

Distributed Bandits: Probabilistic Communication on d -regular Graphs

Udari Madhushani and Naomi Ehrich Leonard

Abstract—We study the decentralized multi-agent multi-armed bandit problem for agents that communicate with probability over a network defined by a d -regular graph. Every edge in the graph has probabilistic weight p to account for the $(1-p)$ probability of a communication link failure. At each time step, each agent chooses an arm and receives a numerical reward associated with the chosen arm. After each choice, each agent observes the last obtained reward of each of its neighbors with probability p . We propose a new Upper Confidence Bound (UCB) based algorithm and analyze how agent-based strategies contribute to minimizing group regret in this probabilistic communication setting. We provide theoretical guarantees that our algorithm outperforms state-of-the-art algorithms. We illustrate our results and validate the theoretical claims using numerical simulations.

I. INTRODUCTION

Distributed learning benefits individual and group decision making in both natural and engineered systems. Communication constraints and security concerns often prevent the possibility of centralized information gathering and centralized decision making. This necessitates the development of effective decentralized communication protocols and decision-making strategies. Decentralized strategies allow individual agents to make decisions using locally available information, i.e., information that they observe or is communicated to them from their neighbors, as defined by a communication network. As a result, decision-making performance is strongly affected by the structure and the reliability of the network.

Consider a group of agents making sequential decisions in an uncertain environment. Each agent is faced with the same problem of repeatedly choosing an option from a given fixed set of options. After every choice, each agent receives a numerical reward drawn from a probability distribution associated with the chosen option. The objective of each agent is to maximize its individual cumulative reward while contributing to maximizing the group cumulative reward. The best strategy in this situation is to repeatedly choose the optimal option, i.e., the option that provides the maximum average reward. However, agents are unaware of the expected reward values of the options. Each individual is required to exploit the options that are known to provide high rewards and explore the lesser known options in order to identify the options that might potentially provide higher rewards. This is known as the explore-exploit dilemma in decision theory. In a networked setting, agents can learn the reward distributions with fewer explorations by sharing information.

This research has been supported in part by ONR grants N00014-18-1-2873 and N00014-19-1-2556 and ARO grant W911NF-18-1-0325.

U. Madhushani and N.E. Leonard are with Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA. {udarim, naomi}@princeton.edu

The multi-armed bandit (MAB) problem is a mathematical model developed to capture the salient features of the explore-exploit trade-off [1]–[3]. Initially, MAB was defined for a single agent making choices among multiple arms, equivalent to options, and receiving a reward after every choice. In the stochastic MAB problem, rewards are drawn from probability distributions. In the non-stochastic MAB problem, each option is assigned a deterministic sequence of rewards. In both problem settings, the goal is to maximize the expected cumulative reward. This is equivalent to minimizing the expected cumulative regret defined as the expected loss incurred by the agent through sampling sub-optimal options. It is well established that for the stochastic MAB the cumulative regret grows as $\Theta(\log T)$, [4] and for the non-stochastic MAB the cumulative regret grows as $\Theta(\sqrt{T})$, [5] where T is the time horizon of the decision-making process. In this paper we focus on stochastic MAB problems. The paper [2] proposed a family of algorithms to achieve a logarithmic upper bound for cumulative regret asymptotically in T . A lower bound that matches the order of the upper bound was also proposed, establishing the result that any asymptotically efficient algorithm samples suboptimal options at least logarithmically in T . The paper [4] proposed a family of Upper Confidence Bound (UCB) algorithms to achieve logarithmic cumulative regret uniformly in T .

The papers [6]–[11] studied distributed bandit problems with static communication protocols. A centralized multi-agent setting was considered in [6] and a decentralized multi-agent setting was considered in [7], [8] in which running consensus algorithms are used to update estimates and provide graph-structure-dependent performance measures. The papers [9]–[11] considered leader-follower settings in which leaders observe the rewards and choices of their followers and followers copy the choices of their leaders. Paper [9] also considered a decentralized bandit problem where agents observe rewards and choices of their neighbors without directly copying their choices.

The papers [12]–[16] considered decentralized bandit problems with time-varying communication. The papers [12]–[14] proposed communication protocols in which agents decide when to communicate depending on the decision-making process. The paper [12] considered an additional constraint in which agents cannot communicate and sample at the same time. The papers [15], [16] proposed stochastic communication protocols. The problem presented in our previous paper [15] is closely related to the problem we study in this paper. In [15] neighbors are defined according to a d -regular graph and each agent k observes each of its neighbors independently with probability p_k .

We provided an index, dependent on p_k 's, to rank agents according to performance. In [16] heterogeneous sampling rules were derived to improve group performance in multi-star networks.

We are motivated by real-world examples modeled by the bandit framework, such as recommender systems [17], user-targeted online advertising [18], [19] and clinical trials [20], in which, to protect the privacy of users or patients, it is desirable not to disclose the history of choices. We study a distributed bandit problem in which agents share information only about the last obtained reward with their neighbors defined according to a fixed d -regular network graph. We consider the case where each edge of the communication graph independently fails with probability $1 - p$. The edge failure can be either considered as an inherent imperfection of the communication links or a decision made by agents to share rewards at random time steps. We first propose a novel UCB based algorithm and then analyze the effect of degree d and edge weight probability p on group performance.

In Section II we formulate the problem mathematically and in Section III we propose a new UCB based algorithm. We analyze the performance of the proposed algorithm and prove theoretical results in Section IV. In Section IV-A we provide specialized results for the case where the communication network is a complete graph. In Section IV-B we compare theoretical results of our algorithm and state-of-the-art algorithms. In Section V we provide numerical simulations illustrating our results and validating theoretical guarantees. In Section V-A we show results for d -regular graphs. In Section V-B and V-C we compare our results with state-of-the-art algorithms. We provide additional discussions in Section VI and concluding remarks in Section VII.

II. PROBLEM FORMULATION

In this section we present the mathematical formulation of the problem. Consider a group of K agents faced with the same N -armed bandit problem for T time steps. The terms arms and options are used interchangeably. Let X_i be a sub-Gaussian random variable, with variance proxy σ_i^2 , that denotes the reward of option $i \in \{1, 2, \dots, N\}$. Sub-Gaussian distributions include Gaussian distributions, Bernoulli distributions, and bounded distributions, which are ubiquitous in real-world applications of the bandit framework. Define $\mathbb{E}(X_i) = \mu_i$ as the expected reward of option i . The option with maximum expected reward is the optimal option $i^* = \arg \max\{\mu_1, \dots, \mu_N\}$. Let $\Delta_i = \mu_{i^*} - \mu_i$ be the expected reward gap between option i^* and option i .

Let the random variable $\varphi_t^k \in \{1, 2, \dots, N\}$ denote the option chosen by agent k at time t . Let $\mathbb{I}_{\{\varphi_t^k=i\}}$ be the indicator random variable that takes value 1 if agent k chose option i at time t and 0 otherwise.

We define the communication network using a family of time-varying graphs denoted by $\{G_t(\mathcal{V}, \mathcal{E}_t)\}_{t=1}^T$. The graph is generated by an underlying fixed d -regular graph $G(\mathcal{V}, \mathcal{E})$ and by letting each edge fail independently with probability $1 - p$. Here each node $k \in \mathcal{V}$ represents an agent and each edge $e(k, j) \in \mathcal{E}$ represents the undirected

communication link between agents k and j . Let $\mathbb{I}_{\{k,j\}}^t$ be the random variable that takes value 1 if agent k observes the reward obtained by agent j at time t and 0 otherwise. Thus, $\mathbb{E}(\mathbb{I}_{\{k,j\}}) = p, \forall k, j$ if $e(k, j) \in \mathcal{E}$. Note that $\mathbb{E}(\mathbb{I}_{\{k,k\}}) = 1, \forall k$. We consider the case where the structure of the communication graph is not known to the agents. When the edge between agents k and j at time t fails, $\mathbb{I}_{\{k,j\}}^t = \mathbb{I}_{\{j,k\}}^t = 0$. Thus, at each time step the communication graph is undirected.

III. ALGORITHM

In this section we present the mathematical formulation of our algorithm. Let $\hat{\mu}_i^k(t)$ be the estimated mean of option i by agent k at time t . Define $n_i^k(t) = \sum_{\tau=1}^t \mathbb{I}_{\{\varphi_\tau^k=i\}}$ as the number of samples of option i obtained by agent k until time t . Let $N_i^k(t)$ be the sum of $n_i^k(t)$ and number of samples of option i communicated to agent k by its neighbors until time t . Then by definition $N_i^k(t) = \sum_{\tau=1}^t \sum_{j=1}^K \mathbb{I}_{\{\varphi_\tau^j=i\}} \mathbb{I}_{\{k,j\}}^\tau$.

The estimated mean value is calculated by taking the simple average of the sum of reward values of option i obtained by, and communicated to, agent k up to time t :

$$\hat{\mu}_i^k(t) = \frac{S_i^k(t)}{N_i^k(t)}$$

where $S_i^k(t) = \sum_{\tau=1}^t \sum_{j=1}^K X_i \mathbb{I}_{\{\varphi_\tau^j=i\}} \mathbb{I}_{\{k,j\}}^\tau$.

The goal of each agent is to maximize its individual cumulative reward while contributing to maximizing the group cumulative reward. In order to realize this goal, each agent employs an agent-based strategy that captures the trade-off between exploring and exploiting by constructing an objective function that strikes a balance between the estimation of the expected reward and the uncertainty associated with the estimate. We consider the case with unknown edge failure probability but known variance proxy as follows.

Assumption 1. We assume that agents know the total number of agents K and the variance proxy σ_i^2 of the rewards associated with each option.

The statement regarding the variance proxy in Assumption 1 is equivalent to knowing the variance when the reward distributions are Gaussian and knowing the upper and lower bounds when the reward distributions are bounded. We make the following assumption about the rewards obtained.

Assumption 2. When more than one agent chooses the same option at the same time they receive rewards independently drawn from the probability distribution associated with the chosen option.

Our proposed algorithm is the sampling rule for each agent k given in Definition 1.

Definition 1. (Sampling Rule) The sampling rule $\{\varphi_t^k\}_{t=1}^T$ for agent k at time $t \in \{1, \dots, T\}$ is defined as

$$\mathbb{I}_{\{\varphi_{t+1}^k=i\}} = \begin{cases} 1 & , i = \arg \max\{Q_1^k(t), \dots, Q_N^k(t)\} \\ 0 & , \text{o.w.} \end{cases}$$

with

$$Q_i^k(t) \triangleq \widehat{\mu}_i^k(t) + C_i^k(t) \quad (1)$$

$$C_i^k(t) \triangleq \sigma_i \sqrt{\frac{2 \log(t^{\xi+1}K)}{N_i^k(t)}} \quad (2)$$

where $\xi > 1$.

Note that the term $C_i^k(t)$ depends on the uncertainty associated with agent k 's estimate of mean reward of option i . At each time step, the agent chooses the option that maximizes the objective function: the corresponding sequence of choices addresses the trade-off between exploiting and exploring. According to the sample rule defined by Definition 1, each agent uses the information it receives from other agents, in addition to the information obtained through its own sampling, to compute the objective functions. Thus, through information sharing, agents reduce the uncertainty associated with estimates of expected reward of the options.

IV. PERFORMANCE ANALYSIS

In this section we analyze the performance of our algorithm. Recall that regret is the loss due to choosing a suboptimal option instead of the optimal option. The expected group cumulative regret can be given as

$$\mathbb{E}(R(T)) = \sum_{i=1}^N \sum_{k=1}^K (\mu_{i^*} - \mu_i) \mathbb{E}(n_i^k(T)). \quad (3)$$

Thus, the expected group cumulative regret can be minimized by minimizing the expected number of samples taken from suboptimal options.

We proceed to upper bound the expected number of samples obtained from suboptimal options as follows. Recall that $n_i^k(t)$ is the number of samples of option i obtained by agent k until time t . Thus for each suboptimal option i we have

$$\begin{aligned} \sum_{k=1}^K \mathbb{E}(n_i^k(T)) &= \sum_{k=1}^K \sum_{t=1}^T \mathbb{P}(\varphi_t^k = i) \\ &\leq \sum_{k=1}^K \sum_{t=1}^T \mathbb{P}(Q_i^k(t) > Q_{i^*}^k(t)). \end{aligned} \quad (4)$$

Let $\{\eta_i(t)\}_1^T$ be a family of nonnegative nondecreasing functions. Then from (4) we have

$$\begin{aligned} \sum_{k=1}^K \mathbb{E}(n_i^k(T)) &\leq \sum_{k=1}^K \sum_{t=1}^T \mathbb{P}(\varphi_t^k = i, N_i^k(t) \leq \eta_i(t)) \\ &\quad + \sum_{k=1}^K \sum_{t=1}^T \mathbb{P}(Q_i^k(t) > Q_{i^*}^k(t), n_i^k(t) > \eta_i(t)). \end{aligned} \quad (5)$$

From (1) in Definition 1 we have

$$\begin{aligned} \{Q_i^k(t) \geq Q_{i^*}^k(t)\} &\subseteq \{\mu_{i^*} < \mu_i + 2C_i^k(t)\} \\ \cup \{\widehat{\mu}_{i^*}^k(t) \leq \mu_{i^*} - C_{i^*}^k(t)\} &\cup \{\widehat{\mu}_i^k(t) \geq \mu_i + C_i^k(t)\}. \end{aligned} \quad (6)$$

Thus from (2) when $\eta_i(t) = \frac{8\sigma_i^2}{\Delta_i^2} \log(t^{\xi+1}K)$ we see that

$$\mathbb{P}(\mu_{i^*} < \mu_i + 2C_i^k(t), n_i^k(t) \geq \eta_i(t)) = 0. \quad (7)$$

Thus from (5)-(7) we get

$$\begin{aligned} \sum_{k=1}^K \mathbb{E}(n_i^k(T)) &\leq \sum_{k=1}^K \sum_{t=1}^T \mathbb{P}(\varphi_t^k = i, N_i^k(t) \leq \eta_i(t)) \\ &\quad + \sum_{k=1}^K \sum_{t=1}^T \mathbb{P}(\widehat{\mu}_{i^*}^k(t) \leq \mu_{i^*} - C_{i^*}^k(t)) \\ &\quad + \sum_{k=1}^K \sum_{t=1}^T \mathbb{P}(\widehat{\mu}_i^k(t) \geq \mu_i + C_i^k(t)). \end{aligned} \quad (8)$$

First, we upper bound the last two summation terms on the right hand side of (8) using the following lemma.

Lemma 1. For any $\xi > 1$, some $\zeta > 1$ and for $\sigma_i > 0$ in the uncertainty $C_i^k(t)$ given by (2), we have

$$\mathbb{P}(|\widehat{\mu}_i^k(t) - \mu_i| > C_i^k(t)) \leq \frac{1}{\log \zeta} \frac{\log((d+1)t)}{t^{\xi+1}K}.$$

Proof. Substituting $d_k = d$ in Lemma 1 of the paper [21] we have for any $\vartheta > 0$

$$\begin{aligned} \mathbb{P}\left(|\widehat{\mu}_i(t) - \mu_i| \geq \sqrt{\frac{\vartheta}{N_i(t)}}\right) \\ \leq \frac{\log((d+1)t)}{\log \zeta} \exp(-2\kappa\vartheta) \end{aligned} \quad (9)$$

where $\kappa = \frac{1}{\sigma_i^2 (\zeta^{\frac{1}{4}} + \zeta^{-\frac{1}{4}})^2}$.

Let $\vartheta = 2\sigma_i^2 \log(t^{\xi+1}K)$. Then the proof of Lemma 1 follows from (2) and (9). \square

Second, we upper bound the first term on the right hand side of (8): $\sum_{k=1}^K \sum_{t=1}^T \mathbb{P}(\varphi_t^k = i, N_i^k(t) \leq \eta_i(t))$ as follows. Recall that $N_i^k(t) = \sum_{\tau=1}^t \sum_{j=1}^K \mathbb{I}_{\{\varphi_\tau^j = i\}} \mathbb{I}_{\{k,j\}}$. Since the communication structure is independent from the decision-making process we have

$$(1 + pd) \sum_{k=1}^K \mathbb{E}(n_i^k(t)) \leq \sum_{k=1}^K \mathbb{E}(N_i^k(t)). \quad (10)$$

Since $N_i^k(t)$ is a nonnegative random variable $N_i^k(t) \leq \eta_i(t) \implies \mathbb{E}(N_i^k(t)) \leq \eta_i(t)$. Since $\eta_i(t)$ is a nondecreasing nonnegative function,

$$\sum_{k=1}^K \sum_{t=1}^T \mathbb{P}(\varphi_t^k = i, N_i^k(t) \leq \eta_i(t)) \leq \frac{K}{1 + pd} \eta_i(T). \quad (11)$$

Now we state the main theoretical result of the paper.

Theorem 1. Consider a group of K agents sharing information over a d -regular graph with an edge failure probability of $1 - p$ and applying the sampling rule of Definition 1. The expected group cumulative regret satisfies

$$\begin{aligned} \mathbb{E}(R(T)) &\leq \sum_{i=1}^N \frac{K}{1 + pd} \frac{8\sigma_i^2}{\Delta_i} ((\xi + 1) \log T + \log K) \\ &\quad + \sum_{i=1}^N \frac{2\Delta_i}{\xi^2 \log \zeta} (\xi^2 \log(d+1) + \xi \log(d+1) + 1). \end{aligned}$$

Proof. From (8), (11) and Lemma 1 we get

$$\sum_{k=1}^K \mathbb{E}(n_i^k(T)) \leq \frac{K}{1+pd} \eta_i(T) + \frac{2}{\log \zeta} \sum_{t=1}^T \frac{\log((d+1)t)}{t^{\xi+1}}. \quad (12)$$

Note that we have

$$\begin{aligned} \sum_{t=1}^T \frac{\log((d+1)t)}{t^{\xi+1}} &\leq \log(d+1) + \int_1^T \frac{\log(d+1)t}{t^{\xi+1}} dt \\ &\leq \log(d+1) + \frac{\log(d+1)}{\xi} + \frac{1}{\xi^2}. \end{aligned} \quad (13)$$

The proof of Theorem 1 follows from (3), (12) and (13). \square

Remark 1. When Δ_i is small and σ_i is large, the case where it is difficult to identify the optimal option, the term $\sum_{i=1}^N \frac{K}{1+pd} \frac{8\sigma_i^2}{\Delta_i} ((\xi+1) \log T + \log K)$ is dominant and the expected cumulative regret of the group decreases as $\frac{1}{1+pd}$ with increasing p . This implies, for any d -regular graph, that expected group cumulative regret decreases monotonically with increasing p ; equivalently, group performance increases monotonically with decreasing edge failure probability $1-p$. Further, this implies that rate of decay of regret with increasing probability p increases monotonically with increasing d . That is, the increasing performance with decreasing edge failure probability is accelerated with greater d . These observations, illustrated in Section V, justify our algorithm as well-defined in the sense that group performance improves with reduced failure probability and increased connectivity.

A. Complete Communication Graph

When the communication graph is complete $d = K - 1$. Thus, from Theorem 1, expected group cumulative regret is

$$\begin{aligned} \mathbb{E}(R(T)) &\leq \sum_{i=1}^N \frac{K}{1+p(K-1)} \frac{8\sigma_i^2}{\Delta_i} ((\xi+1) \log T + \log K) \\ &\quad + \sum_{i=1}^N \frac{2\Delta_i}{\xi^2 \log \zeta} (\xi^2 \log K + \xi \log K + 1). \end{aligned}$$

When $p = 1$ this matches the asymptotic regret bound $O(\log(TK))$, provided in papers [7] and [8].

B. Comparison with Previous Work

In [15] we considered a similar communication structure with degree homogeneity and probabilistic heterogeneity: the underlying communication structure is a d -regular graph and agent k observes each of its neighbors with probability p_k . If we let $p_k = p, \forall k$, then that problem reduces to the problem we study in this paper. In [15] the asymptotic expected group cumulative regret is

$$\mathbb{E}(R(T)) = O\left(K \left(\frac{1}{2} + \frac{1}{2\sqrt{1+p}}\right) \log T + K \log K\right).$$

The present algorithm provides better theoretical guarantees since, by Theorem 1, the same asymptotic group regret is

$$\mathbb{E}(R(T)) = O\left(\frac{K}{1+pd} \log(TK)\right).$$

V. SIMULATION RESULTS

In this section, we provide numerical simulation results illustrating performance of the proposed algorithm and the validity of the theoretical claims. For all the simulations presented in this section, we consider 20 agents ($K = 20$) and 10 options ($N = 10$). For all simulations, we provide results with 1000 time steps ($T = 1000$) using 1000 Monte Carlo simulations with $\xi = 1.01$.

For all simulations provided in Section V-A and V-B we let reward distributions be Gaussian. We consider the expected reward of the optimal option to be $\mu_{i^*} = 11$ and for all sub-optimal options $i \neq i^*$ to be $\mu_i = 10$. We let variance associated with all options i be $\sigma_i^2 = 1$. Since the expected reward gaps $\Delta_i = 1, i \neq i^*$, are equal to the variances $\sigma_i^2 = 1$, it is challenging to distinguish the optimal option from the sub-optimal options.

A. Validating Remark 1

In this section we provide simulation results in the case that observations received by an agent from its neighbors are noisy. We let communication noise be a zero mean Gaussian noise with variance 1. Figure 1 shows the expected group cumulative regret for different values of d when links of the communication graph independently fail with probability 0.5. The orange solid curve corresponds to the circle graph where $d = 2$. The blue dashed curve corresponds to $d = 10$. The green dotted curve corresponds to $d = 19$, which is the complete graph. These results illustrate that the expected group cumulative regret is logarithmic in time and decreases with increasing d .




Fig. 1. Expected group cumulative regret of 20 agents with d -regular underlying communication network and edge failure probability 0.5.




Fig. 2. Expected group cumulative regret of 20 agents with d -regular underlying communication network and edge failure probability $1 - p$.

Figure 2 shows how group regret changes with edge failure probability for $d = 2, d = 10$ and $d = 19$. The expected group cumulative regret at the end of the time horizon $T = 1000$ is plotted. Results show that for each graph expected cumulative group regret decreases monotonically, on average, with increasing p . Further notice that the rate of decay of expected group cumulative regret increases monotonically with increasing d . The results illustrate that group regret decreases with increasing p as $\frac{1}{1+pd}$ thus validating the statement in Remark 1.

B. Performance Comparison for Communication with Edge Failures

In this section we compare experimental results for our algorithm and the algorithm proposed in [15]. We use for both algorithms the tuning parameter $\xi = 1.01$ given in [15].




Fig. 3. Expected group cumulative regret for a group of 20 agents using the algorithm of this paper and the algorithm of paper [15] on d -regular underlying communication networks ($d = 10$ and $d = 19$) with $p = 1 - p = 0.5$ edge failure probability.

Figure 3 shows the expected group cumulative regret for two d -regular graphs, one with $d = 10$ and one with $d = 19$, i.e., the complete graph, and the case in which communication links fail with probability 0.5 ($p = 0.5$). The pink solid (dotted) curve corresponds to the case in which agents use the algorithm proposed in this paper when communicating on a $d = 10$ ($d = 19$)-regular graph. The black solid (dotted) curve corresponds to the case in which agents use the algorithm proposed in the paper [15] when communicating on a $d = 10$ ($d = 19$)-regular graph. Simulation results illustrate that the algorithm proposed in this paper performs better than the algorithm proposed in [15]. The performance gap between the two algorithms increases with decreasing d , i.e., when the number of neighbors decreases. This illustrates that the algorithm proposed in this paper performs significantly better than the algorithm proposed in [15] when the communication graph is sparser.

In Figure 4 we show that this comparative performance holds for all edge failure probability values. The figure shows how expected group cumulative regret varies with edge failure probability $1 - p$. The expected group cumulative regret at the end of the time horizon $T = 1000$ is plotted against probability p . The results illustrate that the algorithm proposed in this paper outperforms the algorithm proposed in the paper [15] under all edge failure probability values.




Fig. 4. Expected group cumulative regret for a group of 20 agents using the algorithm of this paper and the algorithm of paper [15] on d -regular underlying communication networks ($d = 10$ and $d = 19$) with $1 - p$ edge failure probability. Results are provided for time $T = 1000$.

C. Performance Comparison for Communication without Edge Failures

In this section we provide a comparison of results when edge failure probability is zero ($p = 1$). We compare the performance of our algorithm with the performance of the UCB-Network algorithm proposed in the paper [9]. The UCB-Network algorithm does not contain any tuning parameters. We consider reward distributions to be bounded in the range $[0, 1]$. Recall that bounded variables are sub-Gaussian. Let $X \in [a, b]$ be a bounded random variable with mean μ and $b > a \geq 0$. Then $X - \mu$ is sub-Gaussian with variance proxy $\frac{(b-a)^2}{4}$. Thus bounded reward distributions considered in this section are sub-Gaussian with variance proxy $1/4$. We consider the reward distribution of the optimal arm as a triangular distribution with mode 1 and the reward distribution of all the suboptimal arms as a triangular distribution with mode 0.




Fig. 5. Expected group cumulative regret for a group of 20 agents using the algorithm of this paper and the UCB-Network algorithm [9] on d -regular underlying communication networks ($d = 10$ and $d = 19$) and no edge failures.

Figure 5 shows simulation results for expected group cumulative regret for 20 agents and two d -regular graphs: $d = 10$ and $d = 19$. Results for the agents using the algorithm of this paper when communicating on a $d = 10$ ($d = 19$)-regular graph are given in a pink solid (dotted) curve. The blue solid (dotted) curve shows results agents using the UCB-Network

algorithm of [9] when communicating on a $d = 10$ ($d = 19$)-regular graph. Simulation results illustrate that our algorithm outperforms UCB-Network. The performance gap between the two algorithms increases with decreasing d , i.e., when the communication graph becomes sparser our algorithm performs increasingly better than UCB-Network.

VI. DISCUSSION

In this section we discuss the applicability of our framework to social learning. In social learning willingness to share information with neighbors has been explained in evolutionary biology using concepts like reciprocity and intra-specific mutualism [22]. Foraging animals use signalling as a method of communication to track food resources [22], [23]. The paper [23] considered how the bandit framework can be used to model social foraging. In this setting it is reasonable to assume that agents (animal or human) can only communicate their rewards and actions (not estimates). This is because each agent can observe rewards and action of its neighbors or each agent can signal about its rewards and actions to its neighbors. The probability p can be considered as the social effort made by each agent in collective learning. Interpretation of p can be such that at each time step agents observe their neighbors independently with probability p or each agent broadcasts its rewards to its neighbors independently with probability p . Recall that when we consider p as the edge failure probability of the underlying d -regular graph G , the resulting communication graphs $\{G_t\}_{t=1}^T$ are undirected. In the setting discussed in this section agent k not observing (broadcasting) reward value and action of (to) agent j at time t does not necessarily mean that agent j is not observing (broadcasting) reward value and action of (to) agent k , and this yields a directed communication graphs $\{G_t\}_{t=1}^T$. All the results provided in this paper hold for directed communication graphs $\{G_t\}_{t=1}^T$ generated by this alternative implementation of p . This illustrates that our framework is applicable to a range of real-world applications from machine intelligence to evolutionary biology.

VII. CONCLUSIONS

In this paper, we studied the decentralized stochastic bandit problem with probabilistic communication failures. We proposed a new UCB based algorithm to minimize individual expected cumulative regret while contributing to minimizing expected group cumulative regret. We analyzed how agent-based strategies minimize the overall expected regret of the group. We provided improved analytical bounds for the expected group cumulative regret when the underlying communication network is a d -regular graph and the communication links independently fail with probability $1 - p$. We proved that for any d -regular graph expected group cumulative regret monotonically decreases with decreasing $1 - p$ (increasing p) as $\frac{1}{1+pd}$. We proved that when the communication structure is a complete graph with no edge failures, group regret incurred by the proposed algorithm is logarithmic in time and in number of agents. We illustrated the validity of theoretical results using numerical simulations.

REFERENCES

- [1] H. Robbins, "Some aspects of the sequential design of experiments," *Bulletin of the American Mathematical Society*, vol. 58, no. 5, pp. 527–535, 1952.
- [2] T. L. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Advances in Applied Mathematics*, vol. 6, no. 1, pp. 4–22, 1985.
- [3] T. Lattimore and C. Szepesvári, *Bandit Algorithms*. Cambridge University Press, 2020.
- [4] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine Learning*, vol. 47, no. 2-3, pp. 235–256, 2002.
- [5] S. Bubeck and N. Cesa-Bianchi, "Regret analysis of stochastic and nonstochastic multi-armed bandit problems," *arXiv preprint arXiv:1204.5721*, 2012.
- [6] 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.
- [7] D. Martínez-Rubio, V. Kanade, and P. Rebeschini, "Decentralized cooperative stochastic bandits," in *Advances in Neural Information Processing Systems*, 2019, pp. 4531–4542.
- [8] P. Landgren, V. Srivastava, and N. E. Leonard, "Distributed cooperative decision making in multi-agent multi-armed bandits," *Automatica*, vol. 125, p. 109445, 2021, and at arXiv:2003.01312.
- [9] R. K. Kolla, K. Jagannathan, and A. Gopalan, "Collaborative learning of stochastic bandits over a social network," *IEEE/ACM Transactions on Networking*, vol. 26, no. 4, pp. 1782–1795, 2018.
- [10] P. Landgren, V. Srivastava, and N. E. Leonard, "Social imitation in cooperative multiarmed bandits: partition-based algorithms with strictly local information," in *IEEE Conference on Decision and Control (CDC)*, 2018, pp. 5239–5244.
- [11] P.-A. Wang, A. Proutiere, K. Ariu, Y. Jedra, and A. Russo, "Optimal algorithms for multiplayer multi-armed bandits," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 4120–4129.
- [12] M. Chakraborty, K. Y. P. Chua, S. Das, and B. Juba, "Coordinated versus decentralized exploration in multi-agent multi-armed bandits," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2017, pp. 164–170.
- [13] U. Madhushani and N. E. Leonard, "Distributed learning: Sequential decision making in resource-constrained environments," *arXiv:2004.06171*, 2020.
- [14] —, "A dynamic observation strategy for multi-agent multi-armed bandit problem," in *19th European Control Conference (ECC)*, 2020, pp. 1677–1682.
- [15] —, "Heterogeneous stochastic interactions for multiple agents in a multi-armed bandit problem," in *18th European Control Conference (ECC)*, 2019, pp. 3502–3507.
- [16] —, "Heterogeneous explore-exploit strategies on multi-star networks," *IEEE Control Systems Letters*, vol. 5, no. 5, pp. 1603–1608, 2021.
- [17] R. Warlop, A. Lazaric, and J. Mary, "Fighting boredom in recommender systems with linear reinforcement learning," in *Advances in Neural Information Processing Systems*, 2018, pp. 1757–1768.
- [18] A. Durand, C. Achilleos, D. Iacovides, K. Strati, G. D. Mitsis, and J. Pineau, "Contextual bandits for adapting treatment in a mouse model of de novo carcinogenesis," in *3rd Machine Learning for Healthcare Conference (MLHC)*, 2018, pp. 67–82.
- [19] R. Féraud, R. Alami, and R. Laroche, "Decentralized exploration in multi-armed bandits," *arXiv preprint arXiv:1811.07763*, 2018.
- [20] A. C. Tossou and C. Dimitrakakis, "Algorithms for differentially private multi-armed bandits," in *Thirtieth AAAI Conference on Artificial Intelligence*, 2016, p. 2087–2093.
- [21] U. Madhushani and N. Leonard, "It doesn't get better and here's why: A fundamental drawback in natural extensions of ucb to multi-agent bandits," in *"I Can't Believe It's Not Better!" NeurIPS 2020 workshop*, 2020. [Online]. Available: <https://openreview.net/pdf?id=eK034ng005Y>
- [22] C. J. Torney, A. Berdahl, and I. D. Couzin, "Signalling and the evolution of cooperative foraging in dynamic environments," *PLoS Computational Biology*, vol. 7, no. 9, 2011.
- [23] T. Keasar, E. Rashkovich, D. Cohen, and A. Shmida, "Bees in two-armed bandit situations: foraging choices and possible decision mechanisms," *Behavioral Ecology*, vol. 13, no. 6, pp. 757–765, 2002.