

On Distributed Multi-player Multiarmed Bandit Problems in Abruptly Changing Environment

Lai Wei

Vaibhav Srivastava

Abstract—We study the multi-player stochastic multiarmed bandit (MAB) problem in an abruptly changing environment. We consider a collision model in which a player receives reward at an arm if it is the only player to select the arm. We design two novel algorithms, namely, Round-Robin Sliding-Window Upper Confidence Bound# (RR-SW-UCB#), and the Sliding-Window Distributed Learning with Prioritization (SW-DLP). We rigorously analyze these algorithms and show that the expected cumulative group regret for these algorithms is upper bounded by sublinear functions of time, i.e., the time average of the regret asymptotically converges to zero. We complement our analytic results with numerical illustrations.

I. INTRODUCTION

Achieving coordinated behavior of multiple decision-makers in unknown, uncertain, and non-stationary environments without any explicit communication among them is of immense interest in a variety of applications, including robotic swarming, spectrum access, and Internet of Things. Such decision-making problems often require decision-makers to balance tradeoff between learning the environment and the policies of other decision-makers, and maximizing their own utility.

The distributed multi-player MAB problem embodies this fundamental decision-making tradeoff. In these problems, multiple decision-makers sequentially allocate a single resource each by repeatedly choosing one among a set of competing alternative arms (options). The objective of decision-makers is to maximize their total reward, while ensuring that their allocation at each time is not coincident with any other decision-maker. In a non-stationary version of these problems, the utility (reward) of the arms varies with time. These problems provide a rich modeling framework for a variety of interesting settings, including, robotic foraging and surveillance [1–3], acoustic relay positioning for underwater communication [4], and channel allocation in communication networks [5].

Most of the algorithmic solutions to the multi-player MAB problem deal with a stationary environment. In [11, 12], a lower bound on expected cumulative group regret for a centralized policy is derived and algorithms that asymptotically achieve this lower bound are designed. Distributed multi-player MAB problem with no communication among players has been studied in [5, 13–16]. Distributed cooperative multi-player MAB problem in which agents communicate to

improve their estimates of mean rewards are studied in [17–19].

In the context of the single player MAB problem, some classes of non-stationary MAB problem have been studied in the literature. In [6], authors study a non-stochastic MAB problem in which the rewards are deterministic and non-stationary. They propose algorithmic solutions to this problem and study a weaker notion of the regret, wherein the policy generated by the algorithm is compared against the best policy within the policies that select the same arm at each time. These algorithms are adapted to handle a class of non-stationary environments and upper bounds on the standard notion of the regret are derived in [7]. In [8], authors study a class of non-stationary MAB problems in which the mean rewards at arms may switch abruptly at unknown times to unknown values. They design an upper confidence bound (UCB) based algorithm that relies on estimates of the mean rewards from a recent time-window of observations. In [9], authors study the MAB problem in a piecewise-stationary environment. They use active detection algorithms to determine the change-points and restart the UCB algorithm. Authors in [10] develop variants of the algorithm in [8] for non-stationary MAB problems with abruptly-changing and slowly-varying mean rewards. They also develop deterministic sequencing of exploration and exploitation type of algorithms for these environments.

In this paper, we build upon the literature on distributed multi-player MAB problem with no communication among players and design analogous algorithms for a class of non-stationary environments. In particular, we extend ideas in [10] to the multi-player setting and design two algorithms to solve these problems. The major contributions of this work are twofold. First, we design and rigorously analyze two algorithms, namely, RR-SW-UCB# and SW-DLP for the multi-player MAB problem in abruptly changing environments. Second, we augment our analysis with numerical illustrations.

The paper is organized as follows. In Section II, we discuss preliminaries and introduce the problem. In Section III and IV, we design and analyze the RR-SW-UCB# and SW-DLP algorithms, respectively. We illustrate these algorithms with some numerical examples in Section V. Finally, we conclude in Section VI.

II. BACKGROUND & PROBLEM DESCRIPTION

In this section, we formulate the multi-player non-stationary MAB problem, and introduce a class of non-stationary environments: the abruptly-changing environment.

This work has been supported by NSF Award IIS-1734272.

L. Wei and V. Srivastava are with the Department of Electrical and Computer Engineering, Michigan State University, East Lansing, MI 48823 USA. e-mail: weilail1@msu.edu; e-mail: vaibhav@egr.msu.edu

A. The non-stationary multi-player stochastic MAB problem

Consider an N -armed bandit problem, i.e., an MAB problem with N arms. Let the total number of players be $M \in \{1, \dots, N\}$. The reward associated with arm $i \in \{1, \dots, N\}$ is modeled as a random variable with bounded support $[0, 1]$ and an unknown time-varying mean $\mu_i(t) \in [0, 1]$. Each player $j \in \{1, \dots, M\}$ selects a particular arm $s_j(t) \in \{1, \dots, N\}$ at time $t \in \{1, \dots, T\}$ and observes a reward $r_{s_j(t)}(t)$ associated with $s_j(t)$, where T is the time horizon. The adaptive decision rule employed by player j is given by $\rho(\cup_{d=1}^t \{r_{s_j(d)}(t), s_j(d)\})$, which is a distributed policy based on its own observation and decision history.

We assume a collision model \mathcal{M} in which player j receives a reward $r_{s_j(t)}(t)$ from arm $s_j(t)$ if it is the only player to select arm $s_j(t)$ at time t . Then, the expected total reward can be written as

$$\mathbb{E}^\rho[S] = \mathbb{E}^\rho \left[\sum_{t=1}^T \sum_{i=1}^N \sum_{j=1}^M \mu_i(t) \mathcal{O}_{i,j}(t) \right], \quad (1)$$

where $\mathcal{O}_{i,j}(t) = 1$ if only player j selects arm i at time t and is zero otherwise; and the expectation is computed over different realization of $\mathcal{O}_{i,j}(t)$ under policy ρ .

Let σ_t be a permutation of $\{1, \dots, N\}$ at time t such that

$$\mu_{\sigma_t(1)}(t) \geq \dots \geq \mu_{\sigma_t(N)}(t).$$

The maximization of $\mathbb{E}^\rho[S]$ in (1) is equivalent to minimizing the expected cumulative group regret defined by

$$R_M^\rho(T) = \sum_{t=1}^T \sum_{k=1}^M \mu_{\sigma_t(k)}(t) - \mathbb{E}^\rho \left[\sum_{t=1}^T \sum_{i=1}^N \sum_{j=1}^M \mu_i(t) \mathcal{O}_{i,j}(t) \right].$$

$R_M^\rho(T)$ is the expect difference between the cumulative reward of the perfect selection and that of policy ρ under collision model \mathcal{M} . Our main purpose here is to design policy ρ that minimizes $R_M^\rho(T)$.

In this paper, we study the above MAB problem in an abruptly-changing environment, in which the mean rewards from arms switch to unknown values at unknown time instants. We refer to these time instants as *breakpoints*. We assume that the number of breakpoints until time T is $\Upsilon_T \in O(T^\nu)$, where $\nu \in [0, 1)$ and is known a priori. In addition, the minimum difference in mean rewards between any pair of arms at any time is lower bounded by Δ_{\min} .

B. Algorithms for the stationary stochastic MAB problem

In the stationary MAB problem, $\mu_i(t)$ is a constant function of time denoted by μ_i , for each $i \in \{1, \dots, N\}$. In this setting, two algorithms relevant to this paper are : (i) UCB for the single player MAB problem (ii) Distributed Learning algorithm with Prioritization (DLP) for multi-player MAB problem.

The UCB algorithm initializes by selecting each arm once and subsequently selects the arm $s(t)$ at time t defined by

$$s(t) = \arg \max_{i \in \{1, \dots, N\}} \{ \bar{r}_i(t-1) + c_i(t-1) \},$$

where $\bar{r}_i(t-1)$ is the statistical mean of the reward from arm i until time t , $c_i(t-1) = \sqrt{\frac{2 \ln t}{n_i(t-1)}}$ and $n_i(t-1)$ is the number of times arm i has been selected until time t . Auer *et al.* [20] showed that UCB algorithm achieves expected cumulative regret that is within a constant factor of the optimal.

DLP is a distributed algorithm based on the UCB algorithm. Here, distributed algorithm means that a player makes a decision only based upon its own observation history and does not communicate with other players. To avoid regret caused by multiple players selecting the same arm, Selective Learning of the k -th largest expected reward (SL(k)) algorithm is designed aiming at making player k pick the k -th best arm. Towards this end, player k determines the set $A_k(t)$ of k arms associated with the k largest values in the set $\{\bar{r}_i(t-1) + c_i(t-1) \mid i \in \{1, \dots, N\}\}$, and select arm

$$s_k(t) = \arg \min_{i \in A_k(t)} \{ \bar{r}_i(t-1) - c_i(t-1) \}.$$

Gai *et al.* [13] established that DLP achieves order-optimal expected cumulative regret.

C. SW-UCB# for the non-stationary MAB problem

The Sliding-Window UCB (SW-UCB) algorithm [8] was adapted in [10] to develop SW-UCB#. SW-UCB# is an algorithm for single player MAB problem in non-stationary environment. At time t , it maintains an estimate of the mean reward $\bar{r}_i(t, \alpha)$ at each arm i , using only the rewards collected within a sliding-window of observations. Let the width of the sliding-window at time $t \in \{1, \dots, T\}$ be $\tau(t, \alpha) = \min\{\lceil \lambda t^\alpha \rceil, t\}$, where parameter $\alpha \in (0, 1]$ is tuned based on environment characteristics. Let $n_i(t, \alpha) = \sum_{d=t-\tau(t, \alpha)+1}^t \mathbf{1}_{\{s(d)=i\}}$ be the number of times arm i has been selected within the time-window at time t , where $\mathbf{1}_{\{\cdot\}}$ is the indicator function, and

$$\bar{r}_i(t, \alpha) = \frac{1}{n_i(t, \alpha)} \sum_{d=t-\tau(t, \alpha)+1}^t r_i(d) \mathbf{1}_{\{s(d)=i\}}. \quad (2)$$

Based on the above estimate, the SW-UCB# algorithm at each time selects the arm

$$s(t) = \arg \max_{i \in \{1, \dots, N\}} \{ \bar{r}_i(t-1, \alpha) + c_i(t-1, \alpha) \}, \quad (3)$$

where $c_i(t-1, \alpha) = \sqrt{\frac{(1+\alpha) \ln t}{n_i(t-1, \alpha)}}$. In the following, we would refer to $\bar{r}_i(t-1, \alpha) + c_i(t-1, \alpha)$, for $i \in \{1, \dots, N\}$ as the upper confidence bounds on the estimated rewards. It has been shown in [10] that the expected cumulative regret for SW-UCB# under the abruptly-changing environment is upper bounded by a sublinear function of time, i.e., the time average of the regret asymptotically converges to zero.

III. THE ROUND ROBIN SW-UCB# ALGORITHM IN ABRUPTLY-CHANGING ENVIRONMENT

In this section, we present the Round Robin SW-UCB# algorithm (RR-SW-UCB#) for multi-player MAB problem in abruptly-changing environment.

Algorithm 1: The RR-SW-UCB# Algorithm

For abruptly-changing environment

```

Input :  $\nu \in [0, 1]$ ,  $\Delta_{\min} \in (0, 1)$ ,  $\lambda \in \mathbb{R}_{>0}$ ,  $T \in \mathbb{N}$  &  $k$ ;
Set :  $\alpha = \frac{1-\nu}{2}$ ;
Output : sequence of arm selection for player  $k$ ;
```

% Initialization:

- 1 Set $\Omega_k \leftarrow \emptyset$, ordered set $\mathcal{G} \leftarrow ()$, and $t \leftarrow 1$;
- 2 **while** $t \leq T$ **do**
- 3 **if** $t \in \{1, \dots, N\}$ **then**- Pick arm $s_k(t) = \text{mod}(t + k - 2, N) + 1$;

- 4 **else**- Compute Ω_k containing M arms with M largest values in
$$\{\bar{r}_i^k(t-1, \alpha) + c_i^k(t-1, \alpha) \mid i \in \{1, \dots, N\}\};$$
- Ascending sort the arm indices in Ω_k , $\mathcal{G} \leftarrow \text{sort}_{\uparrow}(\Omega_k)$;
- for** $i \in \{1, \dots, M\}$ **do**- Pick arm $s_k(t) = \mathcal{G}(\text{mod}(t - N + k - 2, M) + 1)$;
- $t \leftarrow t + 1$;

A. The RR-SW-UCB# algorithm

In the RR-SW-UCB# algorithm, each player maintains the estimate of mean rewards using a sliding window of observations. Each player $k \in \{1, \dots, M\}$ computes $\bar{r}_i^k(t, \alpha)$ and $c_i^k(t, \alpha)$ as in (2) and (3) using their own observation, and maintains an upper confidence bound on the estimated mean reward $\bar{r}_i^k(t-1, \alpha) + c_i^k(t-1, \alpha)$ from each arm $i \in \{1, \dots, N\}$. For initial N iterations, i.e., $t \in \{1, \dots, N\}$, the player k selects each arm once. Then, at time instants $\{N + \eta M + 1\}_{\eta \in \mathbb{Z}_{\geq 0}}$, it computes the set Ω_k containing M arms with M largest values in the set

$$\{\bar{r}_i^k(t-1, \alpha) + c_i^k(t-1, \alpha) \mid i \in \{1, \dots, N\}\}.$$

These arms are sorted in ascending value of their indices (not using the upper confidence bounds), and the player k selects these arms in a round robin fashion starting with the k -th arm in the set. It will be shown in the following section that the estimated set of M best arms, Ω_k , will be the same for each player with high probability. Details of RR-SW-UCB# is shown in Algorithm 1. The free parameter λ in the algorithm can be used to refine the finite time performance of the algorithm.

B. Analysis of RR-UCB#

Before the analysis, we introduce the following notation. Let $\Omega_*^M(t)$ denote the set of M arms with the M largest mean rewards at time t . Then, the total number of times $\Omega_k(t) \neq \Omega_*^M(t)$ until time T can be defined as

$$\mathcal{N}_k(T) := \sum_{t=1}^T \mathbf{1}_{\{\Omega_k(t) \neq \Omega_*^M(t)\}}.$$

We now upper bound $\mathcal{N}_k(T)$ in the following lemma.

Lemma 1: For the RR-SW-UCB# algorithm and the multi-player MAB problem with N arms and M players

in the abruptly-changing environment with the number of break points $\Upsilon_T = O(T^\nu)$, $\nu \in [0, 1]$, the total number of times $\Omega_k(t) \neq \Omega_*^M(t)$ until time T for any player k satisfies

$$\begin{aligned} \mathcal{N}_k(T) \leq (N - M) & \left[\left(\frac{T^{1-\alpha}}{\lambda(1-\alpha)} + 1 \right) \left(1 + \frac{4M(1+\alpha) \ln T}{\Delta_{\min}^2} \right) \right. \\ & \left. + \frac{\pi^2}{3} \left(\frac{\lambda + M + 1}{M} \right)^2 \right] + \Upsilon_T (\lceil \lambda(T-1)^\alpha \rceil + M - 1) + N. \end{aligned}$$

Proof: We begin by separately analyzing windows with and without breakpoints.

Step 1: We define set \mathcal{T} such that for all $t \in \mathcal{T}$, t is either a breakpoint or there exists a break point in its sliding-window of observations $\{t - \tau(t-1, \alpha), \dots, t-1\}$. For $t \in \mathcal{T}$, the statistical means are biased. It follows that

$$|\mathcal{T}| \leq \Upsilon_T \lceil \lambda(T-1)^\alpha \rceil.$$

Consequently, $\mathcal{N}_k(T)$ can be upper-bounded as

$$\mathcal{N}_k(T) \leq \Upsilon_T (\lceil \lambda(T-1)^\alpha \rceil + M - 1) + \tilde{\mathcal{N}}_k(T), \quad (4)$$

where $\tilde{\mathcal{N}}_k(T) := \sum_{t=1}^T \mathbf{1}_{\{\Omega_k(t) \neq \Omega_*^M(t), t \notin \mathcal{T}\}}$. The term $M-1$ in (4) is due to the fact that Ω_k is computed every M steps. In the following steps, we will bound $\tilde{\mathcal{N}}_k(T)$.

Step 2: If $\Omega_k(t) \neq \Omega_*^M(t)$, there exists at least one arm i such that $i \in \Omega_k(t)$ and $i \notin \Omega_*^M(t)$. Then, it can be shown that

$$\begin{aligned} \tilde{\mathcal{N}}_k(T) \leq N & + \sum_{t=N+1}^T \sum_{i=1}^N \mathbf{1}_{\{i \in \Omega_k(t), i \notin \Omega_*^M(t), t \notin \mathcal{T}, n_i(t-1, \alpha) < l(t, \alpha)\}} \\ & + \sum_{t=N+1}^T \sum_{i=1}^N \mathbf{1}_{\{i \in \Omega_k(t), i \notin \Omega_*^M(t), t \notin \mathcal{T}, n_i(t-1, \alpha) \geq l(t, \alpha)\}}, \end{aligned} \quad (5)$$

where we choose $l(t, \alpha) = \frac{4(1+\alpha) \ln t}{\Delta_{\min}^2}$.

We begin with bounding the second term on the right hand side of inequality (5). First, we partition time instants into G epochs. Let $G \in \mathbb{N}$ be such that

$$[\lambda(1-\alpha)(G-1)]^{\frac{1}{1-\alpha}} < T \leq [\lambda(1-\alpha)G]^{\frac{1}{1-\alpha}}. \quad (6)$$

Then, we have the following epochs

$$\{1 + \phi(g-1), \dots, \phi(g)\}_{g \in \{1, \dots, G\}}, \quad (7)$$

where $\phi(g) = \lfloor [\lambda(1-\alpha)g]^{\frac{1}{1-\alpha}} \rfloor$. Let \tilde{t} be any time instant other than the first instant in the g -th epoch. We will now show that all but one of the time instants in the g -th epoch until \tilde{t} are ensured to be contained in the time-window at \tilde{t} . Towards this end, consider the increasing convex function $f(x) = x^{\frac{1}{1-\alpha}}$ with $\alpha \in (0, 1)$. It follows that $f(x_2) - f(x_1) \leq f'(x_2)(x_2 - x_1)$ if $x_2 \geq x_1$. Then, substituting $x_1 = g-1$ and $x_2 = \frac{\tilde{t}^{1-\alpha}}{\lambda(1-\alpha)}$ in the above inequality and simplifying, we get

$$\tilde{t} - (\lambda(1-\alpha)(g-1))^{\frac{1}{1-\alpha}} \leq \lambda \tilde{t}^\alpha \left(\frac{\tilde{t}^{1-\alpha}}{\lambda(1-\alpha)} - g + 1 \right).$$

Since by definition of the g -th epoch, $\frac{\tilde{t}^{1-\alpha}}{\lambda(1-\alpha)} \leq g$, we have

$$\begin{aligned} \tilde{t} - \lfloor (\lambda(1-\alpha)(g-1))^{\frac{1}{1-\alpha}} \rfloor & \leq \min\{\tilde{t} + 1, \lambda \lceil \tilde{t}^\alpha \rceil + 1\} \\ & = \tau(\tilde{t}, \alpha) + 1. \end{aligned}$$

The only time instant in g -th epoch that is possibly not contained in time window at \tilde{t} is $1 + \phi(g - 1)$. Then for any arm $i \in \{1, \dots, N\}$,

$$\sum_{2+\phi(g-1)}^{\tilde{t}} \mathbf{1}_{\{i \in \Omega_k(t)\}} \leq M n_i(\tilde{t}). \quad (8)$$

Furthermore, in the g -th epoch in the partition, either

$$\sum_{t \in g\text{-th epoch}} \sum_{i=1}^N \mathbf{1}_{\{i \in \Omega_k(t), i \notin \Omega_*^M(t), t \notin \mathcal{T}, n_i(t-1, \alpha) < l(t, \alpha)\}} = 0,$$

or there exist at least one time-instant t in the g -th epoch such that

$$\sum_{i=1}^N \mathbf{1}_{\{i \in \Omega_k(t), i \notin \Omega_*^M(t), t \notin \mathcal{T}, n_i(t-1, \alpha) < l(t, \alpha)\}} > 0.$$

Let the last time instant satisfying this condition in the g -th epoch be

$$t(g) = \max\{t \in g\text{-th epoch} \mid \sum_{i=1}^N \mathbf{1}_{\{i \in \Omega_k(t), i \notin \Omega_*^M(t), t \notin \mathcal{T}, n_i(t-1, \alpha) < l(t, \alpha)\}} > 0\}.$$

Note that $t(g) \notin \mathcal{T}$ indicates, for each $i \in \{1, \dots, N\}$, $\mu_i(s)$ is a constant for all $s \in \{t(g) - \tau(t(g) - 1, \alpha), \dots, t(g)\}$. Then, it follows from (8) that

$$\begin{aligned} & \sum_{t \in g\text{-th epoch}} \sum_{i=1}^N \mathbf{1}_{\{i \in \Omega_k(t), i \notin \Omega_*^M(t), t \notin \mathcal{T}, n_i(t-1, \alpha) < l(t, \alpha)\}} \\ & \leq N - M + \sum_{t=\phi(g-1)+2}^{t(g)} \sum_{i=1}^N \mathbf{1}_{\{i \notin \Omega_*^M(t), t \notin \mathcal{T}, n_i(t-1, \alpha) < l(t, \alpha)\}} \\ & \leq N - M + \sum_{i=1}^N M l(t_i(g), \alpha) \mathbf{1}_{\{i \notin \Omega_*^M(t(g)), t(g) \notin \mathcal{T}\}} \\ & \leq (N - M) \left(1 + \frac{4M(1 + \alpha) \ln T}{\Delta_{\min}^2} \right), \end{aligned} \quad (9)$$

where $t_i(g) = \max\{t \in g\text{-th epoch} \mid i \in \Omega_k(t), i \notin \Omega_*^M(t), t \notin \mathcal{T}, n_i(t-1, \alpha) < l(t, \alpha)\}$ and $t_i(g) \leq t(g)$ for all $i \in \{1, \dots, N\}$. Therefore, from (7) and (9), we have

$$\begin{aligned} & \sum_{t=N+1}^T \sum_{i=1}^N \mathbf{1}_{\{i \in \Omega_k(t), i \notin \Omega_*^M(t), n_i(t-1, \alpha) < l(t, \alpha)\}} \\ & \leq G(N - M) \left(1 + \frac{4M(1 + \alpha) \ln T}{\Delta_{\min}^2} \right). \end{aligned} \quad (10)$$

Step 3: In this step, we bound the expectation of the last term in (5). It can be shown that

$$\begin{aligned} & \sum_{i=1}^N \mathbf{1}_{\{i \in \Omega_k(t), i \notin \Omega_*^M(t), t \notin \mathcal{T}, n_i(t-1, \alpha) \geq l(t, \alpha)\}} \\ & \leq \sum_{i \notin \Omega_*^M(t)} \sum_{\zeta \in \Omega_*^M(t)} \sum_{s_\zeta=1}^{h(t)} \sum_{s_i=l(t, \alpha)}^{h(t)} \mathbf{1}_{\{n_\zeta(t-1, \alpha) = s_\zeta, n_i(t-1, \alpha) = s_i, t \notin \mathcal{T}\}} \\ & \times \mathbf{1}_{\{\bar{r}_\zeta(t-1, \alpha) + c_\zeta(t-1, \alpha) \leq \bar{r}_i(t-1, \alpha) + c_i(t-1, \alpha), n_i(t-1, \alpha) \geq l(t, \alpha)\}}, \end{aligned} \quad (11)$$

where $h(t) := \left\lceil \frac{\lceil \lambda(t-1)^\alpha \rceil}{M} \right\rceil$ is the maximum number of times an arm can be selected within the time window at $t-1$. Note that $\bar{r}_j(t-1, \alpha) + c_j(t-1, \alpha) \leq \bar{r}_i(t-1, \alpha) + c_i(t-1, \alpha)$ means at least one of the following holds.

$$\bar{r}_i(t-1, \alpha) \geq \mu_i(t) + c_i(t-1, \alpha), \quad (12)$$

$$\bar{r}_\zeta(t-1, \alpha) \leq \mu_\zeta(t) - c_\zeta(t-1, \alpha), \quad (13)$$

$$\mu_\zeta(t) - \mu_i(t) < 2c_i(t-1, \alpha). \quad (14)$$

Since $n_i(t-1, \alpha) \geq l(t, \alpha)$, (14) does not hold. Applying Chernoff-Hoeffding inequality [21, Theorem 1] to bound the probability of events (12) and (13), we obtain

$$\mathbb{P}(\bar{r}_i(t-1, \alpha) \geq \mu_i(t) + c_i(t-1, \alpha)) \leq t^{-2(1+\alpha)}, \quad (15)$$

$$\mathbb{P}(\bar{r}_\zeta(t-1, \alpha) \leq \mu_\zeta(t) - c_\zeta(t-1, \alpha)) \leq t^{-2(1+\alpha)}. \quad (16)$$

Since Ω_k is only computed at time instants $\{N + \eta M + 1\}_{\eta \in \mathbb{Z}_{\geq 0}}$, from (11), (15) and (16), we have

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=N+1}^T \sum_{i=1}^N \mathbf{1}_{\{i \in \Omega_k(t), i \notin \Omega_*^M(t), t \notin \mathcal{T}, n_i(t-1, \alpha) \geq l(t, \alpha)\}} \right] \\ & \leq (N - M) M \sum_{s_\zeta=1}^{h(f(\eta))} \sum_{s_i=l(t, \alpha)}^{h(f(\eta))} \sum_{\eta=0}^{\lceil \frac{T-N}{M} \rceil} 2M f(\eta)^{-2(1+\alpha)} \\ & \leq (N - M) M^2 \sum_{\eta=0}^{\lceil \frac{T-N}{M} \rceil} 2f(\eta)^{-2(1+\alpha)} h(f(\eta))^2 \\ & \leq (N - M) \left(\frac{\lambda + M + 1}{M} \right)^2 \sum_{\eta=1}^{\infty} 2\eta^{-2} \\ & = \frac{\pi^2}{3} (N - M) \left(\frac{\lambda + M + 1}{M} \right)^2, \end{aligned} \quad (17)$$

where $f(\eta) := N + \eta M + 1$. Therefore, it follows from (4), (5), (10), and (17) that

$$\begin{aligned} \mathcal{N}_k(T) & \leq N + (N - M) \left[G \left(1 + \frac{4M(1 + \alpha) \ln T}{\Delta_{\min}^2} \right) \right. \\ & \quad \left. + \frac{\pi^2}{3} \left(\frac{\lambda + M + 1}{M} \right)^2 \right] + \Upsilon_T (\lceil \lambda(T-1)^\alpha \rceil + M - 1). \end{aligned}$$

From (6), we have $G \leq \frac{T^{1-\alpha}}{\lambda(1-\alpha)} + 1$, and this yields the conclusion. \blacksquare

Based on Lemma 1, we now establish the order of expected cumulative group regret of RR-SW-UCB# in the abruptly changing environment.

Theorem 2: For the RR-SW-UCB# algorithm and the multi-player MAB problem with N arms and M players in the abruptly-changing environment with the number of break points $\Upsilon_T = O(T^\nu)$, $\nu \in [0, 1)$, under collision model \mathcal{M} , the expected cumulative group regret satisfies

$$R_{\mathcal{M}}^{\text{RR-SW-UCB\#}}(T) \in O(T^{\frac{1+\nu}{2}} \ln T)$$

Proof: If all player identify $\Omega_*^M(t)$ correctly at time t , no expected regret is accrued. It follows from Lemma 1 that $\mathcal{N}_k(T) \in O(T^{\frac{1+\nu}{2}} \ln T)$ for all $k \in \{1, \dots, M\}$. The total number of times that any player misidentifies $\Omega_*^M(t)$ until time T can be upper bounded by $\sum_{k=1}^M \mathcal{N}_k(T)$. Thus, we conclude the proof. \blacksquare

Algorithm 2: The SW-DLP Algorithm

Input, output and parameters are the same as RR-SW-UCB#
% Initialization:
1 Set $A_k \leftarrow \emptyset$ and $t \leftarrow 1$;
2 **while** $t \leq T$ **do**
3 **if** $t \in \{1, \dots, N\}$ **then**
 | Pick arm $s_k(t) = \text{mod}(t + k - 2, N) + 1$;
4 **else**
 | Compute A_k containing k arms with k largest values in
 | $\{\bar{r}_i^k(t-1, \alpha) + c_i^k(t-1, \alpha) \mid i \in \{1, \dots, N\}\}$;
 | Pick arm
 | $s_k(t) = \arg \min_{i \in A_k} \{\bar{r}_i^k(t-1, \alpha) - c_i^k(t-1, \alpha)\}$;

IV. THE SW-DLP ALGORITHM IN ABRUPTLY-CHANGING ENVIRONMENT

In this section, we introduce the SW-DLP algorithm for multi-player MAB problem in abruptly-changing environment.

A. The SW-DLP algorithm

The SW-DLP algorithm combines the ideas from SW-UCB# and DLP. Similar to SW-UCB#, upper confidence bounds on the mean rewards are computed using a sliding window of observations as in (2) and (3). Following the same allocation rule in DLP, at each time instant, player k computes a set A_k containing all k largest values in the set

$$\{\bar{r}_i^k(t-1, \alpha) + c_i^k(t-1, \alpha) \mid i \in \{1, \dots, N\}\},$$

and selects arm

$$s_k(t) = \arg \min_{i \in A_k(t)} \{\bar{r}_i^k(t-1, \alpha) - c_i^k(t-1, \alpha)\}.$$

Details about the SW-DLP is shown in Algorithm 2. The parameters in the SW-DLP algorithm are the same as in the RR-SW-UCB# algorithm. In the following, we will refer to $\bar{r}_i^k(t-1, \alpha) - c_i^k(t-1, \alpha)$ as the lower confidence bound on the estimate reward from arm i .

B. Analysis of SW-DLP

Theorem 3: For the SW-DLP algorithm and the multi-player MAB problem with N arms and M players in the abruptly-changing environment with the number of break points $\Upsilon_T = O(T^\nu)$, $\nu \in [0, 1]$, under collision model \mathcal{M} , the expected cumulative group regret satisfies

$$R_{\mathcal{M}}^{\text{SW-DLP}}(T) \in O(T^{\frac{1+\nu}{2}} \ln T)$$

Proof: The proof is similar to the proof of Theorem 2 and we only present a sketch. Let $\theta_k(t)$ be the k -th best arm at time t . The total number of time instants that $\theta_k(t)$ is not selected by player k with SW-DLP satisfies

$$\hat{\mathcal{N}}_k \leq \Upsilon_T \lceil \lambda(T-1)^\alpha \rceil + \sum_{t=1}^T \mathbf{1}_{\{s_k(t) \neq \theta_k(t), t \notin \mathcal{T}\}}, \quad (18)$$

where $\hat{\mathcal{N}}_k := \sum_{t=1}^T \mathbf{1}_{\{s_k(t) \neq \theta_k(t)\}}$.

We partition the time instants as in (7). Then, similarly to (8) in the proof of Lemma 1, it can be shown that

$$\sum_{t=2+\phi(g-1)}^{\tilde{t}} \mathbf{1}_{\{s_k(t)=i\}} \leq n_i(\tilde{t}), \quad (19)$$

for any arm $i \in \{1, \dots, N\}$ and $\tilde{t} \in g$ -th epoch and $t \neq 1 + \phi(g-1)$.

We study the event that player k does not select arm $\theta_k(t)$ at time t under two scenarios: (i) $A_k \neq \Omega_*^k(t)$, and (ii) $A_k = \Omega_*^k(t)$, where $\Omega_*^k(t)$ is the set with k best arms at time t . Then, we have

$$\begin{aligned} \sum_{t=1}^T \mathbf{1}_{\{s_k(t) \neq \theta_k(t), t \notin \mathcal{T}\}} &\leq \sum_{t=1}^T \mathbf{1}_{\{A_k \neq \Omega_*^k(t), t \notin \mathcal{T}\}} \\ &\quad + \sum_{t=1}^T \mathbf{1}_{\{s_k(t) \neq \theta_k(t), A_k = \Omega_*^k(t), t \notin \mathcal{T}\}}. \end{aligned} \quad (20)$$

Note that unlike RR-SW-UCB#, in SW-DLP, after initialization, $A_k(t)$ is computed every time instead of only at time instants $\{N + \eta M + 1\}_{\eta \in \mathbb{Z}_{\geq 0}}$. However, this difference does not change the order of the total number of times that $\Omega_*^k(t)$ is misidentified. Therefore, with the fact (19), it follows similarly to the proof of Lemma 1 that

$$\sum_{t=1}^T \mathbf{1}_{\{A_k \neq \Omega_*^k(t)\}} \in O(T^{\frac{1+\nu}{2}} \ln T). \quad (21)$$

The Chernoff-Hoeffding inequality is symmetric about the estimated mean and upper tail bound is identical to the lower tail bound. Hence, second term on the right hand side of inequality (20) that involves selecting $s_k(t)$ using lower confidence bounds can be bounded similarly to the first term. Thus, we have

$$\sum_{t=1}^T \mathbf{1}_{\{s_k(t) \neq \theta_k(t), A_k = \Omega_*^k(t), t \notin \mathcal{T}\}} \in O(T^{\frac{1+\nu}{2}} \ln T). \quad (22)$$

Substituting (21) and (22) into (20), and substituting (20) into (18), we conclude that $\hat{\mathcal{N}}_k \in O(T^{\frac{1+\nu}{2}} \ln T)$.

The number of times the group does not receive a reward from arm $\theta_k(t)$ is upper bounded by the number of times player k does not receive a reward from arm $\theta_k(t)$. Player k does not receive a reward from arm $\theta_k(t)$ if one of the following conditions is true (i) arm $\theta_k(t)$ is not selected by player k , and (ii) arm $\theta_k(t)$ is selected by another player $j \neq k$. The total number of times either one of these events occurs at any arm $\theta_k(t)$, for all $k \in \{1, \dots, M\}$, can be upper bounded by $\sum_{k=1}^M 2\hat{\mathcal{N}}_k$. Since $\hat{\mathcal{N}}_k \in O(T^{\frac{1+\nu}{2}} \ln T)$ for all $k \in \{1, \dots, M\}$, we conclude the proof. \blacksquare

Remark 1 (Comparison of RR-SW-UCB# and SW-DLP): In multi-player MAB algorithms, assignment of a player to a targeted arm is crucial to avoid collisions. In RR-SW-UCB#, the indices of arms and the indices of players are employed for this assignment. A round robin policy ensures all players select M -best arms persistently and have an accurate estimate of the associated mean rewards. While in

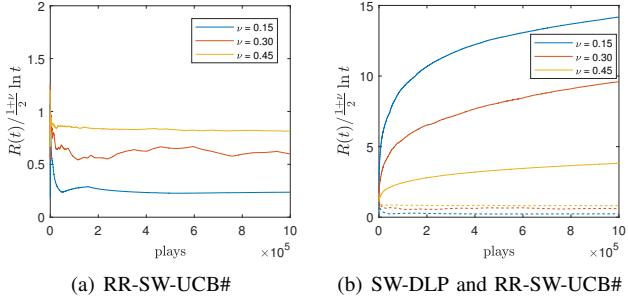


Fig. 1. The performance the two algorithms in the abruptly changing environment. (a) The solid lines show the performance of RR-SW-UCB# (b) The solid lines and dashed lines show the performance of SW-DLP and RR-SW-UCB#, respectively.

SW-DLP, such accurate estimation by all players is driven by the lower confidence bound based assignment of players to the arms. \square

V. NUMERICAL ILLUSTRATION

In this section, we present simulation results for RR-SW-UCB# and SW-DLP in the abruptly changing environment. In the simulations, we consider a multi-player MAB problem with 6 arms and 3 players. We consider three different values $\{0.15, 0.3, 0.45\}$ of parameter ν that describes the number of breakpoints to show the performance the both algorithms. The breakpoints are introduce at time instants where the next element of sequence $\{[t^\nu]\}_{t \in \{1, \dots, N\}}$ is different from current element. We pick them at these time instants to make number of breakpoints $\Upsilon_t \in O(t^\nu)$ uniformly for all $t \in \{1, \dots, T\}$. At each break point, the mean rewards at each arm is randomly selected from $\{0.05, 0.22, 0.39, 0.56, 0.73, 0.90\}$. In both algorithms, we select $\lambda = 12.3$.

As shown in Fig. 1, with either algorithm, the ratio of the empirical cumulative group regret to the order of $t^{1+\nu/2} \ln t$ is upper bounded by a constant. The dashed lines in Fig. 1 (b) are taken directly from (a). The comparison shows that the cumulative regret of RR-SW-UCB# is much lower than SW-DLP. However, if cost of switching between arms is considered, then the round-robin structure of RR-SW-UCB# would incur significant cost, and in such a scenario SW-DLP might be preferred.

VI. CONCLUSION AND FUTURE DIRECTIONS

We studied the multi-player stochastic MAB problem in abruptly environment under a collision model in which a player receives a reward by selecting an arm if it is the only player to select that arm. We designed two novel algorithms, RR-SW-UCB# and SW-DPL to solve this problem. We analyzed these algorithms and characterized their performance in terms of expected cumulative group regret. In particular, we showed that these algorithms incur sublinear expected cumulative regret, i.e., the time average of the regret asymptotically converges to zero.

There are several possible avenues for future research. In this paper, we focused on an abruptly changing environment.

Extension of this work to other classes of non-stationary environments such as slowly-varying environment is of considerable interest. Another avenue of future research is extension of these algorithms to the multi-player Markovian MAB problem.

REFERENCES

- [1] J. R. Krebs, A. Kacelnik, and P. Taylor, "Test of optimal sampling by foraging great tits," *Nature*, vol. 275, no. 5675, pp. 27–31, 1978.
- [2] V. Srivastava, P. Reverdy, and N. E. Leonard, "On optimal foraging and multi-armed bandits," in *Proceedings of the 51st Annual Allerton Conference on Communication, Control, and Computing*, Monticello, IL, USA, 2013, pp. 494–499.
- [3] ———, "Surveillance in an abruptly changing world via multiarmed bandits," in *IEEE Conference on Decision and Control*, 2014, pp. 692–697.
- [4] M. Y. Cheung, J. Leighton, and F. S. Hover, "Autonomous mobile acoustic relay positioning as a multi-armed bandit with switching costs," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Tokyo, Japan, November 2013, pp. 3368–3373.
- [5] A. Anandkumar, N. Michael, A. K. Tang, and A. Swami, "Distributed algorithms for learning and cognitive medium access with logarithmic regret," *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 4, pp. 731–745, 2011.
- [6] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire, "The non-stochastic multiarmed bandit problem," *SIAM Journal on Computing*, vol. 32, no. 1, pp. 48–77, 2002.
- [7] O. Besbes, Y. Gur, and A. Zeevi, "Optimal exploration-exploitation in a multi-armed-bandit problem with non-stationary rewards," *arXiv preprint arXiv:1405.3316*, 2014.
- [8] A. Garivier and E. Moulines, "On upper-confidence bound policies for non-stationary bandit problems," *arXiv preprint arXiv:0805.3415*, 2008.
- [9] F. Liu, J. Lee, and N. Shroff, "A change-detection based framework for piecewise-stationary multi-armed bandit problem," *arXiv preprint arXiv:1711.03539*, 2017.
- [10] L. Wei and V. Srivastava, "On abruptly-changing and slowly-varying multiarmed bandit problems," in *American Control Conference*, Milwaukee, WI, Jun. 2018, to appear, arXiv:1802.08380.
- [11] V. Anantharam, P. Varaiya, and J. Walrand, "Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays-part i: Iid rewards," *IEEE Transactions on Automatic Control*, vol. 32, no. 11, pp. 968–976, 1987.
- [12] ———, "Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays-part ii: Markovian rewards," *IEEE Transactions on Automatic Control*, vol. 32, no. 11, pp. 977–982, 1987.
- [13] Y. Gai and B. Krishnamachari, "Distributed stochastic online learning policies for opportunistic spectrum access," *IEEE Transactions on Signal Processing*, vol. 62, no. 23, pp. 6184–6193, 2014.
- [14] K. Liu and Q. Zhao, "Distributed learning in multi-armed bandit with multiple players," *IEEE Transactions on Signal Processing*, vol. 58, no. 11, pp. 5667–5681, 2010.
- [15] D. Kalathil, N. Nayyar, and R. Jain, "Decentralized learning for multiplayer multiarmed bandits," *IEEE Transactions on Information Theory*, vol. 60, no. 4, pp. 2331–2345, 2014.
- [16] N. Nayyar, D. Kalathil, and R. Jain, "On regret-optimal learning in decentralized multi-player multi-armed bandits," *IEEE Transactions on Control of Network Systems*, vol. 5, no. 1, pp. 597–606, 2016.
- [17] S. Shahrampour, A. Rakhlin, and A. Jadbabaie, "Multi-armed bandits in multi-agent networks," *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017.
- [18] P. Landgren, V. Srivastava, and N. E. Leonard, "On distributed cooperative decision-making in multiarmed bandits," in *European Control Conference*, Aalborg, Denmark, Jun. 2016, pp. 243–248.
- [19] ———, "Distributed cooperative decision-making in multiarmed bandits: Frequentist and Bayesian algorithms," in *IEEE Conference on Decision and Control*, Las Vegas, NV, USA, Dec. 2016, pp. 167–172.
- [20] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine Learning*, vol. 47, no. 2, pp. 235–256, 2002.
- [21] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *Journal of the American Statistical Association*, vol. 58, no. 301, pp. 13–30, 1963.