

# Flooding with Absorption: An Efficient Protocol for Heterogeneous Bandits over Complex Networks

Junghyun Lee   

Kim Jaechul Graduate School of AI, KAIST, Seoul, Republic of Korea

Laura Schmid  

Kim Jaechul Graduate School of AI, KAIST, Seoul, Republic of Korea

Se-Young Yun   

Kim Jaechul Graduate School of AI, KAIST, Seoul, Republic of Korea

---

## Abstract

Multi-armed bandits are extensively used to model sequential decision-making, making them ubiquitous in many real-life applications such as online recommender systems and wireless networking. We consider a multi-agent setting where each agent solves their own bandit instance endowed with a different set of arms. Their goal is to minimize their group regret while collaborating via some communication protocol over a given network. Previous literature on this problem only considered arm heterogeneity and networked agents separately. In this work, we introduce a setting that encompasses both features. For this novel setting, we first provide a rigorous regret analysis for a standard flooding protocol combined with the classic UCB policy. Then, to mitigate the issue of high communication costs incurred by flooding in complex networks, we propose a new protocol called Flooding with Absorption (FwA). We provide a theoretical analysis of the resulting regret bound and discuss the advantages of using FwA over flooding. Lastly, we experimentally verify on various scenarios, including dynamic networks, that FwA leads to significantly lower communication costs despite minimal regret performance loss compared to other network protocols.

**2012 ACM Subject Classification** Theory of computation → Multi-agent learning; Theory of computation → Sequential decision making; Theory of computation → Regret bounds; Networks → Network protocol design

**Keywords and phrases** multi-armed bandits, multi-agent systems, collaborative learning, network protocol, flooding

**Digital Object Identifier** 10.4230/LIPIcs.OPODIS.2023.20

## Supplementary Material

*Software (Source Code):* <https://github.com/nick-jhlee/bandits-network>, archived at `swb:1:dir:41fce125d3e4736eb81fb7771bbf5e795f916c8e`

**Funding** JH and SY were supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grants funded by the Korean government(MSIT) (No.2022-0-00311, Development of Goal-Oriented Reinforcement Learning Techniques for Contact-Rich Robotic Manipulation of Everyday Objects; No.2019-0-00075, Artificial Intelligence Graduate School Program(KAIST)). LS was supported by the Stochastic Analysis and Application Research Center (SAARC) under National Research Foundation of Korea grant NRF-2019R1A5A1028324.

**Acknowledgements** The authors thank Ulrich Schmid (TU Wien) and the anonymous reviewers for helpful comments and suggestions.

## 1 Introduction

Exploration-exploitation dilemmas form the basis of many real-life decision-making tasks [51, 41]. In fact, the trade-off between making a choice to either stay with a current action or explore new possibilities appears as a feature in a variety of well-known applications [6, 52].



© Junghyun Lee, Laura Schmid, and Se-Young Yun;  
licensed under Creative Commons License CC-BY 4.0

27th International Conference on Principles of Distributed Systems (OPODIS 2023).

Editors: Alysson Bessani, Xavier Défago, Junya Nakamura, Koichi Wada, and Yukiko Yamauchi; Article No. 20;  
pp. 20:1–20:25



Leibniz International Proceedings in Informatics  
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany


As a result, the *multi-armed bandit (MAB)* problem, which is designed to reflect dilemmas of this kind, has been intensely studied in a wide range of scenarios [9, 37, 39]. In the baseline setting, an agent must make sequential decisions by choosing from a set of possible actions (the “arms” of the bandit). In the setting of stochastic MABs, each arm gives rewards following an unknown probability distribution [9]. Here, the goal is to minimize the cumulative *regret* over some timespan  $T$ , i.e., the difference between the accumulated reward and the reward that arises from choosing only the best arm. To reach this goal, the agent must balance exploring new actions and choosing already tested ones [2, 37, 9].

For applications involving large-scale decentralized decision-making [38], such as online advertising, search/recommender systems, and wireless channel allocation, collaborative multi-agent multi-armed bandits are a natural modeling choice [1, 58, 43, 44, 12, 3, 32, 67, 42]. In this setting, each agent plays their own bandit instance and communicates some information to others to minimize the *group* regret. One common assumption in the literature is that agents share the same set of arms [34, 22, 65]. However, arm homogeneity does not hold in many large-scale systems (e.g., contextual recommender systems), where agents often have heterogeneous arm sets of available actions [68, 15]. For instance, in a distributed recommender system scenario, arms might correspond to the contents shown to users, such as movies, and rewards to user opinions; then, depending on one’s location, the set of available movies for each agent may be different due to external constraints such as copyright issues. In this case, it would be desirable for the service provider if all the individual systems with partially overlapping contents collaborate with one another to minimize the group regret.

Another common assumption is that agents are connected by a complete network, where agents can directly communicate with every other agent [10, 68]. However, in real large-scale systems, agents are usually connected via a multi-hop communication network, where only adjacent nodes can exchange messages. To disseminate information to agents at larger distances here, agents need to forward messages. This is typically done by means of a flooding protocol [22, 49, 34, 65] or gossiping protocol [59, 56, 14, 63, 15]; that is, a received message is forwarded to all neighbors or only one randomly selected neighbor, respectively. Such a modeling assumption is ubiquitous in many real-life scenarios involving mobile/vehicular networks or social networks, among others.

*Contributions.* To the best of our knowledge, the setting of collaborative, *heterogeneous* multi-agent multi-armed bandits communicating over a general *network* has not been investigated yet. Yet, this setting arises in a wide range of real-life applications, e.g. wireless channel allocation, where not all nodes on an underlying network can access the same channels. For such scenarios, one can neither assume fully connected or particularly regular communication topologies, nor homogeneous arm sets. Distinct from previous work on multi-agent bandits, we here aim at efficient *network protocol design* for this novel setting; specifically, we want to design a *simple* alternative to flooding that achieves low communication complexity while retaining minimal loss in regret. Note that our research objective forces us to consider the bandit instances and the communication protocol in an integrated fashion, which is in stark contrast to the approaches used both in the existing bandit and networking literature, and thus contributes to the novelty of our approach.

To address the significant issue of exploding communication complexity of flooding in our setting, we introduce a new lightweight communication protocol for complex networks, called FLOODING WITH ABSORPTION (FWA). Its design principle is inherently coupled with the given bandit instances. We provide theoretical and experimental results showing that this protocol is highly communication-efficient on a wide range of complex networks, yet induces minimal regret performance loss for complex network topologies, even compared



■ **Figure 1** Communication network and arm heterogeneity.

to standard flooding. Using FWA can also help avoid heavy individual link congestion in complex networks. An important practical advantage of our protocol is that it is fully agnostic to the network structure, and can therefore be deployed on dynamically changing networks without any need for fine-tuning.

## 2 System Model

We now describe our setting of collaborative<sup>1</sup> *heterogeneous* multi-agent multi-armed stochastic bandits over a general communication network; see Figure 1 for an illustration. We assume that there are  $N$  agents connected by an undirected graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , with  $|\mathcal{V}| = N$ . We denote by  $\mathcal{N}_{\mathcal{G}}(v)$  be the neighborhood of  $v$  in  $\mathcal{G}$  *not* including  $v$ , and by  $\mathcal{G}[S]$  the induced subgraph for  $S \subset \mathcal{V}$ . Also, for an integer  $\gamma \geq 1$ , the  $\gamma$ -th order graph power of  $\mathcal{G}$ , denoted as  $\mathcal{G}^\gamma = (\mathcal{V}, \mathcal{E}^\gamma)$ , is defined as the graph on  $\mathcal{V}$  such that  $\{v, w\} \in \mathcal{E}^\gamma$  iff  $d_{\mathcal{G}}(v, w) \leq \gamma$ , where  $d_{\mathcal{G}}(v, w)$  is the length of the shortest path in  $\mathcal{G}$  connecting  $v$  and  $w$ . Each agent  $v \in \mathcal{V}$  has access to a finite set  $\mathcal{K}_v$  of arms with cardinality  $K_v$  that they can pull; let  $\mathcal{K} = \cup_{v \in \mathcal{V}} \mathcal{K}_v$  be the total set of arms of cardinality  $K$ .

Following [39], let  $\mathcal{M}_\sigma$  be a set of  $\sigma$ -sub-Gaussian distributions, and let  $\mu : \mathcal{M}_\sigma \rightarrow \mathbb{R}$  be a function mapping each (reward) distribution to its mean. Each arm  $a \in \mathcal{K}$  is associated with an unknown reward distribution  $P_a \in \mathcal{M}_\sigma$ . For simplicity, let  $\mu_a := \mu(P_a)$ . We note that  $P_a$  is independent of the agents' identities, i.e., each agent  $v$ , regardless of their arm set  $\mathcal{K}_v$ , faces the same distribution of rewards for the same arm  $a$  (whenever  $\mathcal{K}_v$  contains  $a$ ), and receives an i.i.d. reward from this distribution upon pulling this arm. We denote by  $a_\star^v$  the best *local* arm for agent  $v$  that satisfies  $\mu_{a_\star^v} > \mu_a$  for all  $a \in \mathcal{K}_v \setminus \{a_\star^v\}$ , and set  $\mu_\star^v = \mu_{a_\star^v}$ . The main challenge in the regret analysis is that even for the same arm  $a$ , the suboptimality gap may be different across agents containing  $a$ .

The execution of all agents proceeds in a sequence of synchronous rounds  $t = 1, 2, \dots$ . In each round  $t$ , all agents simultaneously (i) pull some arm, (ii) compute and send a message to their neighbors, and (iii) receive and process all messages from their neighbors. From the perspective of agents, let us denote by  $\mathcal{V}_a = \{v \in \mathcal{V} : a \in \mathcal{K}_v\} \subseteq \mathcal{V}$  the set of agents having action  $a$ , and let  $\mathcal{V}_{-a} \subseteq \mathcal{V}_a$  be the set of agents containing  $a$  as a *suboptimal* arm, i.e.,  $a \neq a_\star^v$ .

<sup>1</sup> We remark that we do *not* consider any collisions [46, 40, 65] where two neighbors pulling the same arm do not affect their observed rewards in any way. Rather, we focus on the collaborative setting where the agents are encouraged to cooperate with one another by sharing their own observations.

As done in the classic work on regret minimization in collaborative multi-agent bandits [34, 68, 49], our goal is to minimize the expected *group* regret at time horizon  $T$ ,  $\mathbb{E}[R(T)]$ , defined as:

$$\mathbb{E}[R(T)] := \sum_{v \in \mathcal{V}} \mathbb{E}[R^v(T)], \quad \mathbb{E}[R^v(T)] := \sum_{a \in \mathcal{K}_v} \Delta_a^v \mathbb{E}[N_a^v(T)], \quad (1)$$

where  $\Delta_a^v := \mu_\star^v - \mu_a$  is the agent-specific gap of arm  $a$  and  $N_a^v(t)$  is the number of times agent  $v$  plays the arm  $a$  up to time  $t$ .

We note that due to the two sources of heterogeneity in our setting, the benefit one receives from collaboration differs; furthermore, the same arm  $a$  may be optimal for some agents but suboptimal for others. Of course, even so, we expect that agents communicating and collaborating should lead to a speed-up (in terms of the group regret) compared to the baseline without any information sharing. The question is then how much speed-up one could get from collaboration under our heterogeneous system model. An additional prominent issue for realistic applications and, in particular, *complex* networks concerns the communication complexity of the information-sharing protocol used. Designing communication protocols with *low* communication complexity, defined as the number of messages sent and forwarded, is of paramount importance to not overshadow the improvement in regret due to collaboration.

### 3 Flooding

As is common in much of the previous literature, we will focus on agents that individually run the classic upper-confidence bound (UCB) policy [2, 34, 68] under which each agent pulls the arm associated with the maximum of the so-called UCB index, which is the sum of empirical reward (up to time  $t$ ) and an additional bonus term. Agents can, and should, take advantage of the distributed setting by communicating pulled arms and received rewards amongst each other over the underlying communication network. This is not straightforward to implement or analyze, given that the agents' arm sets are all different, and thus, even when broadcasting over multiple hops by flooding the network is available, the speed-up gained from this is not immediate. In this section, we review (and reformulate) the standard flooding protocol for our setting and discuss its properties.

#### 3.1 Flooding

FLOODING (broadcasting) lets each agent send all messages to all its neighbors in every round, with the number of times the message is forwarded limited by the *time-to-live* (TTL)  $\gamma$  [13, 60, 64, 55]. To account for potential loops in the network and avoid a broadcast storm [61, 66], we explicitly use a sequence number-controlled flooding (SNCF) variant. The pseudocode of FLOODING along with each agent's UCB algorithms is presented in Algorithm 1 with *absorb* = *False*.

FLOODING proceeds as follows: In each round  $t$ , each agent  $v$  pulls an arm  $a^v(t)$  that has the highest upper confidence bound (line 8). Note that  $M_a^v(t)$  is the number of pulls of arm  $a$  available to agent  $v$  by time  $t$ , and  $\hat{\mu}_a^v$  is the estimate of  $\mu_a$  made by agent  $v$  at time  $t$ . In both estimates, agent  $v$  uses all observations available to  $v$  by time  $t$ , including the messages relayed to them. Having received the corresponding reward  $X_{a^v(t)}^v(t)$  from pulling arm  $a^v(t)$ , agent  $v$  creates a message

$$m = \langle a^v(t), X_{a^v(t)}^v(t), \text{HASH}(v, a^v(t), X_{a^v(t)}^v(t)), v, \gamma \rangle, \quad (2)$$

and pushes it to  $\mathcal{M}^v$ , its current queue of messages to be sent (line 11). After UCB has been completed, each agent sends and receives messages to and from its neighbors (lines 14–21).

---

**Algorithm 1** FLOODING, FWA WITH UCB.

---

**Input:**  $\gamma \in \mathbb{N}$ ,  $\alpha > 0$ ,  $\text{absorb} \in \{\text{True}, \text{False}\}$ ,  
 $f: \mathbb{N} \rightarrow [1, \infty)$

```

1 Initialize  $\mathcal{M}^v = \text{queue}()$ ,  $\mathcal{H}^v = \text{queue}(\text{maxlen} = \gamma N)$ ;
2  $M_a^v(0) = 0, \hat{\mu}_a^v = 0$  for all  $(v, a) \in \mathcal{V} \times \mathcal{K}$ ;
3 for  $t \in [T]$  do
4   for  $v \in \mathcal{V}$  do
5     if  $t \leq K_v$  then
6       play the  $t$ -th arm of  $\mathcal{K}_v$ ;
7     else
8       play the arm  $a^v(t)$  chosen as follows:
9       
$$a^v(t) = \arg \max_{a \in \mathcal{K}_v} \left\{ \hat{\mu}_a^v(t) + \sqrt{\frac{2 \log f(t)}{M_a^v(t)}} \right\};$$

10      end
11       $\text{hash} \leftarrow \text{HASH}(v, a^v(t), X^v(t))$ ;
12       $\mathcal{M}^v.\text{push}(\langle a^v(t), X^v(t), \text{hash}, v, \gamma \rangle)$ ;
13    end
14    for  $v \in \mathcal{V}$  do
15      while  $\widetilde{\mathcal{M}}^v$  is not empty do
16         $m \leftarrow \widetilde{\mathcal{M}}^v.\text{pop}()$ ;
17        for  $w \in \mathcal{N}_G(v) \setminus \{m(3)\}$  do
18           $m' \leftarrow \text{copy}(m)$ ;
19          if  $m'(2) \notin \mathcal{H}^w$  then
20             $\mathcal{H}^w.\text{push}(m'(2))$ ;
21             $\text{Receive}(w, m', \text{absorb})$ ;
22          end
23          delete  $m'$ 
24        end
25      end
26    end
27 end

```

---

**Algorithm 2** Subroutines.

---

```

1 Function  $\text{Receive}(w, m, \text{absorb})$ :
2   if  $m(0) \in \mathcal{K}_w$  then
3      $M_{m(0)}^w(t) \leftarrow M_{m(0)}^w(t-1) + 1$ ,
4      $\hat{\mu}_{m(0)}^w(t) \leftarrow \frac{M_{m(0)}^w(t-1)\hat{\mu}_{m(0)}^w(t-1) + m(1)}{M_{m(0)}^w(t)}$ ;
5     if  $\text{absorb}$  is  $\text{True}$  then
6       delete  $m$ ;
7     else
8        $\text{Store}(w, m)$ ;
9     end
10  else
11     $\text{Store}(w, m)$ ;
12  end
13 Function  $\text{Store}(w, m)$ :
14    $m(4) \leftarrow m(4) - 1$ ;
15   if  $m(4) < 0$  then
16     delete  $m$ ;
17   else
18      $m(3) \leftarrow w$ ;
19      $\mathcal{M}^w.\text{push}(m)$ ;
20   end

```

---

Our message  $m = \langle m(0), m(1), m(2), m(3), m(4) \rangle$  consists of the following components:  $m(0)$  and  $m(1)$  are the arm pulled by agent  $v$  and the reward received at time  $t$ , respectively.  $m(2)$  is a hash value of the originating agent  $v$ , the arm pulled, and the obtained reward that acts as a unique identifier of the message. Our protocol uses  $m(2)$  to control flooding by avoiding routing loops that can lead to broadcast storms and improper bias in the reward estimations (the estimation protocol is shown in line 3 of Algorithm 2). Each agent  $v$  keeps track of the hash values of messages that they have seen until time  $t$  via a *queue*<sup>2</sup> of size  $\gamma N$ , denoted as  $\mathcal{H}^v$ . If an already-seen message comes in (line 18), that message is deleted on arrival (line 22). The memory length of  $\gamma N$  is the worst-case space complexity that arises from keeping track of all messages from all agents for the last  $\gamma$  time steps, as all messages can be forwarded at most  $\gamma$  times.  $m(3)$  is the agent that last forwarded the message; if the receiver  $w$  passes on  $m$ , they replace  $m(3)$  with  $w$  and forward  $m$  to all neighbors except the originator  $m(3)$  (line 16). This prevents messages from echoing after one hop.  $m(4)$  keeps track of the remaining life span of the message (TTL), which is initialized to  $\gamma$ . It is decayed by 1 every time a message is forwarded, and the message is discarded when TTL reaches 0. We note that  $\gamma = 1$  is equivalent to INSTANT REWARD SHARING (IRS) [34], where each agent only sends its message to its neighbors, and any message containing arm  $a$  that is sent to agents not containing  $a$  becomes void.

We assume that the nodes have no knowledge of the network topology, or who their neighbors are, which is a realistic assumption in complex networks, and wireless networks. This is in contrast to some of the previous works, e.g., [49], which assumes that each agent

---

<sup>2</sup> The queue operations used in the algorithm (*pop()*, *push()*, *flush()*) are defined as usual [21].

knows its neighborhood in  $\mathcal{G}^\gamma$ , which essentially bypasses any issues regarding communication complexity. Specifically, this also abstracts away any difficulties arising from delayed messages traveling along different paths.

### 3.2 Group Regret Analyses of Flooding

The regret bound of stochastic bandits generally depends on a problem-dependent quantity [2] that quantifies the difficulty of learning. For instance, for a single-agent multi-armed stochastic bandit, the regret bound scales as  $\sum_{a \in \mathcal{K} \setminus \{a^*\}} \frac{\log T}{\Delta_a}$ , where  $\Delta_a$  is the gap between the mean rewards of best arm  $a^*$  and the suboptimal arm  $a$ . The intuition is that if the mean rewards are similar, one would need a much tighter confidence interval to identify the optimal arm, forcing one to pull  $a$  many more times, precisely inversely proportional to its reward gap. This is known to be asymptotically optimal [37].

In our setting, we would expect our problem-dependent quantity to depend on both the underlying network topology and arm distribution. To see this, given a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , we first recall some graph-theoretic quantities [7]:

► **Definition 3.1.** *The **clique covering number** of  $\mathcal{G}$ , denoted as  $\theta(\mathcal{G})$ , is the smallest size of a partition of  $\mathcal{V}$  such that each part induces a clique. Any such partition (not necessarily minimum) is called a **clique cover**. The **independence number** of  $\mathcal{G}$ , denoted as  $\alpha(\mathcal{G})$ , is the maximum size of a subset of  $\mathcal{V}$  that induces no edges.*

We now define our problem-dependent quantity  $\Delta_a^\gamma$  as follows:

$$\Delta_a^\gamma := \max_{\mathcal{C}_{-a}} \left( \sum_{C \in \mathcal{C}_{-a}} \left\{ \frac{2}{\min_{v \in C} \Delta_a^v} - \frac{1}{\max_{v \in C} \Delta_a^v} \right\} \right)^{-1}, \quad (3)$$

where  $\Delta_a^v = \mu_*^v - \mu_a$  is the agent-specific suboptimality gap of arm  $a$ , and  $\max_{\mathcal{C}_{-a}}$  is taken over all possible clique covers  $\mathcal{C}_{-a}$  of  $[\mathcal{G}^\gamma]_{-a} := \mathcal{G}^\gamma[\mathcal{V}_{-a}]$ . If  $\mathcal{V}_{-a} = \emptyset$ , we set  $\Delta_a^\gamma = 0$ .

We now present the nonasymptotic regret upper bound for FLOODING:

► **Theorem 3.2.** *Algorithm 1 with `absorb = False`,  $f(t) = t^\alpha$ ,  $\alpha > \max\left(\frac{1}{2}, \frac{2\sigma^2}{\gamma+1}\right)$ , and  $\gamma \in \{1, \dots, \text{diam}(\mathcal{G})\}$  achieves the group regret upper bound*

$$\mathbb{E}[R(T)] \leq \sum_{\substack{a \in \mathcal{K} \\ \Delta_a^\gamma > 0}} \frac{4\alpha \log T}{\Delta_a^\gamma} + b(\gamma) + \sum_{a \in \mathcal{K}} f_a(\gamma), \quad (4)$$

where

$$b(\gamma) := \left( \frac{\alpha + 1/2}{\alpha - 1/2} \right)^2 \frac{8(\gamma + 1)}{\log \frac{(\gamma+1)(\alpha+1/2)}{4\sigma^2}} \sum_{a \in \mathcal{K}} \sum_{v \in \mathcal{V}_a} \Delta_a^v, \quad f_a(\gamma) = \sum_{v \in \mathcal{V}_{-a}} \Delta_a^v \min \left( 2\gamma, \frac{4\alpha \log T}{(\Delta_a^v)^2} \right).$$

The complete proof, deferred to Appendix B, uses a clique covering argument and an Abel transformation. See Appendix C.1 for a discussion of the main technical challenges when proving the regret bound for our setting, compared to previously considered settings.

By choosing  $\mathcal{C}_{-a}$  as the minimum clique cover for each  $a$  in the definition of  $\Delta_a^\gamma$ , a simplified, asymptotic regret bound can be deduced:

► **Corollary 3.3.** *When  $\max(b_{\alpha, \gamma, \sigma}, \sum_{a \in \mathcal{K}} f_a(\gamma)) = o(\log T)$ ,*

$$\limsup_{T \rightarrow \infty} \frac{\mathbb{E}[R(T)]}{\log T} \leq \sum_{\substack{a \in \mathcal{K} \\ \Delta_a^\gamma > 0}} \frac{4\alpha}{\Delta_a^\gamma} \leq \sum_{\substack{a \in \mathcal{K} \\ \tilde{\Delta}_a > 0}} \frac{8\alpha\theta([\mathcal{G}^\gamma]_{-a})}{\tilde{\Delta}_a}, \quad (5)$$

where  $\tilde{\Delta}_a := \min_{v \in \mathcal{V}_{-a}} \Delta_a^v$ .

Note that  $\tilde{\Delta}_a$  is the suboptimality gap introduced in Yang et al. [68], where they studied the setting of heterogeneous bandits on a fully-connected network. When  $\mathcal{G}$  is a fully connected network,  $\theta(\cdot) = 1$ , and we recover their regret bound, with an improved constant in  $\alpha$ .

Since our general setting also applies to the restricted cases considered in the previous literature, we can compare our regret bounds to existing ones. In the same setting without collaboration, the group regret scales as  $\sum_{a \in \mathcal{K}} \frac{|\mathcal{V}_a| \log T}{\tilde{\Delta}_a}$ ; thus when  $\tilde{\Delta}_a$  is considered to be constant, the regret always scales linearly in  $N$ . Compared to this, depending on the network, the regret bound of FLOODING scales with the clique covering number of subgraphs of  $\mathcal{G}^\gamma$ , which is usually strictly less than  $N$ , and in some cases, even sublinear.

Similarly, our regret bounds and our problem-dependent difficulty quantity  $\Delta_a^\gamma$  also generalize previous literature on collaborative multi-agent multi-armed bandits. When the network is a clique, we recover the regret bound presented in [68] with matching  $\log T$  dependency and an improved leading coefficient<sup>3</sup>. In the homogeneous agents setting with a general network topology [34, 49],  $\Delta_a^\gamma$  reduces to  $\frac{\Delta_a^{Kolla}}{\theta(\mathcal{G}^\gamma)}$ , where  $\Delta_a^{Kolla}$  is the suboptimality gap as defined in [34], satisfying  $\Delta_a^{Kolla} := \Delta_a^v$  for all  $v \in \mathcal{V}$ . As  $\frac{1}{\theta(\mathcal{G}^\gamma)}$  is independent of the arm  $a$ , we have shown that our  $\Delta_a^\gamma$  successfully generalizes the suboptimality gap of [34]. When  $\gamma = \text{diam}(\mathcal{G})$ , we have that  $\Delta_a^\gamma = \Delta_a^{Kolla}$  as  $\theta(\mathcal{G}^\gamma) = 1$ , which results in the same regret bound as in [34].

When  $\gamma = 1$  (IRS), it can be observed that IRS and FwA coincide, yet our bound is a bit worse compared to [34], whose bound depends on  $\alpha(\mathcal{G})$ . We believe that such a gap is an inherent artifact of our proof, and we leave closing this gap to future work. Also, we should remark that the difference between the two regret bounds for IRS is generally small, as the gap depends on the covering gap  $\theta(\mathcal{G}) - \alpha(\mathcal{G})$ , which is known to be small for many classes of graphs, and zero for perfect graphs [27]; see [28, 53] for some recent advances.

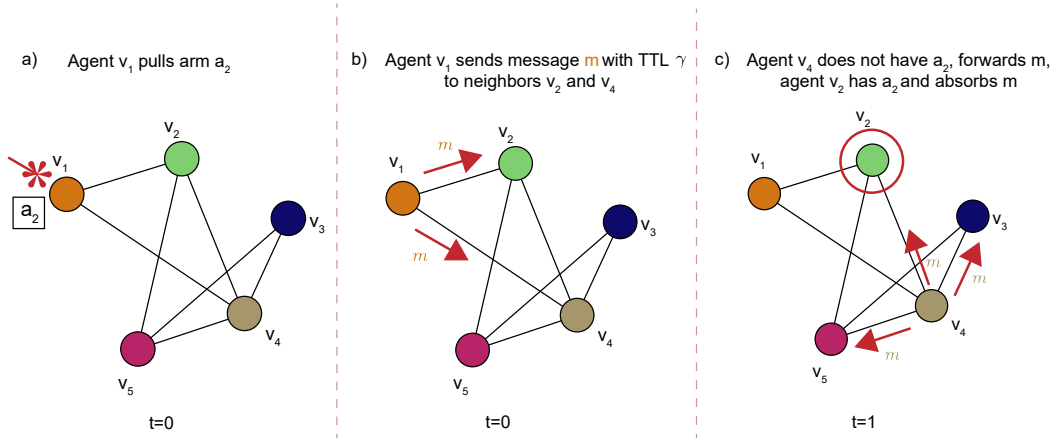
### 3.3 Drawbacks of Flooding

Flooding leads to optimal information dissemination, and thus significantly improves the group regret. Yet, it is very expensive in terms of communication complexity, defined as the *cumulative* number of messages sent by all agents [68]. Indeed, for  $\gamma = \text{diam}(\mathcal{G})$ , the worst-case communication complexity is  $\gamma N |\mathcal{E}| T$ , which is attained when every message created by every agent up to time  $T$  is being passed around at every edge. This can quickly congest the network. One naïve way of controlling the communication complexity is to set the TTL,  $\gamma$ , to a low enough value. However, in our setting, the trade-off between communication complexity and group regret is not trivial due to the arm heterogeneity; for instance, IRS [34, 49], i.e.,  $\gamma = 1$ , has a lower message complexity  $N |\mathcal{E}| T$  but often does not result in good regret guarantees, as immediate neighbors may not share any arms. On the other hand, (uniform) gossiping algorithms for bandit problems [14, 56, 57] suffer from large latencies on networks with sparse links [29, 11].

It is thus desirable to have a simple communication protocol with good regret guarantees when combined with the UCB policy (compared to FLOODING) and low communication complexity. With this, we can target more complex network structures commonly found in real-life applications. For such settings, we introduce a new protocol that interpolates between the communication-efficient nature of IRS and the regret of FLOODING by using the intrinsic heterogeneity of the system caused by network topology and arm distributions.

<sup>3</sup> Theorem 2 of [68] requires  $\alpha > 2$ , and thus with proper scaling, it can be seen that our coefficient is  $8\alpha$  while their coefficient is  $24\alpha$ .





■ **Figure 2 Flooding with Absorption (FwA).** **a.** An agent ( $v_1$ ) pulls one of its arms ( $a_2$ ). **b.**  $v_1$  sends a message  $m$  to its neighbors, with a TTL  $\gamma$ . **c.** Since one receiver of the message ( $v_4$ ) does not have  $a_2$  in its arm set, they forward  $m$  to their neighbors except the originator  $v_1$ . The other receiver ( $v_2$ ) has arm  $a_2$  in their arm set, and thus it absorbs  $m$ .

#### 4 A New Efficient & Effective Protocol on Complex Networks: Flooding with Absorption

In this section, we propose a new approach, which we call **Flooding with Absorption (FwA)** (Figure 2), whose pseudocode is shown in Algorithm 1 with *absorb* = *True*. In contrast to FLOODING, once a message (some copy of it to be precise) containing arm  $a$  reaches an agent whose arm set includes the arm  $a$ , the agent *absorbs* that message, i.e., does not forward it any further. Additionally, as in FLOODING, we retain the TTL  $\gamma$ , meaning that if a message originating at time  $t$  has not found an absorbing agent until  $t' = t + \gamma$ , it gets discarded. It is also dropped if the message hits a “dead end,” i.e., a leaf node.

This seemingly small difference to FLOODING is critical in ensuring low communication complexity, as it prevents messages from circulating for too long. We note that FwA is somewhat reminiscent of the well-studied replication-based epidemic- and other controlled flooding algorithms [62, 55, 45, 24], which were designed for various networking applications, e.g., ad-hoc networks. Our FwA protocol distinguishes itself by *using* the inherent heterogeneity of agents *without* any explicit tuning or need for solving NP-hard combinatorial problems [45]. Furthermore, the goal of FwA is to disseminate information, which is generated at each timestep, to nodes that may benefit from it for its learning, *not* to route packets from an arbitrary source to an arbitrary destination. Despite some outward protocol similarity, it hence also differs from classic P2P systems such as Gnutella [47, 48], where there is no such intrinsic correlation between sender and receiver. In the bandit setting, FwA is advantageous because the sender and final receiver share the arm in question.

► **Remark 4.1.** Each agent must have a sufficiently large memory buffer to store the messages to be sent in the next round and previously seen message identifiers. As all messages expire after  $\gamma$  rounds, this memory requirement is at most  $\gamma N$ . Also, we note that the communication complexity of FwA ranges from  $N|\mathcal{E}|T$  to  $\gamma N|\mathcal{E}|T$ , depending on the underlying network topology and the arm distribution.



### 4.1 Group Regret Bound of FwA

As FwA is algorithmically similar to FLOODING, their regret bounds are also somewhat similar. To formalize this, we first consider a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , and let  $\mathbf{c} : \mathcal{V} \rightarrow 2^{\mathcal{K}}$  be a multi-coloring with overlap allowed, i.e., it may be that  $\mathbf{c}(v) \cap \mathbf{c}(w) \neq \emptyset$  for  $\{v, w\} \in \mathcal{E}$ . Let  $a \in \mathcal{K}$  and  $v, w \in \mathcal{V}$  be such that  $a \in \mathbf{c}(v) \cap \mathbf{c}(w)$ .

► **Definition 4.2.** A path  $v_0 v_1 \cdots v_m$  (of length  $m$ ) is said to  **$a$ -free** if  $a \notin \bigcup_{i \in [m-1]} \mathbf{c}(v_i)$ , where  $[m] = \{0, 1, \dots, m\}$ .

► **Definition 4.3.** For  $\gamma \geq 1$  and  $a \in \mathcal{K}$ , we define the  **$(a, \mathbf{c})$ -non-blocking graph power of  $\gamma$ -th order** of  $\mathcal{G}$ , denoted as  $\mathcal{G}_{(a, \mathbf{c})}^\gamma$ , as a graph on  $\mathcal{V}$  with the edge set  $\mathcal{E}_{(a, \mathbf{c})}^\gamma$  such that  $\{v, w\} \in \mathcal{E}_{(a, \mathbf{c})}^\gamma$  iff there exists an  $a$ -free path from  $v$  to  $w$  in  $\mathcal{G}$  of length at most  $\gamma$ .

► **Remark 4.4.** Def. 4.3 is similar to *color-avoiding percolation (CAP)* in statistical physics [35, 36, 33], albeit there are several differences. We consider a multi-coloring  $\mathbf{c} : v \mapsto \mathcal{K}_v$ , while CAP is only studied for a single color per vertex. Also, CAP only considered the criticality of the connectivity, while in our case, the performance of our algorithm depends on specific graph invariants (e.g., chromatic number) of color-dependent subgraphs.

Defining the suboptimality gap  $\Delta_a^{FwA}$  for FwA as

$$\Delta_a^{FwA, \gamma} := \max_{\mathcal{C}_{(a, \mathbf{c})}^\gamma} \left( \sum_{C \in \mathcal{C}_{(a, \mathbf{c})}^\gamma} \left\{ \frac{2}{\min_{v \in C} \Delta_a^v} - \frac{1}{\max_{v \in C} \Delta_a^v} \right\} \right)^{-1}, \quad (6)$$

where  $\min_{\mathcal{C}_{(a, \mathbf{c})}^\gamma}$  is over all possible clique covers  $\mathcal{C}_{(a, \mathbf{c})}^\gamma$  of  $\mathcal{G}_{(a, \mathbf{c})}^\gamma$ , it is easy to see that the following theorem holds:

► **Theorem 4.5.** With `absorb` = `True`, Theorem 3.2 holds with  $\Delta_a^\gamma$  replaced by  $\Delta_a^{FwA, \gamma}$ .

Similarly, with an appropriate choice of the clique cover, we have the following simplified asymptotic regret bound:

► **Corollary 4.6.** With the same assumption as in Theorem 3.2, we have

$$\limsup_{T \rightarrow \infty} \frac{\mathbb{E}[R(T)]}{\log T} \leq \sum_{\substack{a \in \mathcal{K} \\ \Delta_a^\gamma > 0}} \frac{4\alpha}{\Delta_a^{FwA, \gamma}} \leq \sum_{\substack{a \in \mathcal{K} \\ \Delta_a^\gamma > 0}} \frac{8\alpha\theta \left( [\mathcal{G}_{(a, \mathbf{c})}^\gamma]_{-a} \right)}{\tilde{\Delta}_a}, \quad (7)$$

where  $\tilde{\Delta}_a = \min_{v \in \mathcal{V}_{-a}} \Delta_a^v$ .

As  $\mathcal{G}_{(a, \mathbf{c})}^\gamma$  is always a subgraph of  $\mathcal{G}^\gamma$ , it can be easily seen that the regret upper-bound of FLOODING is always better than that of FwA. But as we will demonstrate later, at the price of *slightly* worse regret, FwA obtains significantly better communication complexity than FLOODING. Corollary 3.3 and 4.6 imply that the gap in the asymptotic regret upper-bounds of FLOODING and FwA roughly scales with  $\sum_{\substack{a \in \mathcal{K} \\ \Delta_a^\gamma > 0}} \frac{\delta_a}{\tilde{\Delta}_a}$ , where  $\delta_a := \theta \left( [\mathcal{G}_{(a, \mathbf{c})}^\gamma]_{-a} \right) - \theta \left( [\mathcal{G}^\gamma]_{-a} \right)$ .

To get some intuition, we consider two extreme cases. First, suppose that the arms are so heterogeneous that no agents of distance at most  $\gamma$  share arm  $a$ . In this case, we have  $\mathcal{G}_{(a, \mathbf{c})}^\gamma = \mathcal{G}^\gamma$ , and  $\delta_a = 0$ . Now suppose that all agents have the same arm set, i.e.,  $\mathcal{K}_v = \mathcal{K}$  for all  $v \in \mathcal{V}$ , in which case FwA is equivalent to IRS, i.e., messages do not get forwarded beyond direct neighbors. Hence, we have that  $\mathcal{G}_{(a, \mathbf{c})}^\gamma = \mathcal{G}$  for all  $\gamma \geq 1$  and  $a \in \mathcal{K}$ , i.e.,  $\delta_a = \theta([\mathcal{G}]_{-a}) - \theta([\mathcal{G}^\gamma]_{-a})$ . Thus, for small  $\gamma$ 's and  $\mathcal{G}$  with *large* average path length [5] between agents containing  $a$ ,  $\delta_a$  is small.

## 4.2 Advantages of Flooding with Absorption

We now informally argue the advantages of using FLOODING WITH ABSORPTION over other protocols such as FLOODING or IRS when we run it on complex network topologies.

*Interpolation between IRS and Flooding.* FWA naturally interpolates between IRS and FLOODING in terms of information propagation, which is advantageous on complex networks that are not particularly regular, e.g., those with both dense and sparse regions. In dense parts of the network, where many nodes share arms, FWA is closer to IRS: a message containing the shared arm and its reward gets absorbed quickly. On the other hand, in regions of the network where the arm that a particular node pulled is rare, FWA acts like FLOODING with  $\gamma \gg 1$ , thereby ensuring that agents at a larger distance get information that is relevant. We additionally note that setting the TTL to larger values in FWA will always be less costly than doing so in FLOODING, as the probability of congestion is much smaller.

*Comparable Regret Guarantees.* As FWA acts as a mix of IRS and FLOODING, its regret should be bounded by IRS (where messages get absorbed in just one step) from above, and FLOODING (where messages are not absorbed until the TTL expires) from below. The combination of Theorems 3.2 and 4.5 gives us an expression for the gap between the regret upper bounds of FLOODING and FWA. From this, we can conclude that for the regret gap between FWA and FLOODING to be small, either the graph is so sparse that the average path length [5] between agents *containing*  $a$  is large, or the graph is dense but the arm distribution is sparse enough such that the same property holds. We emphasize that although the gap may be nonzero, the exploding communication complexity of FLOODING demonstrates a clear trade-off between performance and communication complexity. On the other hand, it is also expected that FWA will outperform (uniform) gossiping in terms of the regret. In fact, in networks with sparse links connecting very dense network regions, the probability that a gossiping protocol hits the sparse link before the TTL expires can be arbitrarily small.

*Communication Efficiency.* Having messages absorbed by agents that can profit from its information implies that the FWA protocol completely falls back to the baseline FLOODING algorithm only in the case of a network where particular arms are very rare. This means that if there is a ball of radius  $\gamma$  in the network in which two agents share an arm, the communication complexity of FWA will already be lower than that of FLOODING,  $O(N \cdot |\mathcal{E}| \cdot \gamma \cdot T)$ . In networks of high density and few arms, the communication complexity of FWA will be close to that of IRS, i.e., lowered by a factor  $\gamma$ . Moreover, due to the arm-dependent absorptions, we expect that FWA will result in much lesser number of messages sent across each individual link per round. Hence, FWA has the advantage of being able to mitigate network overload and heavy link congestion without much overhead or the need to resort to gossiping protocols, which is particularly salient for applications such as large-scale and wireless networks.

*No tuning requirements.* FWA also has an important practical advantage: it has no tunable parameters beyond the TTL  $\gamma$ , which means it is close to network agnostic. This is in contrast to protocols like probabilistic flooding [54], which stops message propagation with some constant probability  $q$ . While such protocols can also reduce communication complexity to a scalable degree, the “optimal” stopping probability  $q$  is highly instance-dependent, making them quite hard to use in unknown, real-life networks, in particular, dynamically changing networks. FWA, on the other hand, can effectively deal with such instances, which we verify in Section 6.

## 5 Regret Lower Bound

We consider a decentralized<sup>4</sup> policy  $\Pi = (\pi^v)_{v \in \mathcal{V}}$ , where  $\pi^v : [T] \rightarrow \mathcal{P}(\mathcal{K})$  is the agent-wise policy followed by agent  $v$ , possibly affected by other policies and the history. For the regret lower bound, we consider a rather general class of policies satisfying the following property, which has been widely adapted in bandit literature [37, 34, 22]:

► **Definition 5.1.**  $\Pi$  is said to be **individually consistent** if, for any agent  $v$  and any  $a \in \mathcal{K}_{-v}$ , we have that  $\mathbb{E}[N_a(T)] = o(T^c)$ ,  $\forall c > 0$ , where  $N_a(T) := \sum_{v \in \mathcal{V}_a} N_a^v(T)$ .

One obtains the following regret lower bound:

► **Theorem 5.2.** For any individually consistent policy  $\Pi$ ,

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}[R(T)]}{\log T} \geq \sum_{\substack{a \in \mathcal{K} \\ \tilde{\Delta}_a > 0}} \frac{\tilde{\Delta}_a}{\inf_{P \in \mathcal{M}_\sigma} \{D_{\text{KL}}(P_a, P) : \mu(P) - \mu(P_a) > \tilde{\Delta}_a\}}, \quad (8)$$

where  $\mathcal{M}_\sigma$  is the set of  $\sigma$ -sub-Gaussian distributions with mean  $\mu$ , and  $\tilde{\Delta}_a = \min_{v \in \mathcal{V}_{-a}} \Delta_a^v$ . When  $\mathcal{M}_\sigma = \{\mathcal{N}(\mu, \sigma^2) : \mu \in \mathbb{R}\}$ , we obtain:

$$\liminf_{T \rightarrow \infty} \frac{\mathbb{E}[R(T)]}{\log T} \geq \sum_{\substack{a \in \mathcal{K} \\ \tilde{\Delta}_a > 0}} \frac{2\sigma^2}{\tilde{\Delta}_a}. \quad (9)$$

The proof is immediate from the classic change-of-measure argument for cooperative multi-agent bandit setting [22, 68]. Note that this asymptotic lower bound matches our asymptotic regret upper bound for both FLOODING and FWA (Theorem 4.5, Cor. 4.6) up to some graph topology and arm distribution-dependent constants; see Appendix C.2 for a more detailed discussion.

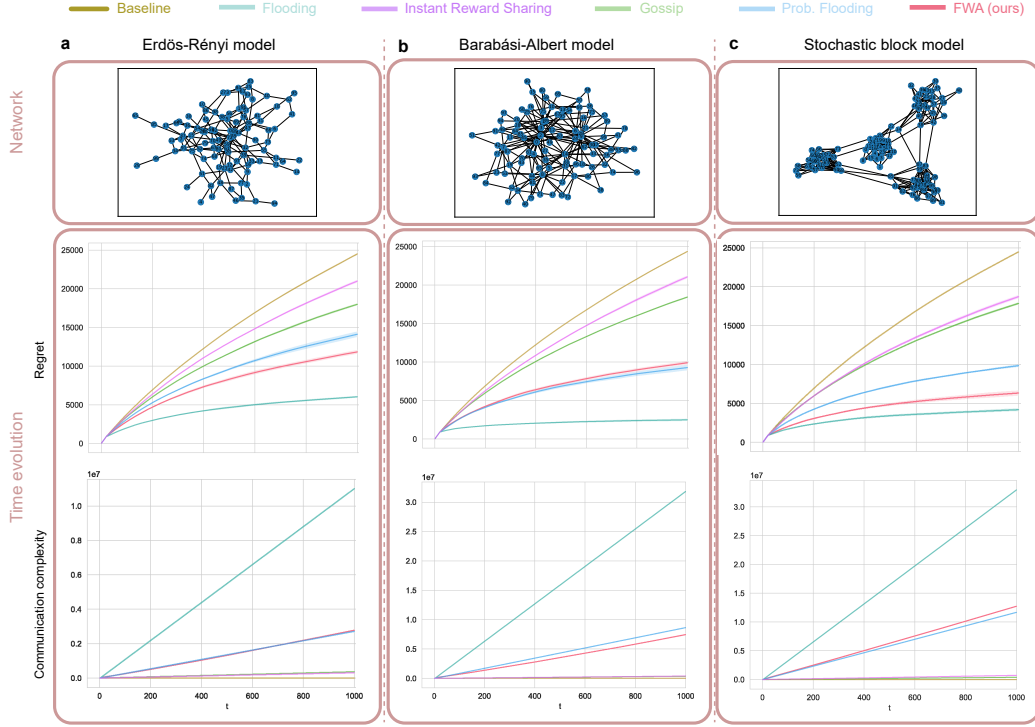
## 6 Experimental Results

In this section, we experimentally compare FWA to several existing algorithmic solutions. The experiments were conducted on three random graph models with  $N = 100$  nodes: the Erdős-Rényi model (ER) [23, 26], the Barabási-Albert model (BA) [4], and the stochastic block model (SBM) [31]. We use the following hyperparameters when generating the random graphs: for ER, the edge probability is set to  $p_{ER} = 0.03$ ; for BA, the preferential attachment constant is set to 2; for SBM, we consider 4 clusters, each with  $N/4$  nodes, with intraccluster and intercluster edge probabilities set to 0.3 and 0.003, respectively.

We set the total number of arms to  $K = 50$ , and the number of arms per agent to be  $k = 20$ . We sample sets of size  $k$  as arm sets for all the agents, uniformly at random. The arm rewards follow Gaussian distributions, with the corresponding means uniformly sampled from  $[0, 1]$  and fixed variance  $\sigma^2 = 1$ . For all experiments, the overall arm distribution among the agents and the reward distributions are fixed.

We compare the baseline UCB with no cooperation between agents (baseline), FLOODING, PROBABILISTIC FLOODING (PROB. FLOODING) [54], (uniform) GOSSIP, IRS, and our FWA. For the GOSSIP algorithm, we assume that each agent forwards messages to only one random

<sup>4</sup> see Appendix A of [22] for the precise measure-theoretic definition of decentralized policies.




■ **Figure 3** Comparing group regret and (cumulative) communication complexity across different topologies and protocols. Note that FWA gives a good trade-off between regret and communication complexity.

neighbor at a time. For **PROB. FLOODING**, assuming that the learner has no prior knowledge of the communication network, we fix  $q = 0.5$  to provide a fairer comparison. We again emphasize that FWA does *not* require any tuning of hyperparameters other than the TTL.

All experiments are repeated 10 times, with time horizon  $T = 1000$ . We set TTL to be  $\gamma = 4$ , as our network instances' diameters are 8, 5, 7 for ER, BA, and SBM, respectively, and we want our  $\gamma$  to be strictly lower than all three values for meaningful results. All codes were written in Python, and we made heavy use of the NetworkX package [30].

### 6.1 Baseline comparison: Group Regret and Communication Complexity

We first compare the group regret and the communication complexity over time, i.e., the (cumulative) total number of messages sent, for all considered protocols. The results are shown in Figure 3. As expected, **FLOODING** achieves the best regret out of the tested protocols, but its communication complexity is the worst. Despite this, the important observation is that our FWA achieves second-best regret, arguably close to that achieved by **FLOODING**, with a significantly reduced communication complexity. Moreover, when compared to **PROB. FLOODING**, FWA exhibits a better tradeoff between regret and communication: with a similar communication complexity, FWA achieves a better regret, or at least at par (for the BA model). We also experimented with **PROB. FLOODING** of other stopping probabilities (not shown here) and observed that they tend to show worse trade-offs between regret and communication complexity. This shows that our proposed FWA protocol is a viable alternative to **FLOODING** if one needs reduced communication complexity and good regret, uniformly across various network topologies.



■ **Figure 4** FWA significantly decreases congestion on sparse network links. We find that, in comparison with other protocols, FWA results in a reduced number of messages sent over such a sparse link (highlighted in the networks).

## 6.2 Link congestion

In a setting where new messages are constantly produced by every agent (as each pulls an arm at each time step), one of the potential issues is link congestion caused by a large number of messages passing through bottleneck links. This can lead to significantly decreased performance and undesirable latency effects - messages may be queued with limited memory in reliable link protocols, or automatically discarded in non-reliable link protocols once more messages are being sent than the link can handle. In Figure 4, we visualize the number of messages sent over a particular link, again with TTL  $\gamma = 4$ . We chose a random link for ER and BA. For SBM, we chose a “sparse” link that connects two dense clusters. Out of all the considered protocols, FWA results in the largest reduction of messages per round while providing good regret; specifically, for ER, BA, and SBM, FWA provides about 77%, 80%, 83% reduction compared to FLOODING, respectively. We note that the reduction is larger even compared to PROB. FLOODING - even though PROB. FLOODING can occasionally have slightly better overall communication complexity (at worse regret). This implies that our protocol exhibits significant benefits regarding individual network link congestion, which would help us avoid latency effects in real-life network applications.

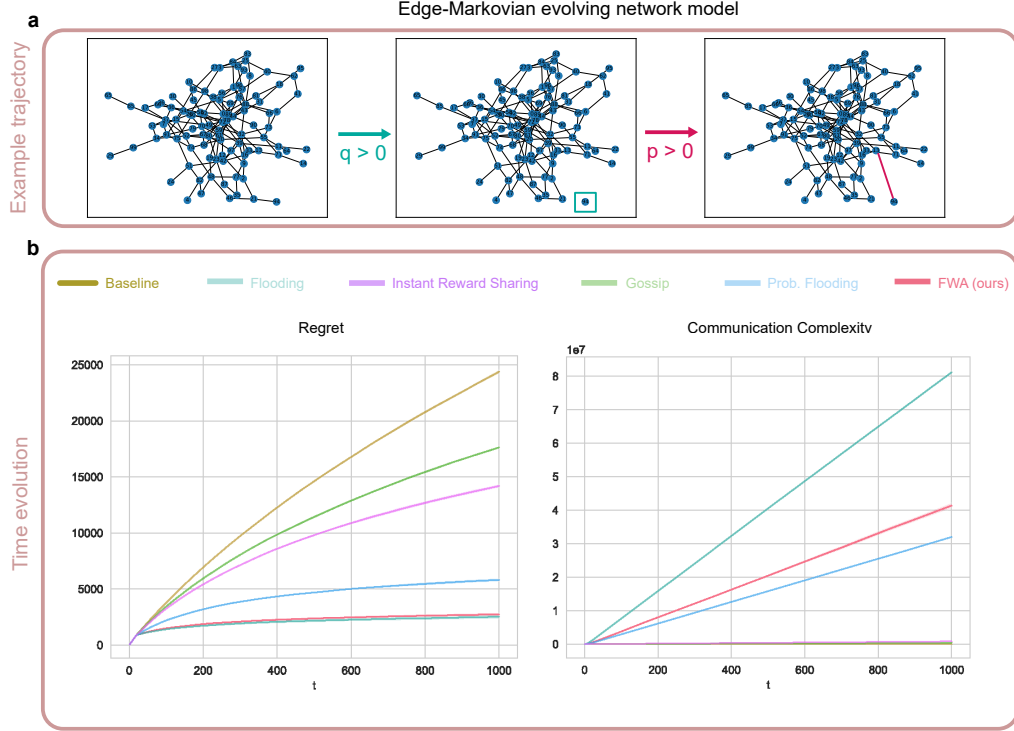
One interesting observation is that FWA produces a spike in the number of messages in the early phase for all network topologies. This is due to the design of the UCB algorithm; in the early phase, most of the agents are exploring the arms, and thus the messages are somewhat “diverse”. But as soon as the agents identify potentially best arms, the arm indices of the messages start to stabilize. They become less diverse, implying that from then on, absorption occurs more frequently under FWA.

## 6.3 Dynamic networks

The advantage of our protocol is especially pronounced when we consider *dynamically changing networks*, specifically, edge-Markovian networks [20, 18, 19, 17], where the network evolves as follows: starting from an arbitrary initial graph  $\mathcal{G}_0 = (\mathcal{V}, \mathcal{E}_0)$ , for  $t \geq 1$ ,  $\mathcal{G}_t = (\mathcal{V}, \mathcal{E}_t)$  is stochastically determined as

$$\mathbb{P}[\{v, w\} \in \mathcal{E}_t | \{v, w\} \notin \mathcal{E}_{t-1}] = p, \quad \mathbb{P}[\{v, w\} \notin \mathcal{E}_t | \{v, w\} \in \mathcal{E}_{t-1}] = q. \quad (10)$$

$p$  and  $q$  are often referred to as edge birth rate and edge death rate, respectively. When  $0 < p, q < 1$ , it is well-known that, regardless of the initial graph  $\mathcal{G}_0$ , the process converges to (stationary) Erdős-Rényi graph  $\mathcal{G}(\mathcal{V}, \frac{p}{p+q})$ . For our experiments, we start from the baseline ER graph and set  $(p, q) = (0.01, 0.1)$ . We plot an example trajectory in Figure 5a.




■ **Figure 5** Comparing group regret and (cumulative) communication complexity in a dynamic network setting. Note that FWA achieves the same regret as FLOODING, with much lesser cumulative communication complexity.

The results are shown in Figure 5b. Observe how well our FWA protocol matches the regret of FLOODING, with *strictly* better communication complexity. This trend is consistent across all considered networks, showing that FWA is the most effective out of the considered protocols in dynamic networks. Similarly to the static case, we experimented with PROB. FLOODING of other stopping probabilities (not shown here) and observed that they tend to show worse trade-offs between regret and communication complexity. This suggests that the arm-dependent absorption mechanism of FWA implicitly regularizes the communication complexity in an efficient manner in dynamic networks, while minimizing the loss in regret.

#### 6.4 Tightness of Theoretical Regret Upper Bounds

Recall that the theoretically derived regret bounds of FLOODING and FWA (Theorem 3.2 and 4.5) depend on both network topology and arm distribution. FLOODING scales with  $\delta^{Flooding} := \sum_{\substack{a \in \mathcal{K} \\ \Delta_a > 0}} \frac{\theta([G^\gamma] - a)}{\Delta_a}$  and FWA scales with  $\delta^{FWA} := \sum_{\substack{a \in \mathcal{K} \\ \Delta_a > 0}} \frac{\theta([G_{(a,c)}^\gamma] - a)}{\Delta_a}$ . To show that the theoretical bounds are tight and match well with practice, we perform an ablation study by varying the underlying edge density and seeing whether the aforementioned quantities scale well with the actual regrets.

We consider  $\mathcal{G} \sim G(n, p)$  of varying edge density  $p$  and compare estimated  $\delta$ 's and regrets of FLOODING and FWA under the same setting as in previous experiments. Computing  $\theta(\cdot)$ 's requires computing the chromatic numbers, which is NP-hard. We thus approximate it with the size of a greedy coloring of the considered graph; for Erdős-Rényi graphs, the greedy coloring asymptotically results in twice the true chromatic number [25].



(a) Regret of FLOODING vs.  $\delta^{FLOODING}$ . Orange line is the best linear fit ( $R^2 = 0.9439$ ). (b) Regret of FwA vs.  $\delta^{FwA}$ . Orange line is the best linear fit ( $R^2 = 0.9813$ ).

**Figure 6** Experimental Results on  $\delta$ . Note that there is a strong linear correlation between the estimated  $\delta^{FLOODING}$ ,  $\delta^{FwA}$  and the final resulting regrets of FLOODING, FwA, respectively.

In Figure 6, we scatter plotted  $(\delta, \text{Regret})$  along with a best linear fit for FLOODING and FwA. Indeed, it can be seen that the relationship is almost linear, with high  $R^2$ , showing that our regret bounds indeed reflect the regrets in practice. There are some deviations from linearity, which we believe is due to small horizon length  $T$  and inaccuracy in estimating  $\delta$ .

## 7 Conclusion

In this work, we described a novel setting for distributed multi-armed bandits, where agents communicate on an underlying network and do not all share the same arm set. We assume that each agent runs a UCB algorithm to identify their local best arm, and communicates the information they receive to their neighbors to minimize cumulative group regret. To deal with the very large communication complexity that however arises from using FLOODING in our setting, we then introduced a new communication protocol for complex networks, FLOODING WITH ABSORPTION (FwA). With FwA, agents forward information only if it pertains to an arm they themselves do not include in their arm set, whereas they absorb a message that gives information about one of their own arms. We provided theoretical upper and lower regret bounds and showed experimentally that FwA incurs only minimal group regret performance loss compared to FLOODING and even PROBABILISTIC FLOODING, while leading to a significantly improved communication complexity. In particular, we showed that FwA can reduce link congestion, which significantly improves upon simple heuristics such as probabilistic flooding. Our protocol is fully network-agnostic and hence does not need any fine-tuning, while still making use of the inherent heterogeneity of the problem instance. This makes it a very suitable choice for dynamically changing networks, or those that suffer from occasional message loss. We believe that our work highlights the importance of integrating network topology and action heterogeneity in the design of distributed bandit algorithms, and provides an efficient way to connect bandit learning and network protocols.



## References

- 1 Animashree Anandkumar, Nithin Michael, and Ao Tang. Opportunistic Spectrum Access with Multiple Users: Learning under Competition. In *2010 Proceedings IEEE INFOCOM*, pages 1–9, 2010. doi:10.1109/INFCOM.2010.5462144.
- 2 Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning*, 47(2):235–256, 2002. doi:10.1023/A:1013689704352.
- 3 Orly Avner and Shie Mannor. Multi-user lax communications: A multi-armed bandit approach. In *IEEE INFOCOM 2016 - The 35th Annual IEEE International Conference on Computer Communications*, pages 1–9, 2016. doi:10.1109/INFOCOM.2016.7524557.
- 4 Albert-László Barabási and Réka Albert. Emergence of Scaling in Random Networks. *Science*, 286(5439):509–512, 1999. doi:10.1126/science.286.5439.509.
- 5 Albert-László Barabási and Réka Albert. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47–97, 2002. doi:10.1103/RevModPhys.74.47.
- 6 Oded Berger-Tal, Jonathan Nathan, Ehud Meron, and David Saltz. The Exploration-Exploitation Dilemma: A Multidisciplinary Framework. *PLOS ONE*, 9:1–8, apr 2014. doi:10.1371/journal.pone.0095693.
- 7 Béla Bollobás. *Modern Graph Theory*, volume 184 of *Graduate Texts in Mathematics*. Springer, 2002. doi:10.1007/978-1-4612-0619-4.
- 8 Sébastien Bubeck. *Bandits Games and Clustering Foundations*. PhD thesis, INRIA Nord Europe, jun 2010.
- 9 Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122, 2012. doi:10.1561/22000000024.
- 10 Swapna Buccapatnam, Jian Tan, and Li Zhang. Information sharing in distributed stochastic bandits. In *2015 IEEE Conference on Computer Communications (INFOCOM)*, pages 2605–2613, 2015. doi:10.1109/INFOCOM.2015.7218651.
- 11 Keren Censor-Hillel, Bernhard Haeupler, Jonathan Kelner, and Petar Maymounkov. Global Computation in a Poorly Connected World: Fast Rumor Spreading with No Dependence on Conductance. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing, STOC '12*, pages 961–970. Association for Computing Machinery, 2012. doi:10.1145/2213977.2214064.
- 12 Nicolò Cesa-Bianchi, Tommaso Cesari, and Claire Monteleoni. Cooperative Online Learning: Keeping your Neighbors Updated. In *Proceedings of the 31st International Conference on Algorithmic Learning Theory*, volume 117 of *Proceedings of Machine Learning Research*, pages 234–250. PMLR, 08 February–11 February 2020. URL: <http://proceedings.mlr.press/v117/cesa-bianchi20a.html>.
- 13 Nicholas B Chang and Mingyan Liu. Controlled Flooding Search in a Large Network. *IEEE/ACM Transactions on Networking*, 15(2):436–449, 2007. doi:10.1145/1279660.1279675.
- 14 Ronshee Chawla, Abishek Sankararaman, Ayalvadi Ganesh, and Sanjay Shakkottai. The Gossiping Insert-Eliminate Algorithm for Multi-Agent Bandits. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 3471–3481. PMLR, 26–28 August 2020. URL: <http://proceedings.mlr.press/v108/chawla20a.html>.
- 15 Ronshee Chawla, Daniel Vial, Sanjay Shakkottai, and R. Srikant. Collaborative Multi-Agent Heterogeneous Multi-Armed Bandits. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 4189–4217. PMLR, 23–29 July 2023. URL: <https://proceedings.mlr.press/v202/chawla23a.html>.
- 16 Yu-Zhen Janice Chen, Lin Yang, Xuchuang Wang, Xutong Liu, Mohammad Hajiesmaili, John C. S. Lui, and Don Towsley. On-Demand Communication for Asynchronous Multi-Agent Bandits. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pages 3903–3930. PMLR, 25–27 April 2023. URL: <https://proceedings.mlr.press/v206/chen23c.html>.

- 17 Andrea Clementi, Pierluigi Crescenzi, Carola Doerr, Pierre Fraigniaud, Francesco Pasquale, and Riccardo Silvestri. Rumor Spreading in Random Evolving Graphs. *Random Structures & Algorithms*, 48(2):290–312, 2016. doi:10.1002/RSA.20586.
- 18 Andrea Clementi, Angelo Monti, Francesco Pasquale, and Riccardo Silvestri. Information Spreading in Stationary Markovian Evolving Graphs. *IEEE Transactions on Parallel and Distributed Systems*, 22(9):1425–1432, 2011. doi:10.1109/TPDS.2011.33.
- 19 Andrea Clementi, Riccardo Silvestri, and Luca Trevisan. Information spreading in dynamic graphs. *Distributed Computing*, 28(1):55–73, feb 2015. doi:10.1007/S00446-014-0219-2.
- 20 Andrea E. F. Clementi, Claudio Macci, Angelo Monti, Francesco Pasquale, and Riccardo Silvestri. Flooding Time of Edge-Markovian Evolving Graphs. *SIAM Journal on Discrete Mathematics*, 24(4):1694–1712, 2010. doi:10.1137/090756053.
- 21 Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. The MIT Press, 4 edition, 2022.
- 22 Abhimanyu Dubey and Alex Pentland. Cooperative Multi-Agent Bandits with Heavy Tails. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 2730–2739. PMLR, 13–18 July 2020. URL: <http://proceedings.mlr.press/v119/dubey20a.html>.
- 23 Paul Erdős and Alfréd Rényi. On random graphs, I. *Publicationes Mathematicae Debrecen*, 6:290–297, 1959.
- 24 P.T. Eugster, R. Guerraoui, A.-M. Kermarrec, and L. Massoulié. Epidemic Information Dissemination in Distributed Systems. *Computer*, 37(5):60–67, 2004. doi:10.1109/MC.2004.1297243.
- 25 Alan Frieze and Colin McDiarmid. Algorithmic Theory of Random Graphs. *Random Structures & Algorithms*, 10(1–2):5–42, feb 1997. doi:10.1002/(SICI)1098-2418(199701/03)10:1/2<%3C5::AID-RSA2%3E3.0.CO;2-Z.
- 26 E. N. Gilbert. Random Graphs. *The Annals of Mathematical Statistics*, 30(4):1141–1144, 1959.
- 27 A. Gyárfás. Problems from the world surrounding perfect graphs. *Applicationes Mathematicae*, 19(3-4):413–441, 1987.
- 28 András Gyárfás, András Sebő, and Nicolas Trotignon. The chromatic gap and its extremes. *Journal of Combinatorial Theory, Series B*, 102(5):1155–1178, 2012. doi:10.1016/J.JCTB.2012.06.001.
- 29 Bernhard Haeupler. Simple, Fast and Deterministic Gossip and Rumor Spreading. *Journal of the ACM*, 62(6), dec 2015. doi:10.1145/2767126.
- 30 Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. Exploring Network Structure, Dynamics, and Function using NetworkX. In *Proceedings of the 7th Python in Science Conference*, pages 11–15, 2008.
- 31 Paul W. Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. Stochastic blockmodels: First steps. *Social Networks*, 5(2):109–137, 1983. doi:10.1016/0378-8733(83)90021-7.
- 32 Long Jin, Shuai Li, Lin Xiao, Rongbo Lu, and Bolin Liao. Cooperative Motion Generation in a Distributed Network of Redundant Robot Manipulators With Noises. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 48(10):1715–1724, 2018. doi:10.1109/TSMC.2017.2693400.
- 33 Andrea Kadović, Sebastian M. Krause, Guido Caldarelli, and Vinko Zlatić. Bond and site color-avoiding percolation in scale-free networks. *Physical Review E*, 98:062308, dec 2018. doi:10.1103/PhysRevE.98.062308.
- 34 Ravi Kumar Kolla, Krishna Jagannathan, and Aditya Gopalan. Collaborative Learning of Stochastic Bandits Over a Social Network. *IEEE/ACM Transactions on Networking*, 26(4):1782–1795, 2018. doi:10.1109/TNET.2018.2852361.
- 35 Sebastian M. Krause, Michael M. Danziger, and Vinko Zlatić. Hidden Connectivity in Networks with Vulnerable Classes of Nodes. *Phys. Rev. X*, 6:041022, oct 2016. doi:10.1103/PhysRevX.6.041022.

- 36 Sebastian M. Krause, Michael M. Danziger, and Vinko Zlatić. Color-avoiding percolation. *Physical Review E*, 96:022313, aug 2017. doi:10.1103/PhysRevE.96.022313.
- 37 T.L Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985. doi:10.1016/0196-8858(85)90002-8.
- 38 Peter Landgren, Vaibhav Srivastava, and Naomi Ehrich Leonard. Distributed cooperative decision making in multi-agent multi-armed bandits. *Automatica*, 125:109445, 2021. doi:10.1016/J.AUTOMATICA.2020.109445.
- 39 Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020. doi:10.1017/9781108571401.
- 40 Marc Lelarge, Alexandre Proutière, and M. Sadegh Talebi. Spectrum bandit optimization. In *2013 IEEE Information Theory Workshop (ITW)*, pages 1–5, 2013. doi:10.1109/ITW.2013.6691221.
- 41 Daniel A Levinthal and James G March. The myopia of learning. *Strategic management journal*, 14(S2):95–112, 1993. doi:10.1002/smj.4250141009.
- 42 Feng Li, Dongxiao Yu, Huan Yang, Jiguo Yu, Holger Karl, and Xiuzhen Cheng. Multi-Armed-Bandit-Based Spectrum Scheduling Algorithms in Wireless Networks: A Survey. *IEEE Wireless Communications*, 27(1):24–30, 2020. doi:10.1109/MWC.001.1900280.
- 43 Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. A Contextual-Bandit Approach to Personalized News Article Recommendation. In *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, pages 661–670. Association for Computing Machinery, 2010. doi:10.1145/1772690.1772758.
- 44 Shuai Li, Ruofan Kong, and Yi Guo. Cooperative Distributed Source Seeking by Multiple Robots: Algorithms and Experiments. *IEEE/ASME Transactions on Mechatronics*, 19(6):1810–1820, 2014. doi:10.1109/TMECH.2013.2295036.
- 45 H Lim and C Kim. Flooding in wireless ad hoc networks. *Computer Communications*, 24(3):353–363, 2001. doi:10.1016/S0140-3664(00)00233-4.
- 46 Keqin Liu and Qing Zhao. Distributed Learning in Multi-Armed Bandit With Multiple Players. *IEEE Transactions on Signal Processing*, 58(11):5667–5681, 2010. doi:10.1109/TSP.2010.2062509.
- 47 Qin Lv, Pei Cao, Edith Cohen, Kai Li, and Scott Shenker. Search and Replication in Unstructured Peer-to-Peer Networks. In *Proceedings of the 16th International Conference on Supercomputing, ICS '02*, pages 84–95, New York, NY, USA, 2002. Association for Computing Machinery. doi:10.1145/514191.514206.
- 48 Qin Lv, Sylvia Ratnasamy, and Scott Shenker. Can Heterogeneity Make Gnutella Scalable? In *Revised Papers from the First International Workshop on Peer-to-Peer Systems, IPTPS '01*, pages 94–103, Berlin, Heidelberg, 2002. Springer-Verlag. doi:10.1007/3-540-45748-8\_9.
- 49 Udari Madhushani, Abhimanyu Dubey, Naomi Leonard, and Alex Pentland. One More Step Towards Reality: Cooperative Bandits with Imperfect Communication. In *Advances in Neural Information Processing Systems*, volume 34, pages 7813–7824. Curran Associates, Inc., 2021. URL: <https://proceedings.neurips.cc/paper/2021/hash/40cb228987243c91b2dd0b7c9c4a0856-Abstract.html>.
- 50 Udari Madhushani and Naomi Ehrich Leonard. Distributed Bandits: Probabilistic Communication on d-regular Graphs. In *2021 European Control Conference (ECC)*, pages 830–835, 2021. doi:10.23919/ECC54610.2021.9655031.
- 51 James G. March. Exploration and Exploitation in Organizational Learning. *Organization Science*, 2(1):71–87, 1991. doi:10.1287/orsc.2.1.71.
- 52 Katja Mehlhorn, Ben R Newell, Peter M Todd, Michael D Lee, Kate Morgan, Victoria A Braithwaite, Daniel Hausmann, Klaus Fiedler, and Cleotilde Gonzalez. Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, 2(3):191–215, 2015. doi:10.1037/dec0000033.
- 53 Andrea Munaro. Bounded clique cover of some sparse graphs. *Discrete Mathematics*, 340(9):2208–2216, 2017. doi:10.1016/J.DISC.2017.04.004.

- 54 Konstantinos Oikonomou, George Koufoudakis, Sonia Aïssa, and Ioannis Stavrakakis. Probabilistic Flooding Performance Analysis Exploiting Graph Spectra Properties. *IEEE/ACM Transactions on Networking*, 31(1):133–146, 2023. doi:10.1109/TNET.2022.3192310.
- 55 Ashikur Rahman, Wlodek Olesinski, and Pawel Gburzynski. Controlled flooding in wireless ad-hoc networks. In *International Workshop on Wireless Ad-Hoc Networks, 2004.*, pages 73–78. IEEE, 2004. doi:10.1109/IWWAN.2004.1525544.
- 56 Abishek Sankararaman, Ayalvadi Ganesh, and Sanjay Shakkottai. Social Learning in Multi Agent Multi Armed Bandits. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 3(3), dec 2019. doi:10.1145/3366701.
- 57 Devavrat Shah. Gossip Algorithms. *Foundations and Trends® in Networking*, 3(1):1–125, 2009. doi:10.1561/13000000014.
- 58 K. Sugawara, T. Kazama, and T. Watanabe. Foraging behavior of interacting robots with virtual pheromone. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 3, pages 3074–3079 vol.3, 2004. doi:10.1109/IR0S.2004.1389878.
- 59 Balazs Szorenyi, Robert Busa-Fekete, Istvan Hegedus, Robert Ormandi, Mark Jelasiy, and Balazs Kegl. Gossip-based distributed stochastic bandit algorithms. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28(3) of *Proceedings of Machine Learning Research*, pages 19–27. PMLR, 17–19 June 2013. URL: <http://proceedings.mlr.press/v28/szorenyi13.html>.
- 60 Andrew S. Tanenbaum, Nick Feamster, and David Wetherall. *Computer Networks*. Pearson, 6 edition, 2021.
- 61 Yu-Chee Tseng, Sze-Yao Ni, Yuh-Shyan Chen, and Jang-Ping Sheu. The Broadcast Storm Problem in a Mobile Ad Hoc Network. *Wireless Networks*, 8(2):153–167, mar 2002. doi:10.1023/A:1013763825347.
- 62 Amin Vahdat and David Becker. Epidemic Routing for Partially-Connected Ad Hoc Networks. Technical Report CS-200006, Duke University, 2000.
- 63 Daniel Vial, Sanjay Shakkottai, and R. Srikant. Robust Multi-Agent Multi-Armed Bandits. In *Proceedings of the Twenty-Second International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing, MobiHoc '21*, pages 161–170. Association for Computing Machinery, 2021. doi:10.1145/3466772.3467045.
- 64 Milan Vojnović and Alexandre Proutière. Hop limited flooding over dynamic networks. In *2011 Proceedings IEEE INFOCOM*, pages 685–693, 2011. doi:10.1109/INFCOM.2011.5935249.
- 65 Po-An Wang, Alexandre Proutière, Kaito Ariu, Yassir Jedra, and Alessio Russo. Optimal Algorithms for Multiplayer Multi-Armed Bandits. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 4120–4129. PMLR, 26–28 August 2020. URL: <http://proceedings.mlr.press/v108/wang20m.html>.
- 66 N. Wisitpongphan, O.K. Tonguz, J.S. Parikh, P. Mudalige, F. Bai, and V. Sadekar. Broadcast storm mitigation techniques in vehicular ad hoc networks. *IEEE Wireless Communications*, 14(6):84–94, 2007. doi:10.1109/MWC.2007.4407231.
- 67 Wenchao Xia, Tony Q. S. Quek, Kun Guo, Wanli Wen, Howard H. Yang, and Hongbo Zhu. Multi-Armed Bandit-Based Client Scheduling for Federated Learning. *IEEE Transactions on Wireless Communications*, 19(11):7108–7123, 2020. doi:10.1109/TWC.2020.3008091.
- 68 Lin Yang, Yu-Zhen Janice Chen, Mohammad H. Hajiemaili, John C. S. Lui, and Don Towsley. Distributed Bandits with Heterogeneous Agents. In *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications*, pages 200–209, 2022. doi:10.1109/INFOCOM48880.2022.9796901.
- 69 Lin Yang, Yu-Zhen Janice Chen, Stephen Pasteris, Mohammad H. Hajiesmaili, John C. S. Lui, and Don Towsley. Cooperative Stochastic Bandits with Asynchronous Agents and Constrained Feedback. In *Advances in Neural Information Processing Systems*, volume 34, pages 8885–8897. Curran Associates, Inc., 2021. URL: <https://proceedings.neurips.cc/paper/2021/hash/4a5876b450b45371f6cfe5047ac8cd45-Abstract.html>.

## A

 Notations

Notations	
$\mathcal{G}$	Communication network
$\mathcal{V}$	Set of vertices in the graph, i.e., agents
$\mathcal{G}[S]$	Induced subgraph over $S \subset \mathcal{V}$
$N =  \mathcal{V} $	The number of agents
$\mathcal{V}_a$	Set of agents having arm $a$
$\mathcal{V}_{-a}$	Set of agents having arm $a$ as the suboptimal arm
$\mathcal{K}$	The set of all arms in the network
$\mathcal{K}_v$	The set of arms agent $v$ has access to
$\mathcal{K}_{-v}$	The set of (locally) suboptimal arms agent $v$ has access to
$a_*^v$	The best local arm for agent $v$
$\mu_{a_*^v}$	Reward of best local arm for agent $v$
$\hat{\mu}_a^v(t)$	Average reward that agent $v$ has computed for arm $a$ at time $t$
$M_a^v(t)$	Number of observations of arm $a$ available to agent $v$ at time $t$
$N_a^v(t)$	Number of times agent $v$ pulls arm $a$ up to time $t$
$N_a(t)$	Number of times all agents pull arm $a$ up to time $t$
$\gamma$	Time-to-live of a message sent by an agent
$\mathcal{G}^\gamma$	Graph power of $\mathcal{G}$ of $\gamma$ -th order: graph over $\mathcal{V}$ such that when $d_{\mathcal{G}}(v, w) \leq \gamma$ , $\{v, w\}$ is an edge
$T$	Time horizon
$R(T)$	Cumulative group regret over the time horizon
$\Delta_a^v$	Agent-specific suboptimality gap

## B

 Proof of Theorem 3.2 – Regret Upper Bound

The basic proof idea is to consider  $\mathcal{G}^\gamma$ , which is equivalent to the collaboration network via flooding with TTL  $\gamma$ . Then we consider a clique covering of  $\mathcal{G}^\gamma$ , and upper bound the group regret of the full collaboration by that in which collaboration happens only intra-clique. The deterministic delays between agents in each clique are precisely their distance, measured in the *original* communication network  $\mathcal{G}$ . We proceed similarly to [68] while filling in some missing details.

For each  $a \in \mathcal{K}$  and  $v \in \mathcal{V}_{-a}$  (i.e.,  $v$  contains  $a$  as a *(locally) suboptimal* arm), define  $B_a^v(t) := \{a^v(t) = a\}$ , which can be rewritten as

$$B_a^v(t) = \left\{ \hat{\mu}_{a'}^v(t) + \sqrt{\frac{\alpha \log t}{M_{a'}^v(t)}} < \hat{\mu}_a^v(t) + \sqrt{\frac{\alpha \log t}{M_a^v(t)}} \quad \forall a' \in \mathcal{K}_v \setminus \{a\} \right\}.$$

Thus, this implies that the following event holds:

$$E_a^v(t) := \left\{ \hat{\mu}_{a_*^v}^v(t) + \sqrt{\frac{\alpha \log t}{M_{a_*^v}^v(t)}} < \hat{\mu}_a^v(t) + \sqrt{\frac{\alpha \log t}{M_a^v(t)}} \right\}.$$

i.e.  $B_a^v(t) \subseteq E_a^v(t)$ .

Next, for  $\ell_a^v := \frac{4\alpha \log T}{(\Delta_a^v)^2}$ , define

$$\tau_a^v := \min \{t = 1, \dots, T : M_a^v(t) > \ell_a^v\}, \quad (11)$$

i.e.,  $\tau_a^v$  is the first time at which agent  $v$ , and its neighbors, observe arm  $a$  at least  $\ell_a^v$  times.

Denoting  $c_a^v(t) = \sqrt{\frac{\alpha \log t}{M_a^v(t)}}$ , we have that [8]

$$E_a^v(t) \subseteq \underbrace{\left\{ \hat{\mu}_{a_*^v}^v(t) + c_{a_*^v}^v(t) \leq \mu_{a_*^v} \right\}}_{=A(t)} \cup \underbrace{\left\{ \hat{\mu}_a^v(t) \geq \mu_a + c_a^v(t) \right\}}_{=B(t)} \cup \{M_a^v(t) < \ell_a^v\}.$$

Then the regret can be rewritten as follows [8]:

$$R(T) = \mathbb{E} \left[ \sum_{a \in \mathcal{K}} \sum_{v \in V_a} \Delta_a^v N_a^v(T) \right] \leq \underbrace{\sum_{a \in \mathcal{K}} \sum_{v \in V_a} \Delta_a^v \mathbb{E}[N_a^v(\tau_a^v)]}_{(a)} + \underbrace{\sum_{a \in \mathcal{K}} \sum_{v \in V_a} \Delta_a^v \sum_{t=1}^T \mathbb{P}[E_a^v(t) \cap \{t > \tau_a^v\}]}_{(b)}.$$

(a) is bounded using clique covering and Abel transformation:

► **Lemma B.1.** *For each  $a \in \mathcal{K}$  with  $\mathcal{V}_{-a} \neq \emptyset$ , we have that*

$$\sum_{v \in V_a} \Delta_a^v \mathbb{E}[N_a^v(\tau_a^v)] \leq \frac{4\alpha \log T}{\Delta_a^\gamma} + f_a(\gamma).$$

(b) is bounded as follows:

► **Lemma B.2.** *The following hold for each  $v \in V$  and locally suboptimal  $a \in \mathcal{K}_v$ :*

$$\sum_{t=1}^T \mathbb{P}[E_t^v(a) \cap \{t > \tau_a^v\}] \leq \left( \frac{\alpha + 1/2}{\alpha - 1/2} \right)^2 \frac{8(\gamma + 1)}{\log \frac{(\gamma+1)(\alpha+1/2)}{4\sigma^2}}.$$

Combining both lemmas gives the desired statement.

## B.1 Proof of Lemma B.1

Let  $\mathcal{C}_{-a}$  be a (vertex-disjoint) clique cover of  $[\mathcal{G}^\gamma]_{-a} = \mathcal{G}^\gamma[\mathcal{V}_{-a}]$ , and let  $C \in \mathcal{C}_{-a}$ . Let  $C = \{v_1, \dots, v_c\}$  be such that  $\Delta_a^{v_1} \geq \dots \geq \Delta_a^{v_c} > 0$ .

We first have that  $\mathbb{E}[N_a^{v_1}(\tau_a^{v_1})] \leq \mathbb{E}[M_a^{v_1}(\tau_a^{v_1})] \leq \frac{4\alpha \log T}{(\Delta_a^1)^2}$ . For each  $k \in \{2, \dots, c\}$ ,

$$\begin{aligned} \sum_{j=1}^k \mathbb{E}[N_a^{v_j}(\tau_a^{v_j})] &\leq \mathbb{E}[N_a^{v_k}(\tau_a^{v_k})] + \sum_{j=1}^{k-1} \mathbb{E}[N_a^{v_j}(\tau_a^{v_j})] \\ &\stackrel{(a)}{\leq} \mathbb{E}[M_a^{v_k}(\tau_a^{v_k})] + \sum_{j=1}^{k-1} \sum_{\iota=1}^{\gamma} \mathbb{1}[d_{\mathcal{G}}(v_k, v_j) = \iota] (\mathbb{E}[N_a^{v_j}(\tau_a^{v_j})] - \mathbb{E}[N_a^{v_j}(\tau_a^{v_k} - \iota)]) \\ &\stackrel{(b)}{\leq} \frac{4\alpha \log T}{(\Delta_a^{v_k})^2} + \sum_{j=1}^{k-1} \sum_{\iota=1}^{\gamma} \mathbb{1}[d_{\mathcal{G}}(v_k, v_j) = \iota] \min \left( \gamma + \iota, \frac{4\alpha \log T}{(\Delta_a^{v_j})^2} \right) \\ &\leq \frac{4\alpha \log T}{(\Delta_a^{v_k})^2} + \sum_{j=1}^{k-1} \min \left( 2\gamma, \frac{4\alpha \log T}{(\Delta_a^{v_j})^2} \right). \end{aligned}$$

Here, (a) follows from the simple decomposition of  $M_a^{v_k}(\tau_a^{v_k})$ :

$$\begin{aligned} M_a^{v_k}(\tau_a^{v_k}) &= N_a^{v_k}(\tau_a^{v_k}) + \sum_{w \in \mathcal{V}_a \setminus \{v_k\}} \sum_{\iota=1}^{\gamma} \mathbb{1}[d_{\mathcal{G}}(v_k, w) = \iota] N_a^w(\tau_a^{v_k} - \iota) \\ &\geq N_a^{v_k}(\tau_a^{v_k}) + \sum_{j=1}^{k-1} \sum_{\iota=1}^{\gamma} \mathbb{1}[d_{\mathcal{G}}(v_k, v_j) = \iota] N_a^{v_j}(\tau_a^{v_k} - \iota). \end{aligned}$$

## 20:22 Flooding with Absorption

(b) follows from the following reasoning: first, we have that  $\tau_a^{v_k} + \gamma \geq \tau_a^{v_j}$  as at time  $\tau_a^{v_k} + \gamma$ , any remaining agent  $v_j$  should have at least  $\ell_a^{v_k} \geq \ell_a^{v_j}$  observations of  $a$ , i.e.,  $\tau_a^{v_j}$  is already reached. Thus, we have that

$$\mathbb{E}[N_a^{v_j}(\tau_a^{v_j})] - \mathbb{E}[N_a^{v_j}(\tau_a^{v_k} - \iota)] \leq \mathbb{E}[N_a^{v_j}(\tau_a^{v_j})] \leq \frac{4\alpha \log T}{(\Delta_a^{v_j})^2},$$

and

$$\mathbb{E}[N_a^{v_j}(\tau_a^{v_j})] - \mathbb{E}[N_a^{v_k}(\tau_a^{v_k} - \iota)] \leq \mathbb{E}[N_a^{v_j}(\tau_a^{v_k} + \gamma)] - \mathbb{E}[N_a^{v_j}(\tau_a^{v_k} - \iota)] \leq \iota + \gamma.$$

Recall that the group regret for our clique  $C$  is  $\sum_{k=1}^c \Delta_a^k \mathbb{E}[N_a^k(\tau_a^k)]$ , where for simplicity we denote  $\Delta_a^k := \Delta_a^k$ . The important observation is that to make the worst-case regret upper-bound, we must “allocate” the most number of pulls to the arms with the largest gap, i.e.,

$$\begin{aligned} & \sum_{k=1}^c \Delta_a^k \mathbb{E}[N_a^k(\tau_a^k)] \\ & \leq \Delta_a^1 \left( \frac{4\alpha \log T}{(\Delta_a^1)^2} \right) + \sum_{k=2}^c \Delta_a^k \left\{ \left( \frac{4\alpha \log T}{(\Delta_a^k)^2} - \frac{4\alpha \log T}{(\Delta_a^{k-1})^2} \right) + \min \left( 2\gamma, \frac{4\alpha \log T}{(\Delta_a^{k-1})^2} \right) \right\} \\ & \stackrel{(*)}{=} \frac{4\alpha \log T}{\Delta_a^1} + \left( \Delta_a^c \frac{4\alpha \log T}{(\Delta_a^c)^2} - \Delta_a^1 \frac{4\alpha \log T}{(\Delta_a^1)^2} \right) + \sum_{k=1}^{c-1} \frac{4\alpha \log T}{(\Delta_a^k)^2} (\Delta_a^k - \Delta_a^{k+1}) + \sum_{k=2}^c \Delta_a^k \min \left( 2\gamma, \frac{4\alpha \log T}{(\Delta_a^{k-1})^2} \right) \\ & \leq \frac{4\alpha \log T}{\Delta_a^c} + \int_{\Delta_a^c}^{\Delta_a^1} \frac{4\alpha \log T}{z^2} dz + \sum_{k=2}^{c+1} \Delta_a^{k-1} \min \left( 2\gamma, \frac{4\alpha \log T}{(\Delta_a^{k-1})^2} \right) \\ & = \frac{8\alpha \log T}{\Delta_a^c} - \frac{4\alpha \log T}{\Delta_a^1} + \sum_{k=1}^c \Delta_a^k \min \left( 2\gamma, \frac{4\alpha \log T}{(\Delta_a^k)^2} \right) \\ & = 4\alpha \log T \left( \frac{2}{\min_{v \in C} \Delta_a^v} - \frac{1}{\max_{v \in C} \Delta_a^v} \right) + \sum_{k=1}^c \Delta_a^k \min \left( 2\gamma, \frac{4\alpha \log T}{(\Delta_a^k)^2} \right), \end{aligned}$$

where  $(*)$  follows from the Abel transformation.

Thus,

$$\begin{aligned} & \sum_{v \in V_a} \Delta_a^v \mathbb{E}[N_a^v(\tau_a^v)] \\ & = \min_{C_{-a}} \sum_{C \in C_{-a}} \sum_{v \in C} \Delta_a^v \mathbb{E}[N_a^v(\tau_a^v)] \\ & \leq \min_{C_{-a}} \sum_{C \in C_{-a}} \left\{ 4\alpha \log T \left( \frac{2}{\min_{v \in C} \Delta_a^v} - \frac{1}{\max_{v \in C} \Delta_a^v} \right) \right\} + \sum_{v \in V_{-a}} \Delta_a^v \min \left( 2\gamma, \frac{4\alpha \log T}{(\Delta_a^v)^2} \right) \\ & = \frac{4\alpha \log T}{\Delta_a^\gamma} + f_a(\gamma). \end{aligned}$$

## B.2 Proof of Lemma B.2

We start by noting that

$$\mathbb{P}[E_a^v(a) \cap \{t > \tau_a^v\}] \leq \mathbb{P}[A(t) \cap \{t > \tau_a^v\}] + \mathbb{P}[B(t) \cap \{t > \tau_a^v\}].$$

We consider  $\mathbb{P}[A(t) \cap \{t > \tau_a^v\}]$  first. For simplicity, denote  $M_\star^v(t) := M_{a_\star^v}^v(t)$ . The initialization phase of Algorithm 1 implies that

$$\sum_{\iota=0}^{\gamma} |\{w \in \mathcal{V}_a : d_{\mathcal{G}}(w, v) = \iota\}| \leq M_\star^v(t) \leq \sum_{\iota=0}^{\gamma} (t - \iota) |\{w \in \mathcal{V}_a : d_{\mathcal{G}}(w, v) = \iota\}|.$$



Denote  $[a, b] := \{[a], [a] + 1, \dots, [b]\}$ ,  $\mathcal{N}_\iota = \{w \in \mathcal{V}_a : d_{\mathcal{G}}(w, v) = \iota\}$ , and  $N_\iota := |\mathcal{N}_\iota|$ . Also, with a slight abuse of notation, here let us denote  $X_{a_\star}^u(k)$  to be the reward received by agent  $v$  when she pulls arm  $a_\star$  for the  $k$ -th time. Then,

$$\begin{aligned}
& \mathbb{P}[A(t) \cap \{t > \tau_a^v\}] \\
& \leq \mathbb{P} \left[ \exists \{s_\iota\}_{\iota \in [0, \gamma]}, s_\iota \in [N_\iota, (t - \iota)N_\iota] \quad \exists \{s_\iota(u)\}_{u \in \mathcal{N}_\iota} : \sum_{u \in \mathcal{N}_\iota} s_\iota(u) = s_\iota, \sum_{\iota=0}^{\gamma} s_\iota = s \right. \\
& \quad \left. \text{s.t. } \frac{1}{s} \sum_{\iota=0}^{\gamma} \sum_{u \in \mathcal{N}_\iota} \sum_{k=1}^{s_\iota(u)} X_{a_\star}^u(k) + \sqrt{\frac{\alpha \log t}{s}} \leq \mu_\star^v \right] \\
& \leq \mathbb{P} \left[ \exists \{s_\iota\}_{\iota \in [0, \gamma]}, s_\iota \in [N_\iota, (t - \iota)N_\iota] \quad \exists \{s_\iota(u)\}_{u \in \mathcal{N}_\iota} : \sum_{u \in \mathcal{N}_\iota} s_\iota(u) = s_\iota, \sum_{\iota=0}^{\gamma} s_\iota = s \right. \\
& \quad \left. \text{s.t. } \sum_{\iota=0}^{\gamma} \sum_{u \in \mathcal{N}_\iota} \sum_{k=1}^{s_\iota(u)} (X_{a_\star}^u(k) - \mu_\star^v) \leq -\sqrt{s\alpha \log t} \right] \\
& \leq \sum_{\iota=0}^{\gamma} \mathbb{P} \left[ \exists s_\iota \in [N_\iota, (t - \iota)N_\iota] \quad \text{s.t. } \sum_{k=1}^{s_\iota} Y(k) \leq -\sqrt{s_\iota(\gamma + 1)\alpha \log t} \right],
\end{aligned}$$

where  $Y(k) \stackrel{d}{=} X_{a_\star}^u(k) - \mu_\star^v$  is *i.i.d.*  $\sigma$ -subGaussian random variable with  $\mathbb{E}[Y(k)] = 0$ .

Before moving forward, we recall a maximal-type concentration result:

► **Theorem B.3** (Theorem 9.2 of [39]). *Let  $Y_1, \dots, Y_n$  be a sequence of independent  $\sigma$ -subGaussian random variables, and let  $S_t = \sum_{s=1}^t Y_s$ . Then, for any  $\varepsilon > 0$ ,*

$$\mathbb{P}[\exists t \leq n : S_t \geq \varepsilon] \leq \exp\left(-\frac{\varepsilon^2}{2n\sigma^2}\right). \quad (12)$$

Using the above concentration result as well as the peeling argument on a geometric grid [8, 34], we have that for any  $\{\beta_\iota\}_{\iota \in [0, \gamma]}$  with  $\beta_\iota \in \left(\frac{2\sigma^2}{(\gamma+1)\alpha}, 1\right)$ ,

$$\begin{aligned}
& \mathbb{P}[A(t)] \\
& \leq \sum_{\iota=0}^{\gamma} \sum_{j=0}^{\frac{\log(t-\iota)}{\log(1/\beta_\iota)}} \mathbb{P} \left[ \exists s_\iota \in [N_\iota \beta_\iota^{j+1}(t - \iota), N_\iota \beta_\iota^j(t - \iota)] \quad \text{s.t. } \sum_{k=1}^{s_\iota} Y(k) \leq -\sqrt{s_\iota(\gamma + 1)\alpha \log t} \right] \\
& \leq \sum_{\iota=0}^{\gamma} \sum_{j=0}^{\frac{\log(t-\iota)}{\log(1/\beta_\iota)}} \mathbb{P} \left[ \exists s_\iota \in [N_\iota \beta_\iota^{j+1}(t - \iota), N_\iota \beta_\iota^j(t - \iota)] \quad \text{s.t. } \sum_{k=1}^{s_\iota} Y(k) \leq -\sqrt{N_\iota \beta_\iota^{j+1}(t - \iota)(\gamma + 1)\alpha \log t} \right] \\
& \leq \sum_{\iota=0}^{\gamma} \sum_{j=0}^{\frac{\log(t-\iota)}{\log(1/\beta_\iota)}} \mathbb{P} \left[ \exists s_\iota \in [1, N_\iota \beta_\iota^j(t - \iota)] \quad \text{s.t. } \sum_{k=1}^{s_\iota} Y(k) \leq -\sqrt{N_\iota \beta_\iota^{j+1}(t - \iota)(\gamma + 1)\alpha \log t} \right] \\
& \leq \sum_{\iota=0}^{\gamma} \sum_{j=0}^{\frac{\log(t-\iota)}{\log(1/\beta_\iota)}} \exp\left(-\frac{N_\iota \beta_\iota^{j+1}(t - \iota)(\gamma + 1)\alpha \log t}{2\sigma^2 N_\iota \beta_\iota^j(t - \iota)}\right) \\
& = \sum_{\iota=0}^{\gamma} \left(1 + \frac{\log(t - \iota)}{\log(1/\beta_\iota)}\right) t^{-\frac{\beta_\iota(\gamma+1)\alpha}{2\sigma^2}} \\
& \leq \sum_{\iota=0}^{\gamma} \left(1 + \frac{\log(t - \iota)}{\log(1/\beta_\iota)}\right) (t - \iota)^{-\frac{\beta_\iota(\gamma+1)\alpha}{2\sigma^2}}.
\end{aligned}$$

One can show the same for  $\mathbb{P}[B(t)]$ . (Note how both of them do not depend on any graph theoretical quantities).

Then,

$$\begin{aligned}
\sum_{t=1}^T \mathbb{P}[E_t^v(a) \cap \{t > \tau_a^v\}] &\leq 2 \sum_{t=\tau_a^v+1}^T \sum_{\iota=0}^{\gamma} \left(1 + \frac{\log(t-\iota)}{\log(1/\beta_\iota)}\right) (t-\iota)^{-\frac{\beta_\iota(\gamma+1)\alpha}{2\sigma^2}} \\
&\leq 2 \sum_{\iota=0}^{\gamma} \sum_{t=\iota+1}^{\infty} \left(1 + \frac{\log(t-\iota)}{\log(1/\beta_\iota)}\right) (t-\iota)^{-\frac{\beta_\iota(\gamma+1)\alpha}{2\sigma^2}} \\
&\leq 2 \sum_{\iota=0}^{\gamma} \int_{\iota+1}^{\infty} \left(1 + \frac{\log(t-\iota)}{\log(1/\beta_\iota)}\right) (t-\iota)^{-\frac{\beta_\iota(\gamma+1)\alpha}{2\sigma^2}} dt \\
&= 2 \sum_{\iota=0}^{\gamma} \int_1^{\infty} \left(1 + \frac{\log t}{\log(1/\beta_\iota)}\right) t^{-\frac{\beta_\iota(\gamma+1)\alpha}{2\sigma^2}} dt \\
&= \frac{8(\gamma+1)}{\left(\frac{\beta(\gamma+1)\alpha}{2\sigma^2} - 1\right)^2 \log(1/\beta)},
\end{aligned}$$

where we've set  $\beta_\iota = \beta$  for all  $\iota$ 's.

Following [8], we choose  $\beta = \frac{4\sigma^2}{\gamma+1} \frac{1}{\alpha+1/2}$ . Then,

$$\sum_{t=1}^T \mathbb{P}[E_t^v(a) \cap \{t > \tau_a^v\}] \leq \left(\frac{\alpha+1/2}{\alpha-1/2}\right)^2 \frac{8(\gamma+1)}{\log \frac{(\gamma+1)(\alpha+1/2)}{4\sigma^2}}.$$

## C Additional Discussions on the Theoretical Results

### C.1 Main Technical Challenges in the Proof of Theorem 3.2 – Regret Upper Bound

One might ask whether it is possible to use a star decomposition-type argument for our setting similar to [34], which could give us an improvement from  $\theta([\mathcal{G}^\gamma]_{-a})$  to  $\alpha([\mathcal{G}^\gamma]_{-a})$ . We believe this is *not* possible, and the reason is as follows.

We note that the main technical challenge is that unlike the homogeneous settings [50, 34], the agents and the arms are intertwined: In a homogeneous setting, one could rewrite the regret as

$$R(T) = \sum_{v \in V} \sum_{a \in \mathcal{K}} \Delta_a^{Kolla} \mathbb{E}[N_a^v(T)] = \sum_{a \in \mathcal{K}} \Delta_a^{Kolla} \sum_{v \in V} \mathbb{E}[N_a^v(T)],$$

i.e., the regret can be decomposed such that one only needs to upper bound the number of visitations of each agent  $v$ . Thus with a star decomposition, for each star, it can be easily seen that the number of visitations of the leaf agents is precisely that of the center agent, allowing for us to further decompose  $\sum_{v \in V} \mathbb{E}[N_a^v(T)]$  to sum of  $\mathbb{E}[M_a^{v^{center}}(T)]$  over all center agents  $v^{center}$ . This is at the core of dealing with more general homogeneous settings such as when there is a communication network [34], possibly with faults [50].

Such a decomposition is *not* possible when the agents are heterogeneous, as the maximal suboptimality gap of  $v$ ,  $\Delta_a^v$ , is agent-dependent. To deal with such heterogeneity in the full information sharing (fully-connected graph) setting, Yang et al. [68] first ordered the agents according to  $\Delta_a^v$ , then bounded the *cumulative* number of times suboptimal arm  $a$  is visited by agents  $v_1, \dots, v_k \in \mathcal{V}_{-a}$  via the design of UCB algorithm. Then, based on the intuition that the worst-case regret bound occurs when the arm  $a$  with the highest  $\Delta_a^v$  is visited at its

maximum, the final regret bound is derived via the Abel transformation. The fact that all agents have access to all other agents' information is crucial in this proof idea. This can be seen from a very simple example; consider a situation in which agent  $v_1$  has a very difficult problem (very small  $\Delta_a^{v_1}$ ) but has lots of connections, and agent  $v_2$  has a somewhat easy problem but has few connections. In this case, from the collaboration, it may be that  $v_1$  learns faster than  $v_2$ , which can impact the ordering of the agents, which in turn impacts the whole Abel transformation-based argument. In a sense, our proof combines these two ideas. We start with a clique decomposition, instead of the star decomposition, of the graph in order to upper-bound the group regret with the sum of regrets of each clique. Then, we apply the Abel transformation-type argument to each clique.

Lastly, we remark that our results can also be easily extended to the setting where the agents asynchronously pull the arms, i.e., each agent  $v$  pulls arms at every  $\omega_v$  round, with  $\omega_v \geq 1$ . It would be an interesting future direction if we could further reduce the communication cost in the asynchronous case based on ideas from recent works [16, 69, 68].

## C.2 Extending Theorem 5.2 – Regret Lower Bound

In order to match the lower bound to the upper bound in terms of graph topology and arm distribution-dependent constant as well, there are two avenues, both of which are inspired by Kolla et al [34], who consider the homogeneous setting with a general graph. On the one hand, it might be valuable to consider a Follow-Your-Leader-type policy, which would tighten the upper bound to match our lower bound. However, it is not clear how to choose the leaders in our heterogeneous setting, let alone how to ensure that followers get relevant information in terms of arms. Another way is to consider a more restricted class of policies, namely NAIC (non-altruistic & individually consistent) policies, which would tighten the lower bound to match our upper bound. Extending such notion to our setting of heterogeneous bandits over a graph *while* taking the communication protocol into account, e.g., whether we use the entire graph (FLOODING) or we use part of the graph (FWA), and seeing whether we can match the lower bound up to the derived upper bounds, is another interesting future direction. On a separate note, deriving a minimax lower bound, as done in Madhushani et al. [49], for our setting is also an interesting future direction.