality.

Now by the similar analysis in the previous section, with probability at least  $1 - \delta$ , the event

$$\begin{aligned} \mathcal{E}_7 &= \left\{ \forall s, a, b, s', k, h, |\hat{\mathbb{P}}_h(s' | s, a, b) - \mathbb{P}_h(s' | s, a, b)| \right. \\ &\leq \left. \sqrt{\frac{2\mathbb{P}_h(s' | s, a, b)}{N_h^k(s, a)} \log\left(\frac{2S^2 ABHK}{\delta}\right)} + \frac{2}{3N_h^k(s, a)} \log\left(\frac{2S^2 ABHK}{\delta}\right) \right\}. \end{aligned}$$

holds with probability at least  $1 - \delta$  by Bernstein's inequality. Then for  $\iota = \log\left(\frac{2S^2 ABHK}{\delta}\right)$

$$\begin{aligned} &(\hat{\mathbb{P}}_h - \mathbb{P}_h)(V_{h+1}^k - V_{h+1}^*)(s_h^k, a_h^k, b_h^k) \\ &\leq \sum_{s'} \left( \sqrt{\frac{2\mathbb{P}_h(s' | s, a, b)}{N_h^k(s, a)}} \iota + \frac{2}{3N_h^k(s, a)} \iota \right) (V_{h+1}^k - V_{h+1}^*)(s') \\ &\leq \sum_{s'} \left( \frac{\mathbb{P}_h(s' | s_h^k, a_h^k)}{H} + \frac{H\iota}{2N_h^k(s_h^k, a_h^k)} + \frac{2\iota}{3N_h^k(s_h^k, a_h^k)} \right) (V_{h+1}^k - V_{h+1}^*)(s') \\ &\leq \frac{1}{H} \mathbb{P}_h(V_{h+1}^k - V_{h+1}^*)(s_h^k, a_h^k) + \frac{2H^2 S \iota}{N_h^k(s_h^k, a_h^k)} \\ &\leq \frac{1}{H} \mathbb{P}_h(V_{h+1}^k - V_{h+1}^{\mu_k \times v_k})(s_h^k, a_h^k) + \frac{2H^2 S \iota}{N_h^k(s_h^k, a_h^k)}. \end{aligned}$$

The last inequality is because we assume  $v_k$  is the best response of  $\mu_k$  and then  $V_{h+1}^*(s_h^k, a_h^k) \geq V_{h+1}^{\mu_k \times v_k}(s_h^k, a_h^k)$ . Now we can get

$$\begin{aligned} &V_h^k(s_h^k) - V_h^{\mu_k \times v_k}(s_h^k) \\ &\leq \left(1 + \frac{1}{H}\right) (V_{h+1}^k(s_h^k) - V_{h+1}^{\mu_k \times v_k}(s_h^k)) + 2\alpha_h^k + 2\gamma_h^k + 2\beta(N_h^k(s_h^k, a_h^k, b_h^k), \delta) + \frac{2H^2 S \iota}{N_h^k(s_h^k, a_h^k)} \end{aligned}$$

and then

$$\begin{aligned} V_1^k(s_1) - V_1^{\mu_k \times v_k}(s_h^k) &\leq \sum_{h=1}^H \left(1 + \frac{1}{H}\right)^H (2\alpha_h^k + 2\gamma_h^k + 2\beta(N_h^k(s_h^k, a_h^k, b_h^k), \delta)) + \frac{2H^2 S \iota}{N_h^k(s_h^k, a_h^k)} \\ &\leq \sum_{h=1}^H e \cdot (2\alpha_h^k + 2\gamma_h^k + 2\beta(N_h^k(s_h^k, a_h^k, b_h^k), \delta)) + \frac{2H^2 S \iota}{N_h^k(s_h^k, a_h^k)}. \end{aligned}$$

By Azuma-hoeffding's inequality, with probability at least  $1 - 2\delta$ ,  $\sum_{k=1}^K \sum_{h=1}^H \alpha_h^k \leq O(\sqrt{H^3 K})$ ,  $\sum_{k=1}^K \sum_{h=1}^H \gamma_h^k = O(\sqrt{H^3 K})$  and

$$\sum_{k=1}^K \sum_{h=1}^H \frac{2H^2 S \iota}{N_h^k(s_h^k, a_h^k)} = O(\log K).$$

Note that

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \beta(N_h^k(s_h^k, a_h^k, b_h^k), \delta) &= \sum_{k=1}^K \sum_{h=1}^H 7H \sqrt{\frac{\iota}{N_h^k(s_h^k, a_h^k, b_h^k)}} \\ &\leq O(H\iota) \sum_{(h,s,a,b) \in [H] \times \mathcal{S} \times \mathcal{A} \times \mathcal{B}} \sqrt{N_h^k(s_h^k, a_h^k, b_h^k)} \\ &\leq \tilde{O}(\sqrt{H^3 SABT}). \end{aligned}$$

For the violation, the argument is similar to the Theorem 4.2. During the learning process, the  $s_{h+1}^k \notin \Delta_h^K(s_h^k, a_h^k, b_h^k)$  can appear at most  $S^2ABH$  times because each time the summation

$$\sum_{(h,s,a,b) \in [H] \times S \times \mathcal{A} \times \mathcal{B}} |\Delta_h^K(s, a, b)|$$

will increase at least 1, and it has a upper bound  $S^2ABH$ . Thus it will lead to at most  $S^2ABH^2$  regret. For other situations, the total violation can be upper bounded by  $\tilde{O}(\sqrt{ST})$ . Thus the final violation bound is  $\tilde{O}(S^2ABH^2 + \sqrt{ST})$ . □

## C Technical Lemmas

### C.1 Lemma C.1

**Lemma C.1.** *If*

$$K\varepsilon/2 \leq a\sqrt{K}\sqrt{\gamma(K, \delta)} + b\gamma(K, \delta) \log(K+1),$$

where  $\gamma(K, \delta) = 2\log(SAHK/\delta) + (S-1)\log(e(1+K/(S-1)))$ , then

$$K = \tilde{O}\left(\left(\frac{a^2}{\varepsilon^2} + \frac{b}{\varepsilon}\right) \log\left(\frac{1}{\delta}\right) + \frac{Sb}{\varepsilon}\right),$$

where  $\tilde{O}(\cdot)$  ignores all the term  $\log S, \log A, \log H, \log 1/\varepsilon$  and  $\log \log(1/\delta)$ .

*Proof.* If  $K \leq 4$ , the lemma is trivial. Now assume  $K \geq 4$ , then  $e(1+K) \leq 4K$  and

$$\begin{aligned} \gamma(K, \delta) &\leq 2\log(SAHK/\delta) + S\log(e(1+K)) \\ &\leq 2\log(SAHK/\delta) + S\log(4K) \\ &\leq 2S\log(4K) + \log(SAH/\delta). \end{aligned}$$

Then

$$\begin{aligned} K\varepsilon/2 &\leq a\sqrt{K}\sqrt{\gamma(K, \delta)} + b\gamma(K, \delta) \log(K+1) \\ &\leq 2\max\{a\sqrt{K}\sqrt{\gamma(K, \delta)}, b\gamma(K, \delta) \log(K+1)\}. \end{aligned}$$

**Case 1:** If  $K\varepsilon/2 \leq 2a\sqrt{K}\sqrt{\gamma(K, \delta)}$ , we can get

$$\begin{aligned} K &\leq \frac{16a^2}{\varepsilon^2} (2S\log(4K) + \log(SAH/\delta)) \\ &\leq 2\max\left\{\frac{32Sa^2}{\varepsilon^2} \log(4K), \frac{16a^2}{\varepsilon^2} \log(SAH/\delta)\right\}. \end{aligned}$$

**Subcase 1:** If  $K \leq \frac{64Sa^2}{\varepsilon^2} \log(4K)$  by Lemma C.2, we can get  $K = \tilde{O}(\frac{Sa^2}{\varepsilon^2})$ .

**Subcase 2:** If  $K \leq \frac{16a^2}{\varepsilon^2} \log(SAH/\delta)$ , we complete the proof.

**Case 2:** If  $K\varepsilon/2 \leq 2b\gamma(K, \delta) \log(K+1)$ , we can get

$$\begin{aligned} K &\leq \frac{4b}{\varepsilon} \log(4K) (2S\log(4K) + \log(SAH/\delta)) \\ &\leq 2\max\left\{\frac{4b}{\varepsilon} \log(4K) 2S\log(4K), \frac{4b}{\varepsilon} \log(4K) \log(SAH/\delta)\right\}. \end{aligned}$$

**Subcase 3:** If  $K \leq \frac{8bS}{\varepsilon} \log^2(4K)$  By Lemma C.2, we can get  $K = \tilde{O}(\frac{8bS}{\varepsilon})$ .

**Subcase 4:** If  $K \leq \frac{8b}{\varepsilon} \log(4K) \log(SAH/\delta)$ , then by Lemma C.2, we can get  $K = \tilde{\mathcal{O}}(\frac{b}{\varepsilon} \log(1/\delta))$ . From the four subcases, we complete the proof of Lemma C.1.  $\square$

## C.2 Lemma C.2

**Lemma C.2.** *If  $K \leq Q \log^2(4K)$  with  $Q \geq 4$ , then we have  $K = \tilde{\mathcal{O}}(Q)$ .*

*Proof.* Define  $f(x) = \frac{x}{\log^2(4x)}$ , then  $f'(x) = \frac{\log^2(4x) - \log(4x)}{\log^4(4x)} > 0$  when  $4x \geq 3$ , and  $f(x)$  is increasing over  $x \geq 1$ . Then we only need to prove there exists a constant  $c_2 > 0$  such that

$$f(c_2 Q \log^2 Q) \geq Q.$$

In fact, if this inequality holds, any  $K$  such that  $K \leq Q \log^2(4K)$  should have  $f(K) \leq Q \leq f(c_2 Q \log^2 Q)$ , and then  $K \leq c_2 Q \log^2 Q = \tilde{\mathcal{O}}(Q)$ . To prove this inequality, observe that

$$\begin{aligned} Q \log^2(4c_2 Q \log^2 Q) &= Q(\log 4c_2 Q + \log \log^2 Q)^2 \\ &\leq 2Q \log^2(4c_2 Q) + 2Q(\log \log^2 Q)^2 \\ &\leq 4Q \log^2(4c_2) + 4Q \log^2(Q) + 2Q \log^2 Q \\ &\leq (4 \log^2(4c_2) + 6)Q \log^2 Q. \end{aligned}$$

where the third inequality is because  $\log^2 Q \leq Q$  for  $Q \geq 4$ . So if we choose  $c_2$  to satisfy  $c_2 \geq 4 \log^2(4c_2) + 6$ , we have  $Q \log^2(4c_2 Q \log^2 Q) \leq c_2 Q \log^2 Q$ , which implies  $f(c_2 Q \log^2 Q) \geq Q$ .  $\square$

## D Detailed Analysis of Assumption 4.1

In this section, we provide a more detailed analysis of the equivalence between  $s_1 \notin \mathcal{U}_1$  and the existence of feasible policy. Indeed, if the agent gets into some state  $s \in \mathcal{U}_h$  at step  $h$ , then any action can lead her to an unsafe next state  $s' \in \mathcal{U}_{h+1}$ . Recursively, any action sequence  $a_h, a_{h+1}, \dots, a_{H-1}$  can lead the agent to unsafe states  $s_{h+1} \in \mathcal{U}_{h+1}, \dots, s_H \in \mathcal{U}_H$ . As a result, the agent cannot avoid getting into an unsafe state. Therefore, all feasible policies  $\pi$  must satisfy that, at any step  $h$ , the probability for the agent in the state  $s \in \mathcal{U}_h$  should be zero under policy  $\pi$ .

$$\mathbb{P}_h^\pi(s) = 0, \quad \forall s \in \mathcal{U}_h. \quad (21)$$

Recall that

$$A_h^{safe}(s) = \{a \in \mathcal{A} \mid \Delta_h(s, a) \cap \mathcal{U}_{h+1} = \emptyset\}.$$

From these definitions, a feasible policy  $\pi$  should satisfy that  $\pi_h(s) \in A_h^{safe}(s)$  for any safe state  $s \notin \mathcal{U}_h$ . Moreover, for any safe state  $s_h \notin \mathcal{U}_h$ , there exists at least one safe action  $a_h \in A_h^{safe}(s_h)$ , and all possible next states  $s_{h+1} \in \Delta(s_h, a_h)$  satisfy  $s_{h+1} \notin \mathcal{U}_{h+1}$ . Recursively, there exists at least one feasible action sequence  $a_h \in A_h^{safe}(s_h), \dots, a_H \in A_H^{safe}(s_H)$  which satisfies that  $s_{h'} \notin \mathcal{U}_{h'}$  for  $h+1 \leq h' \leq H$ . Hence the assumption for the existence of feasible policy can be simplified to  $s_1 \notin \mathcal{U}_1$ .

## E Experiment Setup

We perform each algorithm for 20000 episodes in a grid environment (Chow et al. (2018); Wei et al. (2022)) with 25 states and 4 actions (four directions), and report the average reward and violations across runs. In the experiment, we note that the algorithm Optpess-bouns (Efroni et al., 2020b) needs to solve a linear programming with size  $O(SAH)$  for each episode, which is extremely slow. Thus we only compare all algorithms in a relatively small environment. In the Safe-RL-SW experiments, we choose  $\delta = 0.005$  and binary 0-1 cost function. We set the safety threshold as 0.5 for both the CMDP algorithms and the SUCBVI algorithm.

For the Safe-RFE-SW experiments, we compare our algorithm SRF-UCRL with a state-of-the-art RFE algorithm RF-UCRL (Kaufmann et al., 2021) in an MDP with 11 states and 5 actions. We choose  $\delta = 0.005$  and safety threshold of SRF-UCRL as 0.5. We run 500 episodes 100 times, and then report the average reward in each episode and cumulative step-wise violation. Also at each episode, we calculate the expected violation for the outputted policy and plot the curves. All experiments are conducted with AMD Ryzen 7 5800H with Radeon Graphics and a 16GB unified memory.