

The London School of Economics and Political Science

*Regret-Minimization Algorithms for Multi-Agent
Cooperative Learning Systems*

Jialin Yi

A thesis submitted to the Department of Statistics of the London School of
Economics and Political Science for the degree of Doctor of Philosophy, London,
February 2023.

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgment is made. This thesis may not be reproduced without the prior written consent of the author.

I warrant that this authorization does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 22710 words.

Statement of co-authored work

I confirm that the work in Chapter 3 was jointly co-authored with Professor Milan Vojnović and I contributed 80% of this work (problem setting, algorithms, main proofs, and code implementations). I confirm that the work in Chapter 4 was jointly co-authored with Professor Milan Vojnović and I contributed 90% of this work (problem setting, algorithms, proofs, and code implementations). I confirm that the work in Chapter 5 was authored by me only. I confirm that the work in Chapter 6 was jointly co-authored with Professor Milan Vojnović and I contributed 40% of this work (algorithms, partial proofs, and code implementations).

Acknowledgment

I would like to express my heartfelt gratitude to Milan Vojnović and Clifford Lam, my advisors, who have provided me with invaluable guidance, support, and encouragement throughout my PhD journey. Their unwavering belief in my abilities and their relentless pursuit of excellence have been an inspiration to me. Especially, I want to give my warmest thank to Milan Vojnović for his advice and help on research directions, academic writing, presentation skills, and my career path; to Clément Calauzènes, Chuan Luo and Vianney Perchet, for their countless help on my research projects and job search. The internship at Microsoft Research at Asia was a fantastic experience to me because of the guidance and help from Hang Dong and Chuan Luo.

I would also like to extend my sincere thanks to Laurent Massoulié and Chengchun Shi, for their insightful comments, constructive criticism, and valuable suggestions. Their invaluable contributions have greatly improved the quality of this thesis.

I am deeply grateful to Department of Statistics at London School of Economics and Political Science for providing me with the scholarship, resources and facilities necessary to complete this research. I really appreciate the help from my colleagues in Department of Statistics as well.

I would like to dedicate this thesis to my parents, Li Huang and Xixiang Yi, for their ever-lasting love and support. Their encouragement and trust have been my driving force.

Finally, I want to extend my deepest appreciation to my forever friends, Chuan Cheng, Zezhun Chen, Junyi Liao, Jiaqi Lyu, Zhaoxi Li, Yaqian Wu and Yu Yi. Getting a PhD is a long journey during which I not only received precious memories but also lost important ones. It were those friends whose accompany and care saved me during the hardest time in my PhD journey.

Thank you all.

Abstract

A Multi-Agent Cooperative Learning (MACL) system is an artificial intelligence (AI) system where multiple learning agents work together to complete a common task. Recent empirical success of MACL systems in various domains (e.g. traffic control, cloud computing, robotics) has sparked active research into the design and analysis of MACL systems for sequential decision making problems. One important metric of the learning algorithm for decision making problems is its *regret*, i.e. the difference between the highest achievable reward and the actual reward that the algorithm gains. The design and development of a MACL system with low-regret learning algorithms can create huge economic values.

In this thesis, I analyze MACL systems for different sequential decision making problems. Concretely, the Chapter 3 and 4 investigate the cooperative multi-agent multi-armed bandit problems, with full-information or bandit feedback, in which multiple learning agents can exchange their information through a communication network and the agents can only observe the rewards of the actions they choose. Chapter 5 considers the communication-regret trade-off for online convex optimization in the distributed setting. Chapter 6 discusses how to form high-productive teams for agents based on their unknown but fixed types using adaptive incremental matchings.

For the above problems, I present the regret lower bounds for feasible learning algorithms and provide the efficient algorithms to achieve this bound. The regret bounds I present in Chapter 3, 4 and 5 quantify how the regret depends on the connectivity of the communication network and the communication delay, thus giving useful guidance on design of the communication protocol in MACL systems.

Contents

Chapter 1. Introduction	11
1.1 The structure of this thesis	13
1.2 Contributions	13
1.3 Numerical results	15
Chapter 2. Background	16
2.1 Multi-armed bandit (MAB) problems	16
2.1.1 Concepts	17
2.1.2 Regret minimization	17
2.2 Extensions to MAB problems	20
2.2.1 Multi-agent MAB problems	21
2.2.2 Delayed feedback	22
2.2.3 Combinatorial bandits	22
2.3 Differential equation approximation of discrete stochastic processes	22
Chapter 3. Cooperative Non-stochastic Multi-Armed Bandits	25
3.1 Introduction	25
3.1.1 Related work	26
3.1.2 Organization of this chapter	28
3.2 Problem formulation	28
3.3 Algorithms and regret upper bounds	29
3.3.1 A center-based cooperative algorithm	29
3.3.2 A decentralized cooperative algorithm	32
3.4 Regret lower bounds	34
3.5 Numerical experiments	35
3.5.1 The effect of the number of arms	36
3.5.2 The effect of graph degree on CFTRL	37
3.5.3 The effect of delay on CFTRL and DFTRL	37
3.5.4 The effect of graph sparsity	38
3.6 Proofs	38
3.6.1 An auxiliary lemma	38
3.6.2 Proof of Lemma 3.3.2	39
3.6.3 Proof of Lemma 3.3.3	41
3.6.4 Proof of Theorem 3.3.1	44
3.6.5 Proof of Lemma 3.3.5	49

3.6.6	Proof of Theorem 3.3.4	50
3.6.7	Proof of Theorem 3.4.1	50
3.6.8	CFTRL and the center-based Exp3 when $d = 1$	51
3.6.9	CFTRL and the lower bound when $d = 1$	52
Chapter 4. Doubly Adversarial Federated Bandits		54
4.1	Introduction	54
4.1.1	Related work	55
4.1.2	Organization of this chapter	57
4.2	Problem setting	57
4.3	Lower bounds	58
4.4	FEDEXP3: a federated regret-minimization algorithm	60
4.5	Numerical experiments	64
4.5.1	Synthetic datasets	64
4.5.2	MovieLens dataset: recommending popular movie genres	65
4.6	Proofs	66
4.6.1	Proof of Theorem 4.3.1	67
4.6.2	Proof of Theorem 4.3.2	71
4.6.3	Auxiliary lemmas	72
4.6.4	Proof of Theorem 4.4.1	76
Chapter 5. Communication-Efficient Federated Online Convex Optimization		79
5.1	Introduction	79
5.1.1	Related work	80
5.1.2	Organization of this chapter	81
5.2	Federated online convex optimization	81
5.3	FEDOCO: a communication-efficient stochastic regret-minimization algorithm	83
5.4	Trade-off between communication complexity and regret	84
5.5	Proofs	86
5.5.1	Proof of Theorem 5.4.1	86
Chapter 6. Greedy Bayes Incremental Matching Strategies		89
6.1	Introduction	89
6.1.1	Related work	90
6.1.2	Organization of this chapter	92

6.2	Problem formulation	92
6.3	Algorithms and regret bounds	95
6.3.1	Maximum value function (OR)	96
6.3.2	Minimum value function (AND)	98
6.4	Proofs	106
6.4.1	Proof of Lemma 6.3.1	106
6.4.2	Proof of Theorem 6.3.2	111
6.4.3	Proof of Lemma 6.3.3	117
6.4.4	Proof of Proposition 6.3.4	119
6.4.5	Proof of Theorem 6.3.5	121
6.4.6	Properties of $a_{n,p}$, $b_{n,p}$ and $c_{n,p}$	128
Chapter 7. Conclusions		137
Bibliography		147

List of Tables

1	Regret lower bounds for different value functions.	95
---	--	----

List of Figures

1	Average regret R_T versus T for different algorithms on a 2-regular graph with $N = 3$ agents and edge-delay $d = 1$, and varied number of arms: (top) $K = 20$ (middle) $K = 30$, and (top) $K = 40$. We used 10 independent simulation runs.	36
2	Average regret of CFTRL versus graph degree r , for a r -regular graph with $N = 6$ nodes, $K = 10$ and $d = 1$	37
3	Average regret of CFTRL and DFTRL versus the edge-delay d on a star regular graph with $N = 20$ and $K = 3$	37
4	Average regret R_T versus the number of nodes N for sparse Erdős–Rényi random graphs, for CFTRL and center-based Exp3 algorithms.	38
5	(Top) The average cumulative regret, i.e. $\sum_{v \in \mathcal{V}} R_T^v / N$, versus T , for FEDEXP3 and Exp3 on the grid graph. (Down) The average cumulative regret versus $(1 - \sigma_2(W))^{-1/3}$ for FEDEXP3 on different networks at $T = 3000$	65
6	The average cumulative regret versus horizon time for FEDEXP3 and GUCB in the movie recommendation setting with the communication networks: (top) complete graph, (mid) the grid network, and (down) RGG(0.5).	66
7	An example of an incremental matching sequence: for example, the second matching is derived from the first matching by re-matching node pairs (2,3) and (4,5).	94
8	The asymptotic normalized regret for maximum value function, in the limit of large number of nodes n , versus population parameter p . 98	
9	An illustration of evolution of matching sets of greedy Bayes matching strategy for AND value function.	99
10	Two possible outcomes: merge or remove two matching sets.	103
11	The number of rounds until greedy Bayes incremental matching algorithm terminates for minimum (Boolean AND) value function versus population parameter p , for different numbers of nodes n	105
12	Regret of the greedy Bayes algorithm for the MIN value function versus the population parameter p , for different values of the number of nodes n	107
13	Matchings created in super-epoch s	122

Chapter 1

Introduction

A Multi-Agent Cooperative Learning (MACL) system is an artificial intelligence (AI) system where multiple learning agents work together to complete a common task. These agents communicate with and learn from each other, improving their performance over time. The cooperative nature of these systems is beneficial for solving complex problems, as the agents can divide tasks, share information and thus accelerate their learning processes.

Recently, MACL systems have demonstrated successful applications in various domains, including traffic control (Chu et al., 2019), cloud computing (Mao et al., 2022), robotics (Wu et al., 2022) and so on. This empirical success has sparked active research into the design and analysis of MACL systems for sequential decision making problems. Despite significant advancements in the field, there remain challenges to overcome, even for some fundamental sequential decision making problems: the Multi-Armed Bandit (MAB) problem (Auer et al., 2002a,b) and the Online Convex Optimization (OCO) problem (Hazan et al., 2016). In particular, when learning agents are connected by a communication network,

- the *optimal* distributed regret-minimization algorithm,
- the effect of the communication network's *connectivity* on the regret of an algorithm,
- the effect of *delays* in communication on the regret of an algorithm,
- and the trade-off between the regret and the *communication complexity*

are not yet fully understood. Addressing these questions from a theoretical perspective is important and challenging, and this thesis presents a set of efforts towards this aim.

The study begins with a canonical setting in non-stochastic MAB problems which involve selecting the most rewarding action from a fixed set of actions over a fixed time horizon. In the single-agent non-stochastic MAB problem, near-optimal (Exp3) and optimal (Tsallis-INF) algorithms have been proposed by Auer

et al. (2002b) and Zimmert and Seldin (2021), respectively, as is introduced in Chapter 2. The natural extension to the cooperative multi-agent setting involves each agent individually selecting an action, followed by exchanging information among them. Based on rewards received by agents, cooperative multi-agent MAB problems can be divided into two settings: the homogeneous reward setting and the heterogeneous reward setting.

The homogeneous reward setting involves a scenario where multiple agents receive the same reward for choosing the same action at the same time step, which is commonly observed in the distributed training of reinforcement learning algorithms (Babaeizadeh et al., 2016). In this setting, all agents face the same MAB problem, and the exploration phase of the learning algorithm is executed in parallel by each agent. The challenge lies in devising efficient cooperative algorithms under communication delays.

In contrast to the homogeneous reward setting, the heterogeneous reward setting, motivated by federated learning systems (Kairouz et al., 2021), involves a scenario where multiple agents may choose the same action at the same time step, but receive *different* rewards. The objective for each agent is to determine the globally optimal action in hindsight based on the collective experience of all agents, which requires communication among agents. To achieve this goal, agents must reach a consensus on the globally optimal action by exchanging their private information. The challenge in this setting lies in the design of an appropriate mechanism for exchanging private information, which allows for cooperative learning while preserving the privacy of each individual agent.

Despite borrowing algorithms and analysis from the single-agent sequential decision making, the multi-agent cooperative learning introduces unique challenges not present in the single-agent case. One such challenge is the trade-off between the regret of a cooperative online learning algorithm and the communication complexity incurred by that algorithm (Agarwal et al., 2022; Wan et al., 2020).

Except for continuous optimization problems, efficiently finding a maximum value matching for a given set of nodes based on observations of the previously proposed matching over a sequence of rounds is another fundamental sequential learning problem. This problem arises in many applications including team formation, online social networks, online gaming, online labor, and other online platforms (Kamble et al., 2018; Mehta et al., 2020). The challenge lies in the design of optimal regret-minimization algorithms to address this problem.

The study of MACL systems encompasses a broad area of research, which is beyond the scope of this thesis. This thesis focuses on the theory of non-stochastic multi-agent MAB problems, federated OCO problems and the sequential learning of a maximum value matching, while acknowledging the existence of other significant topics such as stochastic multi-agent MAB problems, which are not discussed in this thesis.

1.1 The structure of this thesis

The preliminary knowledge related to non-stochastic MAB problems is presented in Chapter 2. Subsequently, Chapter 3 and 4 delve deeper into cooperative non-stochastic multi-agent MAB problems, each focusing on a specific setting. Chapter 3 covers the homogeneous reward setting, and Chapter 4 examines the heterogeneous reward setting. Chapter 5 is dedicated to the analysis of the trade-off between the regret and the communication complexity in federated online convex optimization problems. Chapter 6 investigates how to find a maximum value matching in a sequential manner under the minimum change constraint. Finally, the thesis concludes in Chapter 7 with a summary of the findings, and suggestions for future research directions.

1.2 Contributions

The contributions of this thesis can be summarized as follows:

Cooperative non-stochastic multi-armed bandits (Chapter 3) We consider a non-stochastic multi-agent multi-armed bandit problem with agents collaborating via a communication network with delays. We show a lower bound for individual regret of all agents. We show that with suitable regularizers and communication protocols, a collaborative multi-agent *follow-the-regularized-leader* (FTRL) algorithm has an individual regret upper bound that matches the lower bound up to a constant factor when the number of arms is large enough relative to degrees of agents in the communication graph. We also show that an FTRL algorithm with a suitable regularizer is regret optimal with respect to the scaling with the edge-delay parameter. We present numerical experiments validating our theoretical results and demonstrate cases when our algorithms outperform previously proposed algorithms. This chapter is based on joint work with Milan Vojnović (Yi and Vojnovic, 2022).

Doubly adversarial federated bandits (Chapter 4) We study a new non-stochastic federated multi-armed bandit problem with multiple agents collaborating via a communication network. The losses of the arms are assigned by an oblivious adversary that specifies the loss of each arm not only for each time step but also for each agent, which we call “doubly adversarial”. In this setting, different agents may choose the same arm in the same time step but observe different feedback. The goal of each agent is to find a globally best arm in hindsight that has the lowest cumulative loss averaged over all agents, which necessitates the communication among agents. We provide regret lower bounds for any federated bandit algorithm under different settings, when agents have access to full-information feedback, or the bandit feedback. The regret lower bound for the full-information setting also applies to distributed online linear optimization problems, and, to the best of our knowledge, is the first lower bound result for this problem. For the bandit feedback setting, we propose a near-optimal federated bandit algorithm called FEDEXP3. Our algorithm gives a positive answer to an open question proposed in Cesa-Bianchi et al. (2016): FEDEXP3 can guarantee a sub-linear regret without exchanging sequences of selected arm identities or loss sequences among agents. We also provide numerical evaluations of our algorithm to validate our theoretical results and demonstrate its effectiveness on synthetic and real-world datasets. This chapter is based on joint work with Milan Vojnović (Yi and Vojnović, 2023).

Communication-regret trade-off in federated online learning (Chapter 5) The increasing deployment of large-scale federated learning systems has led to a growing interest in research on distributed online convex optimization algorithms (Wangni et al., 2018; Reisizadeh et al., 2020). However, a significant challenge in practical implementation of these algorithms is the high communication complexity which is measured in the cumulative number of messages exchanged between agents. It remains an open question whether it is possible for a distributed online convex optimization algorithm, over T time steps, to concurrently achieve $o(T)$ cumulative regret and $o(\sqrt{T})$ communication complexity. In this chapter, we address this challenge by introducing a communication-efficient stochastic algorithm, FEDOCO.

Greedy Bayes incremental matching strategies (Chapter 6) We consider the problem of finding a maximum value matching in a given set of nodes based on

observed evaluations of node pairs, where the value of a node pair is according to a function of hidden node types—we refer to as a value function. This is considered in a sequential learning setting where in each round, the learner selects a matching of nodes and observes the value of each matched pair of nodes. Importantly, the sequence of matchings must be incremental, meaning that in each round, the matching differs from the matching in the previous round for at most four pairs, i.e. a new matching is derived from the previous matching by re-matching nodes of any two previously matched pairs with each other. This problem is motivated by situations in which matchings need to be explored by making small incremental alternations of existing matches over time. We consider the efficiency of learning with respect to regret minimization for greedy Bayes incremental matching algorithms, which in each round select a matching that has the maximum expected reward in the next round based on the posterior distribution of node types, conditional on data observed up to the current round. Our focus is on the case of binary node types and understanding the efficiency of learning for different value functions, including all symmetric Boolean functions. We identify greedy Bayes incremental matching policies for different value functions and characterize how the regret depends on the population size and distribution of node types. This chapter is based on joint work with Milan Vojnović.

1.3 Numerical results

All experiments in this thesis are run on a desktop with AMD Ryzen 5 2600 Six-Core Processor and 16GB memory. Each experiment took less than 6 hours to finish.

The code is written in Python 3.9 and uses Numpy package (Harris et al., 2020) and NetworkX package (Hagberg et al., 2008) for numerical calculations and graph operations. The Numpy package and the NetworkX package are distributed with the BSD license.

Chapter 2

Background

2.1 Multi-armed bandit (MAB) problems

The multi-armed bandit problem stems from the concept of a “bandit” in gambling. It’s a term used to describe a one-armed bandit machine, also known as a slot machine, where a player pulls a lever to obtain a *reward*. The multi-armed bandit problem models this situation as one where there are multiple levers, (*arms* or *actions*) to choose from and the player wants to maximize his or her rewards over a series of plays by balancing exploration (trying out different arms to learn their rewards) with exploitation (using the knowledge of the best arm to maximize rewards). The problem has been formalized in the field of reinforcement learning and decision theory and has applications in various fields such as online advertising, clinical trials and recommender systems (Li et al., 2010; Lattimore and Szepesvári, 2020).

Depending on the assumption about how the observed rewards are generated, MAB problems can be classified into two main types:

- stochastic bandits: the reward from each arm at any given time is an independent and identically distributed sample from the underlying distribution,
- adversarial bandits: the reward from each arm at any given time is determined by an adversary.

In this thesis, the focus is on adversarial bandits, which are commonly conceptualized as repeated two-player zero-sum games between an agent and an adversary, each with distinct objectives and strategies. Given its strong association with online linear optimization, the adversarial bandit community often refers to the terms “arms” as “actions” and the inverse of “rewards” as “losses”. Consequently, the objective of reward maximization is transformed into loss minimization. This convention is adopted in this thesis.

2.1.1 Concepts

We use $\ell(i)$ to denote the i -th coordinate of the vector ℓ . And \mathcal{S}^* be the set of all sequences whose entries are from the set \mathcal{S} , i.e. $\mathcal{S}^* = \{(s_1, s_2, \dots) : s_i \in \mathcal{S} \text{ for all } i\}$.

Environment Let \mathcal{A} be a set of actions with $|\mathcal{A}| = K$ and T be the number of repeated rounds, i.e. horizon. A K -armed adversarial bandit environment (or instance) is an arbitrary sequence of loss vectors $\{\ell_t\}_{t=1}^T$, where $\ell_t \in [0, 1]^K$. At each time step t , the agent chooses a distribution p_t over the actions. Then an action $a_t \in \mathcal{A}$ is sampled from p_t and the agent receive a loss of $\ell_t(a_t)$.

Policy The agent in the adversarial bandit environment is allowed to make the decision based on the historical information $\mathcal{F}_t = \{a_1, \ell_1(a_1), a_2, \ell_2(a_2), \dots, a_{t-1}, \ell_{t-1}(a_{t-1})\} \in (\mathcal{A} \times [0, 1])^*$. Let \mathcal{P}_{K-1} be the set of probability distributions on the K actions, i.e. a $K - 1$ -simplex $\{p \in [0, 1]^K : \sum_{i=1}^K p_i = 1\}$. A bandit algorithm, also called *policy*, is a function

$$\pi : (\mathcal{A} \times [0, 1])^* \rightarrow \mathcal{P}_{K-1}$$

mapping the historical information to a probability distribution over the actions.

Regret The performance of a policy π is measured by its *regret*, defined as the difference of the expected cumulative loss incurred by following the policy and the expected cumulative loss of a best *fixed* arm in hindsight, i.e.

$$R_T(\pi, \{\ell_t\}) = \mathbb{E} \left[\sum_{t=1}^T \ell_t(a_t) - \min_{i \in \mathcal{A}} \left\{ \sum_{t=1}^T \ell_t(i) \right\} \right]$$

where the expectation is taken over chosen actions under algorithm π on instance L . We will abbreviate $R_T(\pi, \{\ell_t\})$ as R_T when the algorithm π and adversarial bandit environment $\{\ell_t\}$ have no ambiguity in the context.

2.1.2 Regret minimization

The objective of the designer of a bandit algorithm is to identify a policy π that minimizes regret R_T for the agent in any bandit environment $\{\ell_t\}$ selected by the

adversary. This requires the development of a policy that effectively balances *exploration* and *exploitation* in a manner that results in a small regret for the agent over time. On the one hand, exploration is necessary to gather information about different actions and estimate their expected losses. On the other hand, exploitation is necessary to minimize the loss by selecting the best-performing action so far. The difficulty lies in finding the optimal trade-off between exploration and exploitation in an uncertain and dynamic environment.

Minimax lower bound From the optimization perspective, the design of a bandit algorithm with best worst-case regret can be formulated as a minimax optimization problem, i.e.,

$$\min_{\pi \in \Pi} \max_{\{\ell_t \in [0,1]^K\}} R_T(\pi, \{\ell_t\}).$$

where Π is the set of all policies.

We present a lower bound on $\max_{\{\ell_t \in [0,1]^K\}} R_T(\pi, \{\ell_t\})$ which applies to all policy $\pi \in \Pi$, which is so-called *minimax lower bound*.

Theorem 2.1.1 (Lattimore and Szepesvári (2020)). *Let $c, C > 0$ be sufficiently small/large universal constants and $K \geq 2, T \geq 1$ and $\delta \in (0,1)$ be such that $T \geq CK \log(1/(2\delta))$. Then there exists a loss sequence $\ell_t \in [0,1]^K$ such that*

$$\mathbb{P} \left[\sum_{t=1}^T \ell_t(a_t) - \min_{i \in \mathcal{A}} \left\{ \sum_{t=1}^T \ell_t(i) \right\} \geq c \sqrt{KT \log \left(\frac{1}{2\delta} \right)} \right] \geq \delta.$$

Theorem 2.1.1 implies a lower bound on R_T :

$$\begin{aligned} R_T &= \mathbb{E} \left[\sum_{t=1}^T \ell_t(a_t) - \min_{i \in \mathcal{A}} \left\{ \sum_{t=1}^T \ell_t(i) \right\} \right] \\ &= \int_0^\infty \mathbb{P} \left[\sum_{t=1}^T \ell_t(a_t) - \min_{i \in \mathcal{A}} \left\{ \sum_{t=1}^T \ell_t(i) \right\} \geq x \right] dx \\ &\geq \frac{1}{2} \int_0^\infty \exp \left(-\frac{x^2}{c^2 KT} \right) dx \\ &= \Omega \left(\sqrt{KT} \right). \end{aligned}$$

Hence, for any policy $\pi \in \Pi$, there exists a bandit environment such that the regret of the policy is $\Omega \left(\sqrt{KT} \right)$.

Regret minimization algorithms The minimax lower bound in Theorem 2.1.1 provides a benchmark for evaluating regret-minimization algorithms and can be used to determine whether an algorithm is optimal or sub-optimal. Optimal algorithms achieve the $O(\sqrt{KT})$ order in minimax lower bound, while sub-optimal algorithms have a higher regret with an order of $O(\sqrt{KT \log K})$. We present a sub-optimal algorithm (Exp3) and an optimal algorithm (Tsallis-INF) with their regret upper bounds.

Algorithm 1: Exp3 (Auer et al., 2002b)

Input $\gamma \in (0, 1]$.
Initialization: $w_1(i) = 1$ for $i = 1, \dots, K$.
1 **for** each time step $t = 1, 2, \dots, T$ **do**
2 Set

$$p_t(i) = (1 - \gamma) \frac{w_t(i)}{\sum_{j=1}^K w_t(j)} + \gamma_t/K$$

 for $i = 1, \dots, K$;
3 Draw a_t randomly accordingly to the probabilities $p_t(1), \dots, p_t(K)$.;
4 Receive reward $\ell_t(a_t)$;
5 For $j = 1, 2, \dots, K$, set

$$\hat{\ell}_t(j) = \begin{cases} \ell_t(j)/p_t(j) & \text{if } j = a_t, \\ 0 & \text{otherwise} \end{cases}$$

$$w_{t+1}(j) = w_t(j) \exp(\gamma \hat{\ell}_t(j)/K);$$

6 **end**

Theorem 2.1.2 (Auer et al. (2002b)). *For any $K > 0$ and $T > 0$, suppose the agent follows Exp3 algorithm as described in Algorithm 1, with*

$$\gamma = \min \left\{ 1, \sqrt{\frac{K \ln K}{(e-1)T}} \right\},$$

the regret of the agent

$$R_T \leq 2.63 \sqrt{KT \log K}.$$

Theorem 2.1.3 (Zimmert and Seldin (2021)). *For any $K > 0$ and $T > 0$, suppose the agent follows Tsallis-INF algorithm as described in Algorithm 2, with*

$$\eta_t = 2\sqrt{\frac{1}{t}},$$

Algorithm 2: Tsallis-INF (Zimmert and Seldin, 2021)

Input : a sequence of learning rates $\{\eta_t \in (0, 1]\}$.
Initialization: $w_1(i) = 1$ for $i = 1, \dots, K$.

- 1 **for** each time step $t = 1, 2, \dots, T$ **do**
- 2 Set
$$p_t(i) = (1 - \gamma) \frac{w_t(i)}{\sum_{j=1}^K w_t(j)} + \gamma_t / K$$

 for $i = 1, \dots, K$;
- 3 Draw a_t randomly accordingly to the probabilities $p_t(1), \dots, p_t(K)$.;
- 4 Receive reward $\ell_t(a_t)$;
- 5 For $j = 1, 2, \dots, K$, set
$$\hat{\ell}_t(j) = \begin{cases} \ell_t(j) / p_t(j) & \text{if } j = a_t, \\ 0 & \text{otherwise} \end{cases}$$
$$w_t = \arg \max_{w \in \Delta^{K-1}} \left\langle w, - \sum_{s=1}^t \hat{\ell}_s \right\rangle + \frac{4}{\eta_t} \sum_i \sqrt{w_i};$$
- 6 **end**

the regret of the agent

$$R_T \leq 4\sqrt{KT} + 1.$$

2.2 Extensions to MAB problems

Recently, the basic MAB problem, which involves the exploration and exploitation of multiple arms with unknown rewards, has been extended to accommodate more complex scenarios. These modifications provide more nuanced models of real-world systems and offer new insights into the behavior of the MAB problem under different conditions (Lattimore and Szepesvári, 2020).

The extensions of most relevance to this thesis include the multi-agent setting (Cesa-Bianchi et al., 2019; Bar-On and Mansour, 2019; Dubey and Pentland, 2020; Zhu et al., 2021; Shi et al., 2021; Réda et al., 2022), the delayed feedback setting (Joulani et al., 2013; Zimmert and Seldin, 2020) and the combinatorial bandit setting (Cesa-Bianchi and Lugosi, 2012; Combes et al., 2015; Kveton et al., 2015; Chen et al., 2016; Rejwan and Mansour, 2020; Agarwal et al., 2021).

In this section, we provide an overview of previous studies that are relevant to our research, and elucidate unique contributions of this thesis in the context of existing literature.

2.2.1 Multi-agent MAB problems

Multi-agent cooperative adversarial bandits Cesa-Bianchi et al. (2016); Bar-On and Mansour (2019) studied the adversarial case where agents receive the same loss for the same action chosen at the same time step, whose algorithms require the agents to exchange their raw data with neighbors. In Chapter 3, we give a regret lower bound for any learning algorithm in which each agent can only communicate with their neighbors. We present a center-based algorithm whose regret upper bound matches the lower bound when the number of arms is large enough. We present an algorithm that has a regret upper bound with \sqrt{d} dependence on the delay d per edge, which is optimal. All our regret bounds are parametrized with the delay parameter d , which is unlike to Cesa-Bianchi et al. (2019) and Bar-On and Mansour (2019) which considered only the special when $d = 1$. In what follows we summarise our results in more details.

In these works, agents that choose the same action at the same time step receive the same reward or loss value and agents aggregate messages received from their neighbors. Our work in Chapter 4 relaxes this assumption by allowing agents to receive different losses even for the same action in a time step. Besides, we propose a new algorithm that uses a different aggregation of messages than in the aforementioned papers, which is based on distributed dual averaging method in Nesterov (2009); Xiao (2009); Duchi et al. (2011).

Federated bandits Solving bandit problems in the federated learning setting has gained attention in recent years. Dubey and Pentland (2020) and Huang et al. (2021) considered the linear contextual bandit problem and extended the LinUCB algorithm (Li et al., 2010) to the federated learning setting. Zhu et al. (2021) and Shi et al. (2021) studied a federated multi-armed bandit problem where the losses observed by different agents are i.i.d. samples from some common unknown distribution. Réda et al. (2022) considered the problem of identifying a globally best arm for multi-armed bandit problems in a centralized federated learning setting. All these works focus on the stochastic setting, i.e. the reward or loss of an arm is sampled from some unknown but fixed distribution. Our work in Chapter 4 considers the non-stochastic setting, i.e. losses are chosen by an oblivious adversary, which is a more appropriate assumption for non-stationary environments.

2.2.2 Delayed feedback

Multi-armed bandits with delayed feedback have been studied extensively in the single-agent setting (Joulani et al., 2013; Thune et al., 2019; Flaspohler et al., 2021). Specifically, Zimmert and Seldin (2020) considered a setting in which the agent has no prior knowledge about the delays and showed an optimal regret of $O(\sqrt{KT} + \sqrt{d \log(K)T})$ where d is the average delay over T time steps. We present in Chapter 3, in the multi-agent setting, a distributed learning algorithm whose regret upper bound can also achieve this optimal \sqrt{d} dependence on the delay per edge d .

2.2.3 Combinatorial bandits

For combinatorial bandits, in each round, the decision maker selects an arm that corresponds to a feasible subset of base arms. For example, given a set of N base arms, an arm is an arbitrary subset of base arms of cardinality K . Under full bandit feedback, the reward of the selected arm is observed, which is according to an underlying function of the corresponding base arms' values. For example, this may be sum of values (Cesa-Bianchi and Lugosi, 2012; Combes et al., 2015; Rejwan and Mansour, 2020), or some non-linear function of values (Chen et al., 2016; Agarwal et al., 2021). Under semi-bandit feedback, individual values of all selected base arms are observed. In our work in Chapter 5, we may think of node pairs to be base arms, matchings to be arms, and we have semi-bandit feedback. Alternatively, we may regard nodes to be base arms, matchings to be arms, and we have full-bandit feedback for each match in a matching. Note that requiring for sequence of matching selections to be according to an incremental matching sequence is not studied in the line of work on combinatorial bandits.

2.3 Differential equation approximation of discrete stochastic processes

Here we provide background results on differential equation approximation for discrete random processes by Wormald (1995), which we use in the proof of Theorem 6.3.2.

Consider a discrete random process on probability space Ω , denoted by X_0, X_1, \dots , where each X_t takes values in some set S . We use \mathcal{F}_t to denote (X_0, X_1, \dots, X_t) , the history of the process up to time t . For a function y of history

\mathcal{F}_t , random variable $y(\mathcal{F}_t)$ is denoted by Y_t .

We consider a sequence $\Omega_n, n = 1, 2, \dots$ of random processes. For a random variable X , we say $X = o(f(n))$ always if

$$\max\{x : \mathbb{P}[X = x] \neq 0\} = o(f(n)).$$

An event is said to hold *almost surely* if its probability in Ω_n is $1 - o(1)$. Let S_n^+ denote the set of all functions $f_n(x_0, x_1, \dots, x_t)$ where each $x_i \in S_n$, for $t \geq 0$.

A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is said to satisfy a Lipschitz condition on $D \subseteq \mathbb{R}^d$ if there exists a constant $L > 0$ such that $|f(x) - f(y)| \leq L\|x - y\|_1$ for all x and y in D .

Recall the definition of the extension of the solution to a differential equations

$$\frac{d}{dt}x(t) = f(t, x)$$

where f and f_x are continuous functions.

Definition 2.3.1 (extension of solutions). *Given two solutions $u : J \subset \mathbb{R} \rightarrow \mathbb{R}^{d-1}$ and $v : I \subset \mathbb{R} \rightarrow \mathbb{R}^{d-1}$, u is called an extension of v if $I \subset J$ and*

$$u(x) = v(x) \quad x \in I.$$

Let d be a fixed positive integer. For $i \in [d] = \{1, \dots, d\}$, let

$$y^{(i)} : \cup_n S_n^+ \rightarrow \mathbb{R}$$

and $f_i : \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ be such that for some constant C and all $i \in [d]$,

$$|y^{(i)}(\mathcal{F}_t)| \leq Cn \text{ for all } \mathcal{F}_t \in S_n^+ \text{ for all } n.$$

Assume that for some sequence $m(n)$, the following conditions hold:

(C1) there is a constant C' such that for all $t < m(n)$ and all $i \in [d]$, it holds always,

$$|Y_{t+1}^{(i)} - Y_t^{(i)}| \leq C'.$$

(C2) for all $i \in [d]$ and uniformly over $t < m(n)$, it holds always,

$$\mathbb{E}[Y_{t+1}^{(i)} - Y_t^{(i)} \mid \mathcal{F}_t] = f_i \left(\frac{t}{n}, \frac{Y_t^{(1)}}{n}, \dots, \frac{Y_t^{(d)}}{n} \right) + o(1).$$

(C3) for all $i \in [d]$, f_i is continuous and satisfies a Lipschitz condition on D where D is some bounded connected open set containing the intersection of $\{(t, y^{(1)}, \dots, y^{(d)}) : t \geq 0\}$ with some neighborhood of $\{(0, y^{(1)}, \dots, y^{(d)}) : \mathbb{P}[\cap_{i \in [d]} \{Y_0^{(i)} = y^{(i)}n\}] \neq 0 \text{ for some } n\}$.

Theorem 2.3.2 (Wormald (1995)). *Assume that conditions (C1)-(C3) hold. Then,*

1. *For $(0, \hat{x}^{(1)}, \dots, \hat{x}^{(d)}) \in D$, the system of differential equations*

$$\frac{d}{dt} x_t^v = f_i(t, x_1(t), \dots, x_d(t)), \text{ for } i \in [d]$$

has a unique solution in D for all $x_i : \mathbb{R} \rightarrow \mathbb{R}$ such that

$$x_i(0) = \hat{x}^{(i)} \text{ for } i \in [d]$$

and which extends to points arbitrarily close to the boundary of D .

2. *The following holds almost surely*

$$Y_t^{(i)} = nx_i(t/n) + o(n)$$

uniformly for $0 \leq t \leq \min\{\sigma(n)n, m(n)\}$, for all $i \in [d]$, where x_t^v is the solution of the differential equation with $\hat{x}^{(i)}(0) = Y_0^{(i)}/n$ and $\sigma(n)$ is the supremum of t for which the solution of the differential equation can be extended.

Chapter 3

Cooperative Non-stochastic Multi-Armed Bandits

In this chapter, we consider the multi-agent version of a multi-armed bandit problem which is one of the most fundamental decision making problems under uncertainty. In this problem, a learning agent needs to consider the exploration-exploitation trade-off, i.e. balancing the exploration of various actions in order to learn how much rewarding they are and selecting high-rewarding actions. In the multi-agent version of this problem, multiple agents collaborate with each other trying to maximize their individual cumulative rewards, and the challenge is to design efficient cooperative algorithms under communication constraints, e.g. the topology of the network and communication delays.

3.1 Introduction

We consider the non-stochastic (adversarial) multi-armed bandit problem in a cooperative multi-agent setting, with $K \geq 2$ arms and $N \geq 1$ agents. In each time step, each agent selects an arm and then observes the incurred loss corresponding to its selected arm. The losses of arms are according to an arbitrary loss sequence, which is commonly referred to as the non-stochastic or adversarial setting. Each agent observes only the loss of the arm this agent selected in each time step. The agents are allowed to cooperate by exchanging messages, which is constrained by a communication graph G such that any two agents can exchange a message directly between themselves only if they are neighbors in graph G . Each exchange of a message over an edge has delay of d time steps. The goal of each agent is to minimize its cumulative loss over a time horizon of T time steps. We study the objective of minimizing the individual regret of agents, i.e. the difference between the expected cumulative loss incurred by an agent and the cumulative loss of the best arm in hindsight. We also study the average regret of all agents.

The multi-agent multi-armed bandit problem formulation that we study captures many systems that use a network of learning agents. For example, in peer-to-peer recommender systems, the agents are users and the arms are products that can be recommended to users (Baraglia et al., 2013). The delay corresponds to the time it takes for a message to be transmitted between users. Note that in

this application scenario, the number of products (i.e. arms) may be much larger than the number of users (i.e. agents).

The collaborative multi-agent multi-armed bandit problem was studied, e.g., in Cesa-Bianchi et al. (2019) and Bar-On and Mansour (2019), where each edge has unit delay. Our setting is more general in allowing for arbitrary delay d per edge. Cesa-Bianchi et al. (2019) showed that when each agent selects arms according to a cooperative Exp3 algorithm (Exp3-Coop), the average regret is $O(\sqrt{(\alpha(G)/N + 1/K) \log(K)KT})$ for large enough T , where $\alpha(G)$ is the independence number of graph G . Bar-On and Mansour (2019) have shown that individual regret of each agent v is $O(\sqrt{(1/|\mathcal{N}(v)| + 1/K) \log(K)KT})$ when $T \geq K^2 \log(K)$, where $\mathcal{N}(v)$ is the set of neighbors of agent v and itself in graph G . This regret bound is shown to hold for an algorithm where some agents, referred to as center agents, select arms using the Exp3-Coop policy and other agents copy the actions of center agents. These bounds reveal the effect of collaboration on the learning and what graph properties effect the efficiency of learning. However, some fundamental questions still remain. For example, to the best of our knowledge, it is unknown from the previous literature what is the lower bound for this problem. Moreover, it is unknown whether better algorithms can be designed whose regret matches a lower bound under certain conditions.

In this work, we give a regret lower bound for any learning algorithm in which each agent can only communicate with their neighbors. We present a center-based algorithm whose regret upper bound matches the lower bound when the number of arms is large enough. We present an algorithm that has a regret upper bound with \sqrt{d} dependence on the delay per edge, which is optimal. All our regret bounds are parametrized with the delay parameter d , which is unlike to Cesa-Bianchi et al. (2019) and Bar-On and Mansour (2019) which considered only the special when $d = 1$.

3.1.1 Related work

The multi-armed bandit problem in a multi-agent setting, where agents collaborate with each other subject to some communication constraints, has received considerable attention in recent years. Awerbuch and Kleinberg (2008) introduced the cooperative non-stochastic multi-armed bandit problem setting where communication is through a public channel (corresponding to a complete graph) and some agents may be dishonest. Kar et al. (2011) considered a special collaboration network in which only one agent can observe the loss of the selected arm

in each time step. Szorenyi et al. (2013) discussed two specific P2P networks in which at each time step, each agent can send messages to only two other agents. Cesa-Bianchi et al. (2020) studied an online learning problem where only a subset of agents play in each time step. They showed that an optimal average regret bound for this problem is $\Theta(\sqrt{\alpha(G)T})$ when the set of agents that play in each time step is chosen randomly, while $\Omega(T)$ bound holds when the set of agents can be chosen arbitrarily in each time step. Kolla et al. (2018), Landgren et al. (2016) and Martínez-Rubio et al. (2019) considered a setting in which communication is constrained by a communication graph such that any two agents can communicate *instantly* if there is an edge connecting them.

The communication model considered in this chapter was introduced by Cesa-Bianchi et al. (2019). Here, agents communicate via messages sent over edges of a fixed connected graph and sending a message over an edge incurs a delay of value d . Cesa-Bianchi et al. (2019) considered the case when $d = 1$ whereas in this chapter, we consider $d \geq 1$.

They proposed an algorithm, referred to as Exp3-Coop, in which each agent constructs loss estimators for each arm using an importance-weighted estimator. The Exp3-Coop algorithm has an upper bound of $O(\sqrt{(\alpha(G)/N + 1/K) \log(K)KT} + \log(T))$ on the average regret.

Bar-On and Mansour (2019) combines the idea of center-based communication from Kolla et al. (2018) with the Exp3-Coop algorithm, showing that the center-based Exp3 algorithm has a regret upper bound of $O(\sqrt{(1/|\mathcal{N}(v)| + 1/K) \log(K)KT})$ for each individual agent when $d = 1$. We show that a better regret bound can be guaranteed with respect to the scaling with the number of arms K .

Multi-armed bandits with delayed feedback have been studied extensively in the single-agent setting (Joulani et al., 2013; Thune et al., 2019; Flaspohler et al., 2021). Specifically, Zimmert and Seldin (2020) considered a setting in which the agent has no prior knowledge about the delays and showed an optimal regret of $O(\sqrt{KT} + \sqrt{d \log(K)T})$ where d is the average delay over T time steps. We present, in the multi-agent setting, a distributed learning algorithm whose regret upper bound can also achieve this optimal \sqrt{d} dependence on the delay per edge d .

We use the *Tsallis entropy* family of regularizers proposed by Tsallis (1988). Zimmert and Seldin (2021) have shown that an online mirror descent algorithm with a Tsallis entropy regularizer achieves optimal regret for the single-agent ban-

dit problem. We show a distributed learning algorithm for the multi-agent bandit setting, which uses a Tsallis entropy regularizer.

3.1.2 Organization of this chapter

Section 3.2 provides problem formulation. Section 3.3 presents our two algorithms and their regret bounds. In Section 3.4, we present a lower bound on individual regret of each agent. Section 3.5 contains numerical results. Finally, proofs of our results are available in the Section 3.6.

3.2 Problem formulation

We consider a multi-armed bandit problem with a finite set $\mathcal{A} = \{1, \dots, K\}$ of actions (arms) played by N agents. The agents can communicate through a communication network $G = (\mathcal{V}, \mathcal{E})$ where \mathcal{V} is the set of N agents and \mathcal{E} is the set of edges such that $(u, v) \in \mathcal{E}$ if, and only if, agent u can send/receive messages to/from agent v . We denote the neighbors of the agent v and itself by the set $\mathcal{N}(v) = \{u \in \mathcal{V} : (u, v) \in \mathcal{E}\} \cup \{v\}$. Sending a message over edge $e \in \mathcal{E}$ incurs a delay of value $d_e \geq 0$ time steps. We consider the *homogeneous* setting under which $d_e = d \geq 1$ for every edge $e \in \mathcal{E}$. Note that the delayed communication network model in Cesa-Bianchi et al. (2019) and Bar-On and Mansour (2019) is restricted to the special case $d = 1$.

At each time step $t = 1, 2, \dots, T$, each agent $v \in \mathcal{V}$ chooses an action $I_t(v) \in \mathcal{A}$ according to distribution p_t^v over \mathcal{A} and then observes the loss value, $\ell_t(I_t(v)) \in [0, 1]$. Notice that the loss does not depend on the agent, but only on the time step and the chosen action. Hence, if two agents choose the same action at the same time step, they incur the same loss. We consider the *non-stochastic setting* where the losses are determined by an oblivious adversary, meaning that the losses do not depend on the agent's realized actions.

At the end of each time step t , each agent $v \in \mathcal{V}$ sends a message $S_t(v)$ of size $b_t(v)$ information bits to all its neighbors and after this, each agent $v \in \mathcal{V}$ has messages $\cup_{u \in \mathcal{N}(v)} \{S_s(u) : s + d = t\}$.

We assume that at each time step t , each agent v can send to each of its neighbors a message $S_t(v) = \langle v, t, I_t(v), \ell_t(I_t(v)), p_t^v \rangle$, i.e. the agent id, the time step, the chosen arm id, the instant loss received and the instant action distribution. We denote with $b_t(v)$ the number of information bits to encode $S_t(v)$. The total communication cost in each time step is $\sum_{v \in \mathcal{V}} \sum_{u: (u, v) \in \mathcal{E}} b_t(u)$ information bits.

The *individual regret* of each agent v is defined as the difference between its expected accumulated loss and the loss of the best action in hindsight, i.e.

$$R_T^v = \mathbb{E} \left[\sum_{t=1}^T \ell_t(I_t(v)) \right] - \min_{i \in \mathcal{A}} \sum_{t=1}^T \ell_t(i).$$

The *average regret* of N agents is defined as

$$R_T = \frac{1}{N} \sum_{v \in \mathcal{V}} R_T^v.$$

Let $\alpha(G)$ be the size of a maximal independent set of graph G , where the maximal independent set is the largest subset of nodes such that no two nodes in this set are connected by an edge.

3.3 Algorithms and regret upper bounds

In this section, we propose two collaborative multi-agent bandit algorithms, the center-based cooperative follow-the-regularized-leader (CFTRL) algorithm and the decentralized cooperative follow-the-regularized-leader (DFTRL) algorithm. The first algorithm has optimal regret up to a constant factor when the number of arms is large enough. The second algorithm has optimal dependence on the delay parameter d .

3.3.1 A center-based cooperative algorithm

We consider an algorithm where some agents, referred to as *centers*, run a FTRL algorithm, and each other agent copies the action selection distribution from its nearest center. The strategy based on using center agents was proposed in Bar-On and Mansour (2019), where agents played the Exp3 strategy instead. These center agents collaboratively update their strategies by exchanging messages with other agents, and each non-center agent copies the strategy of its nearest center agent. The center agents are selected such that they have a sufficiently large degree, which can be shown to reduce individual regret of center agents. Moreover, the center agents are selected such that each non-center agent is within a small distance to a center agent.

Let $\mathcal{C} \subseteq \mathcal{V}$ be the set of centers. The set of agents \mathcal{V} is partitioned into disjoint components \mathcal{V}_c , $c \in \mathcal{C}$. Each non-center agent v belongs to a unique component.

For each agent v , let $C(v)$ denote its center agent, $c = C(v)$ if and only if $v \in \mathcal{V}_c$. Let $d(v)$ be the distance between a non-center agent v and its center $C(v)$. The set of centers \mathcal{C} and the partitioning $\{\mathcal{V}_c : c \in \mathcal{C}\}$ are computed according to Algorithms 3 and 4 in Bar-On and Mansour (2019).

Let $\mathcal{J}_t(v) = \{I_t(v') : v' \in \mathcal{N}(v)\}$ be the set of actions chosen by agent v or its neighbors at time step t . Each center agent $c \in \mathcal{C}$ runs a FTRL algorithm with the collaborative importance-weighted loss estimators observable up to time step t , for each action i ,

$$\hat{L}_t^{c,obs}(i) = \sum_{s=1}^{t-1} \hat{\ell}_s^{c,obs}(i)$$

and

$$\hat{\ell}_t^{c,obs}(i) = \begin{cases} \frac{\ell_{t-d}(i)}{q_{t-d}^c(i)} \mathbb{I}\{i \in \mathcal{J}_{t-d}(c)\} & \text{if } t > d \\ 0 & \text{otherwise} \end{cases}$$

where

$$q_t^c(i) = 1 - \prod_{v \in \mathcal{N}(c)} (1 - p_t^v(i))$$

is the *neighborhood-aggregated importance weight*.

In each time step, the center agents update their action selection distributions according to the FTRL algorithm, i.e.

$$p_t^c = \operatorname{argmin}_{p \in \mathcal{P}_{K-1}} \left\{ \left\langle p, \hat{L}_t^{c,obs} \right\rangle + F_t(p) \right\}.$$

where $F_t(p)$ is the *Tsallis entropy regularizer* (Zimmert and Seldin, 2021) with the learning rate $\eta(c)$,

$$F_t(p) = -2 \sum_{i=1}^K \sqrt{p_i} / \eta(c). \quad (3.3.1)$$

Each non-center agent $v \in \mathcal{V} \setminus \mathcal{C}$ selects actions according to the uniform distribution until time step $t > d(v)d$. Then the non-center agent copies the action selection distribution from its center, i.e. $p_t^v = p_{t-d(v)d}^{C(v)}$.

The details of the CFTRL algorithm are described in Algorithm 3.

Individual regret upper bound We show an individual regret upper bound for Algorithm 3 in the following theorem.

Theorem 3.3.1. *Assume that $K \geq \max_{v \in \mathcal{V}} |\mathcal{N}(v)|$ and $T \geq 36(d+1)^2 K \max_c |\mathcal{N}(c)|$, and agents follow the CFTRL algorithm with each center agent $c \in \mathcal{C}$ using the learning rate $\eta(c) = \sqrt{|\mathcal{N}(c)| / (3T)}$. Then, the individual regret of each agent*

Algorithm 3: Center-based cooperative FTRL (CFTRL)

Input : Tsallis regularizer Eq. (3.3.1), learning rate $\eta(c)$ and the delay d .

Initialization: $\hat{L}_1^{c,obs}(i) = 0$ for all $i \in \mathcal{A}$ and $c \in \mathcal{C}$, $p_1^v(i) = 1/K$ for all $i \in \mathcal{A}$ and $v \notin \mathcal{C}$.

- 1 **for** each time step $t = 1, 2, \dots, T$ **do**
- 2 Each $c \in \mathcal{C}$ updates $p_t^c = \operatorname{argmin}_{p \in \mathcal{P}_{K-1}} \{\langle p, \hat{L}_t^{c,obs} \rangle + F_t(p)\}$;
- 3 Each $c \in \mathcal{C}$ chooses $I_t(c) = i$ with probability $p_t^c(i)$ and receives the loss $\ell_t(I_t(c))$;
- 4 Each $c \in \mathcal{C}$ sends the message $S_t(c) = \langle c, t, I_t(c), \ell_t(I_t(c)), p_t^c \rangle$ to all their neighbors;
- 5 Each $c \in \mathcal{C}$ receives messages $\{S_{t-d}(v) : v \in \mathcal{N}(c)\}$ and computes $\hat{L}_{t+1}^{c,obs}$;
- 6 Each $v \in \mathcal{V} \setminus \mathcal{C}$ updates $p_t^v = p_{t-d(v)d}^{C(v)}$ when $t > d(v)d$ and $p_t^v = p_1^v$ otherwise;
- 7 Each $v \in \mathcal{V} \setminus \mathcal{C}$ chooses $I_t(v) = i$ with probability $p_t^v(i)$ and receives $\ell_t(I_t(v))$;
- 8 Each $v \in \mathcal{V} \setminus \mathcal{C}$ sends $S_t(v) = \langle v, t, I_t(v), \ell_t(I_t(v)), p_t^v \rangle$ to all its neighbors.
- 9 **end**

$v \in \mathcal{V}$ is bounded as

$$R_T^v = O\left(\frac{1}{\sqrt{|\mathcal{N}(v)|}}\sqrt{KT} + d \log(K)\sqrt{|\mathcal{N}(C(v))|T}\right).$$

The proof of the theorem is provided in Section 3.6.4.

In the following, we provide a proof sketch. The proof relies on two key lemmas which are shown next.

Lemma 3.3.2. *Assume that the delay of each edge is $d \geq 1$, then the individual regret of each center agent v with the regularizer $F_t(p) = \sum_{i=1}^K f_t(p_i)$ satisfies*

$$\begin{aligned} R_T^v \leq & M + \frac{1}{2} \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} \frac{1}{q_t^v(i) f_t''(p_t^v(i))} \right] \\ & + d \cdot \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} \frac{1}{f_t''(p_t^v(i))} \right] \end{aligned}$$

where

$$M = \max_{x \in \mathcal{P}_{K-1}} -F_1(x) + \sum_{t=2}^T \max_{x \in \mathcal{P}_{K-1}} (F_{t-1}(x) - F_t(x)).$$

Lemma 3.3.3. *For any $\delta > 1$, assume that agent v runs a FTRL algorithm with Tsallis entropy (3.3.1) and learning rate $\eta(v) \leq (1 - 1/\sqrt{\delta})/(\delta^{3d/2}\sqrt{K})$ and $K \geq 2$, then for all $t \geq 1$ and $i \in \mathcal{A}$*

$$(1 - (1 + \delta)\eta(v)\hat{\ell}_t^{v,obs}(i))p_t^v(i) \leq p_{t+1}^v(i) \leq \delta p_t^v(i).$$

The proofs of the two lemmas are provided in Section 3.6.2 and 3.6.3. Similar property as in Lemma 3.3.3 was known to hold for the Exp3 algorithm by a result in Cesa-Bianchi et al. (2019). To the best of our knowledge, this property was previously not known to hold for the FTRL algorithm with Tsallis entropy.

Lemma 3.3.2 bounds the individual regret by the sum of a constant, the regret due to the *instant* bandit feedback, and the regret due to the *delayed* full-information feedback, which is $O(\sqrt{KT/|\mathcal{N}(c)|})$. Since the action selection distributions of non-center agents are copied from their centers in the past rounds, Lemma 3.3.3 bounds the difference between action selection distributions of non-center agent v and its center $C(v)$ in the same rounds when $T > d(v)d$. Consequently, the difference between the individual regret of a non-center agent v and its center $C(v)$ is bounded by $O(d(v)d\eta(C(v))T) = O(d \log(K) \sqrt{|\mathcal{N}(C(v))|T})$.

3.3.2 A decentralized cooperative algorithm

Theorem 3.3.1 provides a bound for individual regrets, which increases linearly in the edge-delay parameter d . This can be problematic when the delay in the communication network is large. We show that the effect of delays on regret can be reduced by using a decentralized follow-the-regularized-leader (DFTRL) algorithm.

In the DFTRL algorithm, each agent runs a FTRL algorithm with a *hybrid regularizer* $F_t(p)$ defined in Zimmert and Seldin (2020) as follows

$$F_t(p) = \sum_{i=1}^K \left(\frac{-2\sqrt{p_i}}{\eta_t} + \frac{p_i \log(p_i)}{\zeta_t} \right) \quad (3.3.2)$$

where η_t and ζ_t are some non-increasing sequences.

As is shown in Theorem 3.4.1, there is a regret lower bound that consists of two parts: the first part is the regret lower bound of the multi-armed bandit prob-

Algorithm 4: Decentralized cooperative FTRL (DFTRL)

Input : Hybrid regularizer Eq. (3.3.2), learning rates η_t , ζ_t , and delay d .

Initialization: $\hat{L}_1^{v,obs}(i) = 0$ for all $i \in \mathcal{A}$ and $v \in \mathcal{V}$.

- 1 **for** each time step $t = 1, 2, \dots, T$ **do**
- 2 Each $v \in \mathcal{V}$ updates $p_t^v = \operatorname{argmin}_{p \in \mathcal{P}_{K-1}} \{\langle p, \hat{L}_t^{v,obs} \rangle + F_t(p)\}$;
- 3 Each $v \in \mathcal{V}$ chooses $I_t(v) = i$ with probability $p_t^v(i)$ and receives the loss $\ell_t(I_t(v))$;
- 4 Each $v \in \mathcal{V}$ sends the message $S_t(v) = \langle v, t, I_t(v), \ell_t(I_t(v)), p_t^v \rangle$ to all their neighbors;
- 5 Each $v \in \mathcal{V}$ receives messages $\{S_{t-d}(v') : v' \in \mathcal{N}(v)\}$ and computes $\hat{L}_{t+1}^{v,obs}$;
- 6 **end**

lem and the second part is the regret lower bound of the bandit problem with full-information but delayed feedback (Joulani et al., 2013). The hybrid regularizer combines the Tsallis entropy regularizer with an optimal regularizer in the full-information setting, the negative entropy regularizer. The learning rates of the two regularizers can be tuned separately to minimize the regret from the two parts.

The details of the DFTRL are described in Algorithm 4.

Average regret upper bound We show a bound on the average regret for the DFTRL algorithm in the following theorem. The sequences η_t and ζ_t are assumed to be set as

$$\eta_t = (1/(1 - 1/e))(\alpha(G)/N + 1/K)^{-1/4} \sqrt{2/T}$$

and

$$\zeta_t = \sqrt{\log(K)/(dt)}.$$

Theorem 3.3.4. *Assume that each agent follows the DFTRL algorithm and the delay of each edge is $d \geq 1$, then the average regret over N agents is bounded as*

$$R_T = O\left(\left(\frac{\alpha(G)}{N} + \frac{1}{K}\right)^{1/4} \sqrt{KT} + \sqrt{d \log(K)T}\right).$$

Proof of the theorem is provided in Section 3.6.6. We note that the average regret scales as \sqrt{d} which is better than linear scaling of the CFTRL algorithm.

For the special case when $d = 1$, as in Cesa-Bianchi et al. (2019), Theorem 3.3.4 shows that when the number of arms K is large enough, then the DFTRL algorithm has an $O((\alpha(G)/N)^{1/4}\sqrt{KT})$ regret, which is better than $O((\alpha(G)/N)^{1/2}\sqrt{KT}\sqrt{\log(K)})$ of Exp3-Coop from Cesa-Bianchi et al. (2019). Specifically, this improvement holds when $K = \Omega(\exp(\sqrt{N/\alpha(G)}))$.

In what follows we provide a proof sketch of Theorem 3.3.4. First we present a key lemma whose proof is provided in Section 3.6.5.

Lemma 3.3.5. *For every agent $v \in \mathcal{A}$ and any probability distribution p^v over \mathcal{A} , it holds*

$$\sum_{i \in \mathcal{A}} \sum_{v \in \mathcal{V}} \frac{p^v(i)^{\frac{3}{2}}}{q^v(i)} \leq N \sqrt{\frac{1}{1-1/e} \left(\frac{\alpha(G)}{N} + \frac{1}{K} \right) K}$$

where $q^v(i) = 1 - \prod_{v' \in \mathcal{N}(v)} (1 - p^{v'}(i))$.

Lemma 3.3.5 shows that the average regret from the FTRL algorithm with the hybrid regularizer with instant feedback is $O((\alpha(G)/N + 1/K)^{1/4}\sqrt{KT})$. For the regularizer in Eq. (3.3.2), the delay effect term is $O(\sqrt{dT \log(K)})$. Lemma 3.3.2 bounds the average regret by the sum of two terms.

3.4 Regret lower bounds

We present lower bounds on individual regret R_T^v for every agent $v \in \mathcal{V}$ and average regret R_T .

Theorem 3.4.1. *The worst-case individual regret of each agent $v \in \mathcal{V}$, R_T^v , is*

$$\Omega \left(\max \left\{ \min \left\{ T, \frac{1}{\sqrt{|\mathcal{N}(v)|}} \sqrt{KT} \right\}, \sqrt{d \log(K) T} \right\} \right)$$

and the worst-case average regret, R_T , is

$$\Omega \left(\max \left\{ \min \left\{ T, c_G \sqrt{KT} \right\}, \sqrt{d \log(K) T} \right\} \right)$$

where $c_G = (1/N) \sum_{v \in \mathcal{V}} 1/\sqrt{|\mathcal{N}(v)|}$.

The proof is provided in Section 3.6.7. The lower bounds contain two parts. The first part is derived from the lower bounds in Shamir (2014) for a class of online algorithms. The second part handles the effect of delays by showing that

the individual regret of each agent cannot be smaller than the regret of a single agent with delayed full information.

We note that the individual regret of Algorithm 3 is optimal with respect to scaling with the number of arms K and the average regret of Algorithm 4 is optimal with respect to scaling with delay d .

3.5 Numerical experiments

In this section, we present results of numerical experiments whose goal is to compare performance of CFTRL and DFTRL algorithms with some state-of-the-art algorithms and demonstrate the tightness of our theoretical bounds. We consider the classic stochastic multi-armed bandit problem with agents communicating via different networks.

The stochastic multi-armed bandit problem is defined as follows: each arm i is associated with a Bernoulli distribution with mean μ_i for $i = 1, 2, \dots, K$. The loss $\ell_t(i)$ from choosing arm i at time step t is sampled independently from the corresponding Bernoulli distribution. In our experiments, we set $\mu_i = (1 + 8(i - 1)/(K - 1))/10$ so that μ_1, \dots, μ_K is a linearly decreasing sequence. Each problem instance is specified by a tuple (K, G, d) . The two baseline algorithms we choose are the center-based Exp3 algorithm in Bar-On and Mansour (2019) and the Exp3-Coop algorithm in Cesa-Bianchi et al. (2019) whose regret upper bounds are suboptimal as discussed in the introduction.

The numerical results are for four experiments whose goals are as follows:

- the first experiment compares the performance of CFTRL, DFTRL and the baselines when the number of arms increases,
- the second experiment validates the effect of the graph degree in the regret upper bound on CFTRL,
- the third experiment validates the effect of the delay parameter on the regret upper bounds of CFTRL and DFTRL, and, finally,
- the fourth experiment compares CFTRL and the center-based Exp3 algorithm on some sparse random graphs.

In summary, our numerical results validate theoretical results and demonstrate that CFTRL and DFTRL can achieve significant performance gains over some previously proposed algorithms. The code for producing our experimental results is available online in the GitHub repository: [link].

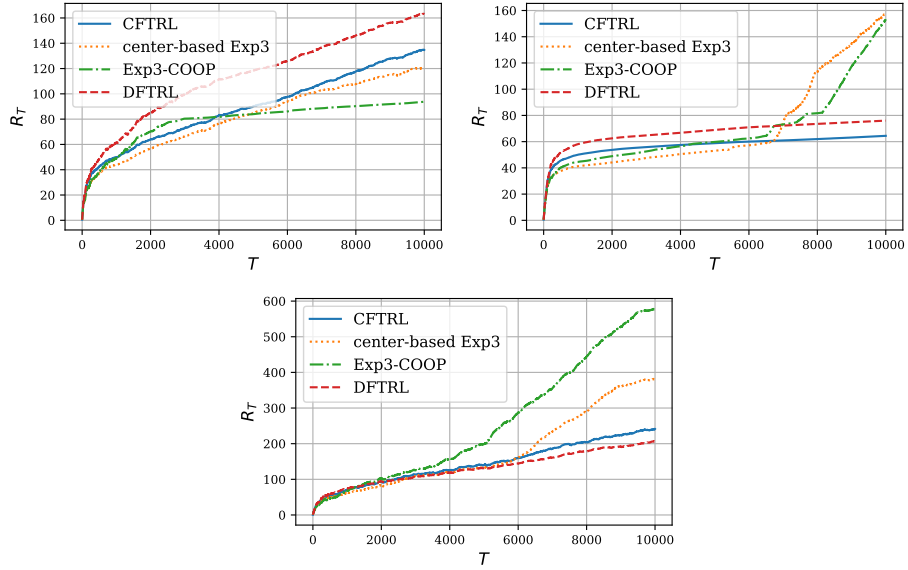


Figure 1: Average regret R_T versus T for different algorithms on a 2-regular graph with $N = 3$ agents and edge-delay $d = 1$, and varied number of arms: (top) $K = 20$ (middle) $K = 30$, and (top) $K = 40$. We used 10 independent simulation runs.

3.5.1 The effect of the number of arms

In the first experiment, we evaluate the performance of CFTRL and DFTRL against the baselines on a r -regular graph (all nodes have the same degree of r). Note that in a regular graph, each agent has equal probability to be a center agent. For a r -regular graph, CFTRL has an individual regret upper bound of $O(\sqrt{(1/r)}\sqrt{KT})$ and DFTRL has an average regret upper bound of $O(\sqrt[4]{1 - r/N}\sqrt{KT})$ when the number of arms K is large enough according to our analysis.

We first demonstrate numerical results showing that CFTRL and DFTRL can achieve significant performance gains over the center-based Exp3 algorithm whose individual regret upper bound is $O(\sqrt{(1/r)\log(K)}\sqrt{KT})$ and the Exp3-Coop algorithm whose average regret upper bound is $O(\sqrt{(1 - r/N)\log(K)}\sqrt{KT})$ when the number of arms K is large enough.

Figure 1 shows the regret R_T versus T for different number of arms, namely 20, 30 and 40. The results demonstrate that CFTRL and DFTRL achieve better regret than Exp3-COOP and center-based Exp3 when the number of arms is large enough.




Figure 2: Average regret of CFTRL versus graph degree r , for a r -regular graph with $N = 6$ nodes, $K = 10$ and $d = 1$.




Figure 3: Average regret of CFTRL and DFTRL versus the edge-delay d on a star regular graph with $N = 20$ and $K = 3$.

3.5.2 The effect of graph degree on CFTRL

In the second experiment, we validate the scaling of the graph degree in the regret upper bound of CFTRL. On the r -regular graph, CFTRL has a regret that scales as $O(1/\sqrt{r})$ according to Theorem 3.3.1. We run CFTRL on the problem instances with fixed number of arms K , delay d and increasing node degree r . The results in Figure 2 shows that the averaged regret decreases as the graph degree increases and the rate of decrease is approximately $O(1/\sqrt{r})$.

3.5.3 The effect of delay on CFTRL and DFTRL

In the third experiment, we run CFTRL and DFTRL algorithms on a fixed star graph G with a fixed number of arms K and varied edge-delay d . Figure 3 shows that the normalized regret of CFTRL is $R_T/\sqrt{d} = O(\sqrt{d})$ while the normalized regret of DFTRL is $R_T/\sqrt{d} = O(1)$. Hence, when the delay d is large enough, CFTRL has a linearly increasing regret with respect to d which is in contrast to the sub-linear increasing regret of DFTRL. This is consistent with our theoretical analysis, which states that CFTRL has a regret upper bound of $O(d)$ and DFTRL has a regret upper bound of $O(\sqrt{d})$.

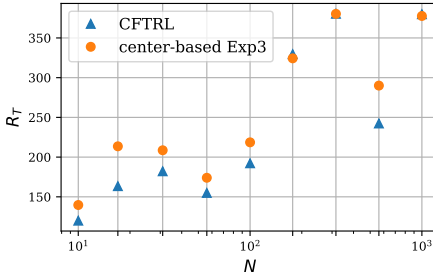


Figure 4: Average regret R_T versus the number of nodes N for sparse Erdős–Rényi random graphs, for CFTRL and center-based Exp3 algorithms.

3.5.4 The effect of graph sparsity

In the fourth experiment, we validate that our CFTRL algorithm can outperform the center-based Exp3 algorithm on some random graphs. We consider Erdős–Rényi random graphs of N nodes with probability of an edge equal to $2 \log(N)/N$. This condition ensures that the graph is connected and $|\mathcal{N}(v)| = O(\log(N))$ for all $v \in \mathcal{V}$, almost surely (Blum et al., 2020, Corollary 8.2). This random graph allows us to evaluate performance of algorithms for a large sparse random graph. We fix K and d and vary the number of nodes N and compare the performance of CFTRL and the center-based Exp3 algorithm on these graphs. We fix T to 1000 time steps. According to our analysis, CFTRL has a lower individual regret bound than the center-based Exp3 algorithm when the number of arms is large enough relative to the number of agents. The results in Figure 4 indicate that CFTRL has at least as good performance as the center-based Exp3 algorithm when K varies, and can have significantly better performance when the number of arms is large relative to the number of agents.

3.6 Proofs

3.6.1 An auxiliary lemma

We show and prove the next lemma which is used for proving some other lemmas.

Lemma 3.6.1. *For all $1 \leq s \leq t$, $v \in \mathcal{V}$ and $i \in \mathcal{A}$, it holds*

1. $\mathbb{E}_s[\hat{\ell}_t^v(i)] = \ell_t(i)$
2. $\mathbb{E}_t[\hat{\ell}_t^v(i)^2] = \ell_t(i)^2 / q_t^v(i)$, and

$$3. \mathbb{E}_t[\hat{\ell}_t^v(i)\hat{L}_t^{v,miss}(i)] = \ell_t(i)\hat{L}_t^{v,miss}(i).$$

Proof. Consider an arbitrary $v \in \mathcal{V}$. By definition, p_t^v is determined by $\hat{L}_t^{v,obs}$ which depends on the realization of $I_1(v), I_2(v), \dots, I_{t-d-1}(v)$. Note that the following equations hold, for all $i \in \mathcal{A}$,

$$\begin{aligned} \mathbb{E}_t[\hat{\ell}_t^v(i)] &= \frac{\ell_t(i)}{q_t^v(i)} \mathbb{E}_t[\mathbb{I}\{i \in \mathcal{J}_t(v)\}] = \frac{\ell_t(i)}{q_t^v(i)} q_t^v(i) = \ell_t(i) \\ \mathbb{E}_t[\hat{\ell}_t^v(i)^2] &= \frac{\ell_t(i)^2}{q_t^v(i)^2} \mathbb{E}_t[\mathbb{I}\{i \in \mathcal{J}_t(v)\}] = \frac{\ell_t(i)^2}{q_t^v(i)^2} q_t^v(i) = \frac{\ell_t(i)^2}{q_t^v(i)}. \end{aligned}$$

By the tower law of conditional expectation, we have $\mathbb{E}_s[\hat{\ell}_t^v(i)] = \ell_t(i)$ for all $s \leq t$. Note that $\hat{L}_t^{v,obs}(i) = \sum_{s=1}^{t-d-1} \hat{\ell}_s^{v,obs}(i)$, thus, $\hat{L}_t^{v,miss}(i) = \sum_{s=t-d}^{t-1} \hat{\ell}_s^v(i)$, which depends on the realization of $I_1(v), I_2(v), \dots, I_{t-1}(v)$ for all $v \in \mathcal{V}$. Hence, for all $i \in \mathcal{A}$, we have

$$\mathbb{E}_t[\hat{\ell}_t^v(i)\hat{L}_t^{v,miss}(i)] = \mathbb{E}_t[\hat{\ell}_t^v(i)]\hat{L}_t^{v,miss}(i) = \ell_t(i)\hat{L}_t^{v,miss}(i).$$

□

3.6.2 Proof of Lemma 3.3.2

Let $i^* = \arg \min_{i \in \mathcal{A}} \sum_{t=1}^T \ell_t(i)$ be the best action in hindsight. Consider an arbitrary agent $v \in \mathcal{V}$. Following arguments in Theorem 3 from Zimmert and Seldin (2020), we have

$$\begin{aligned} R_T^v &= \mathbb{E} \left[\sum_{t=1}^T \ell_t(I_t(v)) - \sum_{t=1}^T \ell_t(i^*) \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \langle p_t^v, \hat{\ell}_t(v) \rangle - \hat{L}_{T+1}^v(i^*) \right] \\ &= \mathbb{E} [\alpha + \beta + \gamma] \end{aligned}$$

where

$$\begin{aligned} \alpha &= \sum_{t=1}^T \left[\bar{F}_t^* \left(-\hat{L}_t^{v,obs} - \hat{\ell}_t^v \right) - \bar{F}_t^* \left(-\hat{L}_t^{v,obs} \right) + \langle p_t^v, \hat{\ell}_t^v \rangle \right] \\ \beta &= \sum_{t=1}^T \left[\bar{F}_t^* \left(-\hat{L}_t^{v,obs} \right) - \bar{F}_t^* \left(-\hat{L}_t^{v,obs} - \hat{\ell}_t^v \right) - \bar{F}_t^* \left(-\hat{L}_t^v \right) + \bar{F}_t^* \left(-\hat{L}_{t+1}^v \right) \right] \\ \gamma &= \sum_{t=1}^T \left[\bar{F}_t^* \left(-\hat{L}_t^v \right) - \bar{F}_t^* \left(-\hat{L}_{t+1}^v \right) \right] - \hat{L}_{T+1}^v(i^*). \end{aligned}$$

Next, we bound α, β and γ , as follows.

First, for α we have

$$\begin{aligned}\mathbb{E}[\alpha] &\leq \frac{1}{2} \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{J}_t(v)} \frac{\hat{\ell}_t^v(i)^2}{f_t''(p_t^v(i))} \right] \\ &\leq \frac{1}{2} \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} \frac{\hat{\ell}_t^v(i)^2}{f_t''(p_t^v(i))} \right] \\ &\leq \frac{1}{2} \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} \frac{1}{q_t^v(i) f_t''(p_t^v(i))} \right]\end{aligned}$$

where the first inequality comes from the proof in Lemma 1 in Zimmert and Seldin (2020) and the last inequality comes from Lemma 3.6.1.

Second, by Lemma 3 in Zimmert and Seldin (2020),

$$\begin{aligned}&\bar{F}_t^* \left(-\hat{L}_t^{v,obs} \right) - \bar{F}_t^* \left(-\hat{L}_t^{v,obs} - \hat{\ell}_t^v \right) - \bar{F}_t^* \left(-\hat{L}_t^v \right) + \bar{F}_t^* \left(-\hat{L}_{t+1}^v \right) \\ &\leq \sum_{i \in \mathcal{J}_t(v)} \int_0^1 \hat{\ell}_t^v(i) f_t''(p_t^v(i))^{-1} \hat{L}_t^{v,miss}(i) dx \\ &= \sum_{i \in \mathcal{J}_t(v)} \hat{\ell}_t^v(i) f_t''(p_t^v(i))^{-1} \hat{L}_t^{v,miss}(i).\end{aligned}$$

Hence, for β we have

$$\begin{aligned}\mathbb{E}[\beta] &\leq \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} \frac{\ell_t(i) \hat{L}_t^{v,miss}(i)}{f_t''(p_t^v(i))} \right] \leq \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} \frac{\sum_{s=1}^d \hat{\ell}_{t-s}^v(i)}{f_t''(p_t^v(i))} \right] \\ &\leq d \sum_{t=1}^T \mathbb{E} \left[\sum_{i \in \mathcal{A}} \frac{1}{f_t''(p_t^v(i))} \right]\end{aligned}$$

where the first and third inequalities come from Lemma 3.6.1.

Finally, by Lemma 2 in Zimmert and Seldin (2020), we have

$$\gamma \leq \max_{x \in \mathcal{P}_{K-1}} -F_1(x) + \sum_{t=2}^T \max_{x \in \mathcal{P}_{K-1}} (F_{t-1}(x) - F_t(x)).$$

3.6.3 Proof of Lemma 3.3.3

Consider an arbitrary agent $v \in \mathcal{V}$. The arm selection distribution p_t^v is the solution of the following convex optimization problem:

$$\begin{aligned} & \text{minimize} && \eta(v) \sum_{s=1}^{t-1} \langle x, \hat{\ell}_t^{v,obs} \rangle + \left(\sum_{i=1}^K -2\sqrt{x_i} \right) \\ & \text{over} && x \in \mathbf{R}^K \\ & \text{subject to} && x_i \geq 0, \text{ for all } i \in \mathcal{A} \\ & && \sum_{i=1}^K x_i = 1. \end{aligned}$$

The distribution p_t^v satisfies

$$p_t^v(i) = \frac{1}{\left(\eta(v) \sum_{s=1}^{t-1} \hat{\ell}_t^{v,obs}(i) + \lambda_t^v \right)^2}, \text{ for } i \in \mathcal{A} \quad (3.6.1)$$

where $\lambda_t^v \geq 0$ is the Lagrange multiplier associated with the constraint $\sum_{i=1}^K p_t^v(i) = 1$, which is the solution of the following fixed-point equation

$$\sum_{i=1}^K \frac{1}{\left(\eta(v) \sum_{s=1}^{t-1} \hat{\ell}_t^{v,obs}(i) + \lambda_t^v \right)^2} = 1. \quad (3.6.2)$$

From (3.6.2) it follows that λ_t^v is a decreasing sequence in t . To see this, note

$$\sum_{i=1}^K \frac{1}{\left(\eta(v) \sum_{s=1}^{t-1} \hat{\ell}_t^{v,obs}(i) + \lambda_t^v + (\eta(v) \hat{\ell}_t^{v,obs}(i) + \lambda_{t+1}^v - \lambda_t^v) \right)^2} = 1$$

If $\lambda_{t+1}^v - \lambda_t^v > 0$, then the left-hand side in the last equation is less than 1 and, hence, the equation cannot hold. Since $\hat{\ell}_t^{v,obs}(i) > 0$ if, and only if, $I_t(v) = i$, it follows that $\hat{\ell}_t^{v,obs}(I_t(v)) + \lambda_{t+1}^v - \lambda_t^v \geq 0$ and $\lambda_{t+1}^v \leq \lambda_t^v$.

From (3.6.1), for all $i \in \mathcal{A}$,

$$p_{t+1}^v(i) = \frac{p_t^v(i)}{\left(1 + \sqrt{p_t^v(i)} (\eta(v) \hat{\ell}_t^{v,obs}(i) + \lambda_{t+1}^v - \lambda_t^v) \right)^2}.$$

From this and the fact $\hat{\ell}_t^{v,obs}(I_t(v)) + \lambda_{t+1}^v - \lambda_t^v \geq 0$, it follows

$$p_{t+1}^v(I_t(v)) \leq p_t^v(I_t(v)).$$

Let $\gamma = \lambda_{t+1}^v - \lambda_t^v$. It holds

$$\sum_{i \in \mathcal{A} \setminus \{I_t(v)\}} p_t^v(i) \frac{1}{\left(1 + \sqrt{p_t^v(i)} \gamma\right)^2} + p_t^v(I_t(v)) \frac{1}{\left(1 + \sqrt{p_t^v(I_t(v))} (\eta \hat{\ell}_t^{v,obs}(I_t(v)) + \gamma)\right)^2} = 1.$$

By Jensen's inequality,

$$\frac{1}{\left(1 + \sum_{i \in \mathcal{A} \setminus \{I_t(v)\}} p_t^v(i) \sqrt{p_t^v(i)} \gamma + p_t^v(I_t(v)) \sqrt{p_t^v(I_t(v))} (\eta \hat{\ell}_t^{v,obs}(I_t(v)) + \gamma)\right)^2} \leq 1.$$

This is equivalent to

$$\sum_{i \in \mathcal{A} \setminus \{I_t(v)\}} p_t^v(i) \sqrt{p_t^v(i)} \gamma + p_t^v(I_t(v)) \sqrt{p_t^v(I_t(v))} (\eta \hat{\ell}_t^{v,obs}(I_t(v)) + \gamma) \geq 0$$

from which it follows

$$\gamma \geq -\eta(v) \frac{p_t^v(I_t(v))^{3/2}}{\sum_{i \in \mathcal{A}} p_t^v(i)^{3/2}} \hat{\ell}_t^{v,obs}(I_t(v)) \geq -\eta(v) \frac{p_t^v(I_t(v))^{3/2}}{q_{t-d}^v(I_t(v)) \sum_{i \in \mathcal{A}} p_t^v(i)^{3/2}}$$

where the second inequality comes from $\ell_{t-d}(I_t(v)) \mathbb{I}\{I_t(v) \in \mathcal{J}_{t-d}(v)\} \leq 1$.

Here we show that for any $\delta > 1$, there exists a learning rate $\eta(v) > 0$ such that

$$p_t^v(i) \leq \delta p_{t-1}^v(i) \quad \text{for all } i \in \mathcal{A}$$

for all $t = 2, \dots, T$ by induction over t .

Base step: it is trivial to see that the above inequalities hold for $t = 2, \dots, d+1$ by noticing that

$$\hat{\ell}_t^{v,obs}(i) = 0 \quad \text{for all } 1 \leq t \leq d$$

from which, combined with (3.6.1) and (3.6.2), it follows that

$$p_1^v(i) = \dots = p_d^v(i) = p_{d+1}^v(i) \quad \text{for all } i \in \mathcal{A}.$$

Hence, for $\delta > 1$,

$$p_t^v(i) = p_{t-1}^v(i) \leq \delta p_{t-1}^v(i) \quad \text{for all } i \in \mathcal{A}$$

holds for $t = 2, \dots, d+1$.

Induction step: suppose that $p_s^v(i) \leq \delta p_{s-1}^v(i)$ holds for all $i \in \mathcal{A}$ and for all

$s = 2, \dots, t$ where $1 + d \leq t \leq T - 1$, then we have

$$p_t^v(i) \leq \delta p_{t-1}^v(i) \leq \dots \leq \delta^d p_{t-d}^v(i).$$

Specifically, we have $p_t^v(I_t(v)) \leq \delta^d p_{t-d}^v(I_t(v))$.

From this, it follows

$$\frac{p_t^v(I_t(v))^{3/2}}{q_{t-d}^v(I_t(v)) \sum_{i \in \mathcal{A}} p_t^v(i)^{3/2}} \leq \delta^{3d/2} \frac{p_{t-d}^v(I_t(v))^{3/2}}{q_{t-d}^v(I_t(v)) \sum_{i \in \mathcal{A}} p_t^v(i)^{3/2}} \leq \delta^{3d/2} \frac{\sqrt{p_{t-d}^v(I_t(v))}}{\sum_{i \in \mathcal{A}} p_t^v(i)^{3/2}}$$

where the second inequality comes from $p_{t-d}^v(I_t(v)) \leq q_{t-d}^v(I_t(v))$.

This implies

$$\begin{aligned} \sqrt{p_t^v(i)} \gamma &\geq -\eta(v) \delta^{3d/2} \frac{\sqrt{p_{t-d}^v(I_t(v))} \sqrt{p_t^v(i)}}{\sum_{i \in \mathcal{A}} p_t^v(i)^{3/2}} \\ &\geq -\eta(v) \delta^{3d/2} \frac{1}{\sum_{i \in \mathcal{A}} p_t^v(i)^{3/2}} \\ &\geq -\eta(v) \delta^{3d/2} \sqrt{K} \end{aligned}$$

where the last inequality is because $\sum_{j=1}^K p_t^v(j)^{3/2} = K \sum_{j=1}^K (1/K) p_t^v(j)^{3/2} \geq K (\sum_{j=1}^K (1/K) p_t^v(j))^{3/2} = 1/\sqrt{K}$. From (3.6.1), it follows that for every $i \neq I_t(v)$,

$$p_{t+1}^v(i) = \frac{1}{(1 + \sqrt{p_t^v(i)} \gamma)^2} p_t^v(i) \leq \frac{1}{(1 - \eta(v) \delta^{3d/2} \sqrt{K})^2} p_t^v(i) \leq \delta p_t^v(i)$$

whenever $\eta(v) \leq (1 - 1/\sqrt{\delta}) / (\delta^{3d/2} \sqrt{K})$.

Clearly, since $\delta > 1$, we have

$$p_{t+1}^v(I_t(v)) \leq p_t^v(I_t(v)) \leq \delta p_t^v(I_t(v)).$$

Hence,

$$p_s^v(i) \leq \delta p_{s-1}^v(i) \quad \text{for all } i \in \mathcal{A}$$

holds for $s = t + 1$.

Now we have shown both the base step and induction step.

From the induction, for any $\delta > 1$,

$$p_{t+1}^v(i) \leq \delta p_t^v(i)$$

whenever $\eta(v) \leq (1 - 1/\sqrt{\delta}) / (\delta^{3d/2} \sqrt{K})$.

From (3.6.1) and $\gamma \leq 0$, for all $i \in \mathcal{A}$,

$$p_t^v(i) \leq \left(1 + \sqrt{p_t^v(i) \eta(v) \hat{\ell}_t^{v,obs}(i)}\right)^2 p_{t+1}^v(i).$$

Now note that

$$\left(1 + \sqrt{p_t^v(i) \eta(v) \hat{\ell}_t^{v,obs}(i)}\right)^2 \leq \left(1 + \eta(v) \hat{\ell}_t^{v,obs}(i)\right)^2 \leq \frac{1 + \eta(v) \hat{\ell}_t^{v,obs}(i)}{1 - \eta(v) \hat{\ell}_t^{v,obs}(i)}$$

from which it follows that

$$p_t^v(i) \leq \frac{1 + \eta(v) \hat{\ell}_t^{v,obs}(i)}{1 - \eta(v) \hat{\ell}_t^{v,obs}(i)} p_{t+1}^v(i).$$

This is equivalent to

$$\left(1 - \eta(v) \hat{\ell}_t^{v,obs}(i)\right) p_t^v(i) \leq \left(1 + \eta(v) \hat{\ell}_t^{v,obs}(i)\right) p_{t+1}^v(i) = p_{t+1}^v(i) + \eta(v) \hat{\ell}_t^{v,obs}(i) p_{t+1}^v(i).$$

It follows that when $\eta(v) \leq (1 - 1/\sqrt{\delta}) / (\delta^{3d/2} \sqrt{K})$,

$$\left(1 - \eta(v) \hat{\ell}_t^{v,obs}(i)\right) p_t^v(i) \leq p_{t+1}^v(i) + \eta(v) \hat{\ell}_t^{v,obs}(i) p_{t+1}^v(i) \leq p_{t+1}^v(i) + \delta \eta(v) \hat{\ell}_t^{v,obs}(i) p_t^v(i).$$

Rearranging the terms, we have

$$\left(1 - (1 + \delta) \eta(v) \hat{\ell}_t^{v,obs}(i)\right) p_t^v(i) \leq p_{t+1}^v(i).$$

Hence, we have

$$\left(1 - (1 + \delta) \eta(v) \hat{\ell}_t^{v,obs}(i)\right) p_t^v(i) \leq p_{t+1}^v(i) \leq \delta p_t^v(i)$$

whenever $\eta(v) \leq (1 - 1/\sqrt{\delta}) / (\delta^{3d/2} \sqrt{K})$. Letting $\delta = 2$ completes the proof.

3.6.4 Proof of Theorem 3.3.1

Recall from the proof of Lemma 3.3 that for any $\delta > 1$, when $\eta(v) \leq (1 - 1/\sqrt{\delta}) / (\delta^{3d/2} \sqrt{K})$, then it holds

$$\left(1 - (1 + \delta) \eta(v) \hat{\ell}_t^{v,obs}(i)\right) p_t^v(i) \leq p_{t+1}^v(i) \leq \delta p_t^v(i).$$

Let $\delta = 1 + 1/d$. Then, we have that for each center agent $c \in \mathcal{C}$,

$$p_t^c(i) \leq (1 + 1/d)p_{t-1}^c(i) \leq \cdots \leq (1 + 1/d)^d p_{t-d}^c(i) \leq 3p_{t-d}^c(i) \quad (3.6.3)$$

and

$$p_t^c(i) \leq p_{t+1}^c(i) + (1 + \delta)\eta(c)p_t^c(i)\hat{\ell}_t^{c,obs}(i) = p_{t+1}^c(i) + (2 + 1/d)\eta(c)p_t^c(i)\hat{\ell}_t^{c,obs}(i) \quad (3.6.4)$$

when

$$\eta(c) \leq \frac{1 - \sqrt{d/(d+1)}}{(1 + 1/d)^{3d/2}\sqrt{K}}.$$

From Algorithm 3 and (3.6.3), it follows that

$$\begin{aligned} q_t^c(i) &= 1 - \prod_{v \in \mathcal{N}(c)} (1 - p_t^v(i)) \\ &= 1 - \prod_{v \in \mathcal{N}(c)} (1 - p_{t-d}^c(i)) \\ &\geq 1 - \left(1 - \frac{1}{3}p_t^c(i)\right)^{|\mathcal{N}(c)|} \\ &\geq 1 - \exp\left(-\frac{1}{3}p_t^c(i)|\mathcal{N}(c)|\right) \\ &\geq (1 - e^{-1}) \min\left\{1, \frac{1}{3}p_t^c(i)|\mathcal{N}(c)|\right\} \end{aligned}$$

from which it follows

$$\frac{1}{q_t^c(i)} \leq \frac{1}{(1 - e^{-1}) \min\{1, \frac{1}{3}p_t^c(i)|\mathcal{N}(c)|\}} \leq 2 + \frac{6}{|\mathcal{N}(c)|p_t^c(i)}. \quad (3.6.5)$$

For the Tsallis entropy regularizer (1) in Algorithm 1, we have

$$f_t''(x) = \frac{1}{2\eta(c)x^{3/2}}. \quad (3.6.6)$$

From Lemma 3.3.2 and (3.6.6), we have that the individual regret of center c is

bounded as

$$\begin{aligned}
R_T^c &\leq M + \frac{1}{2} \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} \frac{1}{q_t^c(i) f_t''(p_t^c(i))} \right] + d \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} \frac{1}{f_t''(p_t^c(i))} \right] \\
&\leq 2 \frac{1}{\eta(c)} \sqrt{K} + \eta(c) \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} \frac{p_t^c(i)^{\frac{3}{2}}}{q_t^c(i)} \right] + 2d \eta(c) \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} p_t^c(i)^{\frac{3}{2}} \right] \quad (3.6.7)
\end{aligned}$$

Since $p_t^c(i) \leq 1$, we have

$$\sum_{t=1}^T \sum_{i \in \mathcal{A}} p_t^c(i)^{\frac{3}{2}} \leq \sum_{t=1}^T \sum_{i \in \mathcal{A}} p_t^c(i) = T. \quad (3.6.8)$$

Let us define $M : \mathcal{V} \rightarrow \mathbf{R}$ as follows

$$M(c) = \min\{|\mathcal{N}(c)|, K\}, \text{ for } c \in \mathcal{C} \text{ and } M(v) = e^{-\frac{1}{6}d(v)} M(\mathcal{C}(v)) \text{ for } v \in \mathcal{V} \setminus \mathcal{C}. \quad (3.6.9)$$

Now, note

$$\begin{aligned}
\mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} \frac{p_t^c(i)^{\frac{3}{2}}}{q_t^c(i)} \right] &\leq \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} p_t^c(i)^{\frac{3}{2}} \left(2 + \frac{6}{|\mathcal{N}(c)| p_t^c(i)} \right) \right] \\
&\leq 2T + \frac{6}{M(c)} \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} \sqrt{p_t^c(i)} \right] \\
&\leq 2T + \frac{6T\sqrt{K}}{M(c)} \quad (3.6.10)
\end{aligned}$$

where the first inequality is by (3.6.5), the second inequality is by (3.6.8) and (3.6.9), and the last inequality is by Cauchy-Schwartz inequality.

From (3.6.7), (3.6.8), and (3.6.10), we have

$$R_T^c \leq 2 \frac{1}{\eta(c)} \sqrt{K} + \eta(c) \frac{6T\sqrt{K}}{M(c)} + 2(d+1)\eta(c)T.$$

By letting $\eta(c) = \sqrt{M(c)/(3T)}$, it follows that

$$R_T^c \leq 4\sqrt{3} \sqrt{\frac{KT}{M(c)}} + \frac{2(d+1)}{\sqrt{3}} \sqrt{M(c)T}. \quad (3.6.11)$$

Note that when

$$T \geq 36(d+1)^2 K \max_c M(c) > 9K \max_c M(c) \left(\frac{1 + \sqrt{1 - \frac{1}{d+1}}}{\frac{1}{d+1}} \right)^2 = \frac{9K \max_c M(c)}{\left(1 - \sqrt{\frac{d}{d+1}}\right)^2},$$

it follows that

$$\eta(c) = \sqrt{\frac{M(c)}{3T}} \leq \frac{1 - \sqrt{\frac{d}{d+1}}}{3^{3/2}\sqrt{K}} \leq \frac{1 - \sqrt{\frac{d}{d+1}}}{(1 + 1/d)^{3d/2}\sqrt{K}}.$$

Hence (3.6.3) and (3.6.4) holds for every center c .

Now, note that for any non-center agent v , $p_t^v = p_{t-d(v)d}^{C(v)}$. Then, from (3.6.4), for all $t > d(v)d$,

$$\begin{aligned} p_t^v(i) &= p_{t-d(v)d}^{C(v)}(i) \\ &\leq p_{t-d(v)d+1}^{C(v)}(i) + (2 + 1/d)\eta(C(v))p_{t-d(v)d}^{C(v)}(i)\hat{\ell}_{t-d(v)d}^{C(v),obs}(i) \\ &\dots \\ &\leq p_t^{C(v)}(i) + (2 + 1/d)\eta(C(v)) \sum_{s=1}^{d(v)d} p_{t-s}^{C(v)}(i)\hat{\ell}_{t-s}^{C(v),obs}(i). \end{aligned} \quad (3.6.12)$$

For every non-center agent v , the individual regret is bounded as

$$\begin{aligned} R_T^v &\leq (d(v) + 1)d + \mathbb{E} \left[\sum_{t=(d(v)+1)d+1}^T \sum_{i \in \mathcal{A}} p_t^v(i) \ell_t(i) - \min_{i \in \mathcal{A}} \sum_{t=(d(v)+1)d+1}^T \ell_t(i) \right] \\ &\leq (d(v) + 1)d + \mathbb{E} \left[\sum_{t=(d(v)+1)d+1}^T \sum_{i \in \mathcal{A}} p_t^{C(v)}(i) \ell_t(i) - \min_{i \in \mathcal{A}} \sum_{t=(d(v)+1)d+1}^T \ell_t(i) \right] \\ &\quad + (2 + 1/d)\eta(C(v)) \mathbb{E} \left[\sum_{t=(d(v)+1)d+1}^T \sum_{i \in \mathcal{A}} \sum_{s=1}^{d(v)d} p_{t-s}^{C(v)}(i) \hat{\ell}_{t-s}^{C(v),obs}(i) \ell_t(i) \right] \\ &= (d(v) + 1)d + \mathbb{E} \left[\sum_{t=(d(v)+1)d+1}^T \sum_{i \in \mathcal{A}} p_t^{C(v)}(i) \ell_t(i) - \min_{i \in \mathcal{A}} \sum_{t=(d(v)+1)d+1}^T \ell_t(i) \right] \\ &\quad + (2 + 1/d)\eta(C(v)) \mathbb{E} \left[\sum_{t=(d(v)+1)d+1}^T \sum_{i \in \mathcal{A}} \sum_{s=1}^{d(v)d} p_{t-s}^{C(v)}(i) \hat{\ell}_{t-s-d}^{C(v)}(i) \ell_t(i) \right] \end{aligned} \quad (3.6.13)$$

where the second inequality follows from (3.6.12) and the equality uses $\hat{\ell}_{t-d}^v =$

$\hat{\ell}_t^{v,obs}$ when $t > d$.

From Lemma 3.6.1, we have

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=(d(v)+1)d+1}^T \sum_{i \in \mathcal{A}} \sum_{s=1}^{d(v)d} p_{t-s}^{C(v)}(i) \hat{\ell}_{t-s-d}^{C(v)}(i) \ell_t(i) \right] \\ & \leq \mathbb{E} \left[\sum_{t=(d(v)+1)d+1}^T \sum_{i \in \mathcal{A}} \sum_{s=1}^{d(v)d} p_{t-s}^{C(v)}(i) \mathbb{E}_{t-s-d} \left[\hat{\ell}_{t-s-d}^{C(v)}(i) \right] \right] \\ & \leq d(v)dT. \end{aligned}$$

Hence, it follows that

$$R_T^v \leq R_T^{C(v)} + (d(v) + 1)d + (2 + 1/d)\eta(C(v))d(v)dT.$$

Combining with Eq. (3.6.11), we have

$$\begin{aligned} R_T^v & \leq 4\sqrt{3} \sqrt{\frac{KT}{M(C(v))}} + \frac{2(d+1)}{\sqrt{3}} \sqrt{M(C(v))T} + (d(v) + 1)d + (2 + 1/d)\eta(C(v))d(v)dT \\ & = 4\sqrt{3} \sqrt{\frac{KT}{M(C(v))}} + \frac{2(d+1)}{\sqrt{3}} \sqrt{M(C(v))T} + (d(v) + 1)d + (2d+1)d(v) \sqrt{M(C(v))T/3} \\ & \leq 4\sqrt{3} \sqrt{\frac{KT}{M(v)}} + \frac{2(d+1)}{\sqrt{3}} \sqrt{M(C(v))T} + (d(v) + 1)d + (2d+1)d(v) \sqrt{M(C(v))T/3} \\ & \leq 4\sqrt{3} \sqrt{\frac{KT}{M(v)}} + \frac{2(d+1)}{\sqrt{3}} \sqrt{M(C(v))T} + 6d \log(K) + 2\sqrt{3}(2d+1) \log(K) \sqrt{M(C(v))T} \end{aligned}$$

where the penultimate inequality follows from Eq. (3.6.9) and the last inequality follows from Lemma 7 in Bar-On and Mansour (2019).

From Theorem 8 in Bar-On and Mansour (2019), we have for all agent $v \in \mathcal{V}$

$$M(v) \geq (1/e) \min\{|\mathcal{N}(v)|, K\}$$

Hence, we have

$$\begin{aligned}
R_T^v &\leq 4\sqrt{3}\sqrt{e}\sqrt{\frac{KT}{\min\{|\mathcal{N}(v)|, K\}}} \\
&\quad + \frac{2(d+1)}{\sqrt{3}}\sqrt{M(C(v))T + 6d\log(K) + 2\sqrt{3}(2d+1)\log(K)}\sqrt{M(C(v))T} \\
&\leq 12\sqrt{\frac{KT}{|\mathcal{N}(v)|}} + \frac{2(d+1)}{\sqrt{3}}\sqrt{|\mathcal{N}(C(v))|T + 6d\log(K) + 2\sqrt{3}(2d+1)\log(K)}\sqrt{|\mathcal{N}(C(v))|T}
\end{aligned}$$

where the second inequality comes from the fact that $K \geq \max_v |\mathcal{N}(v)|$.

3.6.5 Proof of Lemma 3.3.5

The proof follows by the following sequence of relations:

$$\begin{aligned}
\sum_{i \in \mathcal{A}} \sum_{v \in \mathcal{V}} \frac{p^v(i)^{3/2}}{q^v(i)} &= \sum_{i \in \mathcal{A}} \sum_{v \in \mathcal{V}} \left(\frac{p^v(i)}{q^v(i)} \cdot \sqrt{p^v(i)} \right) \\
&\leq \sum_{i \in \mathcal{A}} \sqrt{\sum_{v \in \mathcal{V}} \frac{p^v(i)^2}{q^v(i)^2}} \sqrt{\sum_{v \in \mathcal{V}} p^v(i)} \\
&\leq \sum_{i \in \mathcal{A}} \sqrt{\sum_{v \in \mathcal{V}} \frac{p^v(i)}{q^v(i)}} \sqrt{\sum_{v \in \mathcal{V}} p^v(i)} \\
&\leq \sum_{i \in \mathcal{A}} \sqrt{\frac{1}{1-1/e} \left(\alpha(G) + \sum_{v \in \mathcal{V}} p^v(i) \right)} \sqrt{\sum_{v \in \mathcal{V}} p^v(i)} \\
&\leq \sqrt{\sum_{i \in \mathcal{A}} \frac{1}{1-1/e} \left(\alpha(G) + \sum_{v \in \mathcal{V}} p^v(i) \right)} \sqrt{\sum_{i \in \mathcal{A}} \sum_{v \in \mathcal{V}} p^v(i)} \\
&= \sqrt{\frac{N}{1-1/e} (K\alpha(G) + N)}
\end{aligned}$$

where the first inequality comes from Cauchy-Schwarz inequality, the second inequality comes from the fact that $q^v(i) \geq p^v(i)$, the third inequality comes from Lemma 3 in Cesa-Bianchi et al. (2020) and the last inequality comes from Cauchy-Schwarz inequality again.

3.6.6 Proof of Theorem 3.3.4

Note that

$$\max_{x \in \mathcal{P}_{K-1}} \{-F_1(x)\} + \sum_{t=2}^T \max_{x \in \mathcal{P}_{K-1}} \{F_{t-1}(x) - F_t(x)\} = -F_T(e_{[K]}/K) = 2\frac{1}{\eta_T}\sqrt{K} + \frac{1}{\zeta_T}\log(K),$$

and

$$f_t''(x) \geq \max \left\{ \frac{1}{2\eta_t} \frac{1}{x^{3/2}}, \frac{1}{\zeta_t} \frac{1}{x} \right\}.$$

Hence, from Lemma 3.5, we have

$$\begin{aligned} \frac{1}{2N} \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} \sum_{v \in \mathcal{V}} \frac{1}{q_t^v(i) f_t''(p_t^v(i))} \right] &\leq \frac{1}{N} \mathbb{E} \left[\sum_{t=1}^T \eta_t \sum_{i \in \mathcal{A}} \sum_{v \in \mathcal{V}} \frac{p_t^v(i)^{3/2}}{q_t^v(i)} \right] \\ &\leq \sqrt{\frac{1}{1-1/e} \left(\frac{K}{N} \alpha(G) + 1 \right)} \sum_{t=1}^T \eta_t \end{aligned}$$

and

$$\begin{aligned} \frac{d}{N} \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} \sum_{v \in \mathcal{V}} \frac{1}{f_t''(p_t^v(i))} \right] &\leq \frac{d}{N} \mathbb{E} \left[\sum_{t=1}^T \zeta_t \sum_{i \in \mathcal{A}} \sum_{v \in \mathcal{V}} p_t^v(i) \right] \\ &= d \sum_{t=1}^T \zeta_t. \end{aligned}$$

Combining with Lemma 3.2, we have

$$\begin{aligned} R_T &= \frac{1}{N} \sum_{v \in \mathcal{V}} R_T^v \\ &\leq M + \frac{1}{2N} \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} \sum_{v \in \mathcal{V}} \frac{1}{q_t^v(i) f_t''(p_t^v(i))} \right] + \frac{d}{N} \mathbb{E} \left[\sum_{t=1}^T \sum_{i \in \mathcal{A}} \sum_{v \in \mathcal{V}} \frac{1}{f_t''(p_t^v(i))} \right] \\ &\leq 2\frac{1}{\eta_T}\sqrt{K} + \sqrt{\frac{1}{1-1/e} \left(\frac{K}{N} \alpha(G) + 1 \right)} \sum_{t=1}^T \eta_t + \frac{1}{\zeta_T} \log(K) + d \sum_{t=1}^T \zeta_t. \end{aligned}$$

Plugging $\eta_t = (1/(1-1/e))(\alpha(G)/N + 1/K)^{-1/4} \sqrt{2/T}$ and $\zeta_t = \sqrt{\log(K)/(dt)}$ into the inequality above, we complete the proof.

3.6.7 Proof of Theorem 3.4.1

In each round t , every agent $v \in \mathcal{V}$ receives $b_t(v) = O(|\mathcal{N}(v)|)$ bits since its neighboring agents can choose at most $\mathcal{N}(v)$ distinct actions. By Theorem 4 in

Shamir (2014), there exists some distribution \mathcal{D} over $[0, 1]^K$ such that loss vectors $\ell_t \sim \mathcal{D}$ for all $t = 1, 2, \dots, T$ independently and $\min_{i \in \mathcal{A}} \mathbb{E} [\ell_t(I_t(v)) - \ell_t(i)] = \Omega\left(\min\{T, \sqrt{KT/|\mathcal{N}(v)|}\}\right)$. Hence, the worst-case individual regret of agent v is

$$\begin{aligned}
\sup_{\ell_1, \dots, \ell_T} R_T^v &= \sup_{\ell_1, \dots, \ell_T} \mathbb{E} \left[\sum_{t=1}^T \ell_t(I_t(v)) - \min_{i \in \mathcal{A}} \sum_{t=1}^T \ell_t(i) \right] \\
&\geq \mathbb{E}_{\ell_t \sim \mathcal{D}} \left[\mathbb{E} \left[\sum_{t=1}^T \ell_t(I_t(v)) \right] \right] - \mathbb{E}_{\ell_t \sim \mathcal{D}} \left[\min_{i \in \mathcal{A}} \sum_{t=1}^T \ell_t(i) \right] \\
&\geq \mathbb{E}_{\ell_t \sim \mathcal{D}} \left[\mathbb{E} \left[\sum_{t=1}^T \ell_t(I_t(v)) \right] \right] - \min_{i \in \mathcal{A}} \mathbb{E}_{\ell_t \sim \mathcal{D}} \left[\sum_{t=1}^T \ell_t(i) \right] \\
&= \max_{i \in \mathcal{A}} \mathbb{E} \left[\sum_{t=1}^T \ell_t(I_t(v)) - \sum_{t=1}^T \ell_t(i) \right] \\
&\geq \min_{i \in \mathcal{A}} \mathbb{E} \left[\sum_{t=1}^T \ell_t(I_t(v)) - \sum_{t=1}^T \ell_t(i) \right] \\
&\geq \Omega\left(\min\{T, \sqrt{KT/|\mathcal{N}(v)|}\}\right).
\end{aligned}$$

Note that with delay d , the individual regret of the agent v can not be smaller than the regret of the agent v with access to full information of $\ell_1, \dots, \ell_{t-d}$ at each round t . Following the argument in the proof of Corollary 15 in Cesa-Bianchi et al. (2019), the worst-case regret of a single-agent with delayed full information is $\Omega(\sqrt{dT \log(K)})$. Hence,

$$\sup_{\ell_1, \dots, \ell_T} R_T^v = \Omega\left(\max\left\{\min\left\{T, \sqrt{\frac{KT}{|\mathcal{N}(v)|}}\right\}, \sqrt{dT \log(K)}\right\}\right).$$

We complete the proof by noting that $R_T = \frac{1}{N} \sum_{v \in \mathcal{V}} R_T^v$.

3.6.8 CFTRL and the center-based Exp3 when $d = 1$

As noted in Theorem 3.1 with $d = 1$ and Bar-On and Mansour (2019), when $K \geq \max_{v \in \mathcal{V}} |\mathcal{N}(v)|$ and $T \geq \max\{144K \max_{c \in \mathcal{C}} |\mathcal{N}(c)|, K^2 \log(K)\}$, the individual regret of CFTRL is

$$O\left(\sqrt{\left(\frac{1}{|\mathcal{N}(v)|} + \frac{|\mathcal{N}(\mathcal{C}(v))|}{K} \log(K)^2\right) KT}\right),$$

and the individual regret of the center-based Exp3 is

$$\tilde{O} \left(\sqrt{\left(1 + \frac{K}{|\mathcal{N}(v)|}\right) T} \right) = O \left(\sqrt{\left(\frac{1}{|\mathcal{N}(v)|} + \frac{1}{K}\right) \log(K) KT} \right)$$

If we want the individual regret of CFTRL lower than the center-based Exp3, it suffices to show

$$\frac{1}{|\mathcal{N}(v)|} + \frac{|\mathcal{N}(C(v))|}{K} \log(K)^2 \leq \left(\frac{1}{|\mathcal{N}(v)|} + \frac{1}{K} \right) \log(K)$$

which is equivalent to

$$\frac{|\mathcal{N}(C(v))|}{K} \log(K)^2 - \frac{\log(K)}{K} \leq \frac{\log(K) - 1}{|\mathcal{N}(v)|}$$

Note that for $K \geq 9$, we have

$$\log(K) - 1 \geq \frac{1}{2} \log(K)$$

It suffices to show that

$$\frac{|\mathcal{N}(C(v))|}{K} \log(K)^2 \leq \frac{\log(K)}{2|\mathcal{N}(v)|}$$

which is equivalent to

$$\frac{K}{\log(K)} \geq 2|\mathcal{N}(v)||\mathcal{N}(C(v))|.$$

3.6.9 CFTRL and the lower bound when $d = 1$

As noted in Theorem 4.1, for $T \geq K/|\mathcal{N}(v)|$, the individual regret lower bound is

$$\Omega \left(\sqrt{\left(\frac{1}{|\mathcal{N}(v)|} + \frac{1}{K} \log(K)\right) KT} \right).$$

From Theorem 3.1 with $d = 1$, when $K \geq \max_{v \in V} |\mathcal{N}(v)|$ and $T \geq 144K \max_{c \in \mathcal{C}} |\mathcal{N}(c)|$, the individual regret of CFTRL is

$$O \left(\sqrt{\left(\frac{1}{|\mathcal{N}(v)|} + \frac{|\mathcal{N}(C(v))|}{K} \log(K)^2\right) KT} \right).$$

Hence when $K \geq \max_{v \in V} |\mathcal{N}(v)|$ and $T \geq \max\{144K \max_{c \in \mathcal{C}} |\mathcal{N}(c)|, K/|\mathcal{N}(v)|\}$, the ratio between the individual regret of CFTRL and the individual regret lower bound is

$$O\left(\sqrt{\frac{\frac{1}{|\mathcal{N}(v)|} + \frac{|\mathcal{N}(\mathcal{C}(v))|}{K} \log(K)^2}{\frac{1}{|\mathcal{N}(v)|} + \frac{1}{K} \log(K)}}\right).$$

Note that

$$\begin{aligned} \sqrt{\frac{\frac{1}{|\mathcal{N}(v)|} + \frac{|\mathcal{N}(\mathcal{C}(v))|}{K} \log(K)^2}{\frac{1}{|\mathcal{N}(v)|} + \frac{1}{K} \log(K)}} &\leq 1 + \sqrt{\frac{|\mathcal{N}(\mathcal{C}(v))| \log(K) - 1}{\frac{K}{|\mathcal{N}(v)| \log(K)} + 1}} \\ &\leq 1 + \sqrt{\frac{|\mathcal{N}(\mathcal{C}(v))| |\mathcal{N}(v)| \log(K)^2}{K}}. \end{aligned}$$

When $|\mathcal{N}(\mathcal{C}(v))| |\mathcal{N}(v)| \log(K)^2 / K = O(1)$, the individual regret upper bound of CFTRL matches the lower bound up to a constant factor.

Chapter 4

Doubly Adversarial Federated Bandits

There is a rising trend of research on federated learning, which coordinates multiple *heterogeneous* agents to collectively train a learning algorithm, while keeping the raw data decentralized (Kairouz et al., 2021). We consider the federated learning variant of a multi-armed bandit problem which is one of the most fundamental sequential decision making problems. In standard multi-armed bandit problems, a learning agent needs to balance the trade-off between exploring various arms in order to learn how much rewarding they are and selecting high-rewarding arms. In federated bandit problems, multiple heterogeneous agents collaborate with each other to maximize their cumulative rewards. The challenge here is to design decentralized collaborative algorithms to find a *globally* best arm for all agents while keeping their raw data decentralized.

4.1 Introduction

Finding a globally best arm with raw arm or loss sequences stored in a distributed system has ubiquitous applications in many systems built with a network of learning agents. One application is in recommender systems where different recommendation app clients (i.e. agents) in a communication network collaborate with each other to find news articles (i.e. arms) that are popular among all users within a specific region, which can be helpful to solve the *cold start* problem (Li et al., 2010; Yi et al., 2021). In this setting, the system avoids the exchange of users' browsing history (i.e. arm or loss sequences) between different clients for better privacy protection. Another motivation is in international collaborative drug discovery research, where different countries (i.e. agents) cooperate with each other to find a drug (i.e. arm) that is uniformly effective for all patients across the world (Varatharajah and Berry, 2022). To protect the privacy of the patients involved in the research, the exact treatment history of specific patients (i.e. arm or loss sequences) should not be shared during the collaboration.

The federated bandit problems are focused on identifying a globally best arm (pure exploration) or maximizing the cumulative group reward (regret minimization) in face of heterogeneous feedback from different agents for the same arm,

which has gained much attention in recent years (Dubey and Pentland, 2020; Zhu et al., 2021; Huang et al., 2021; Shi et al., 2021; Réda et al., 2022). In prior work, heterogeneous feedbacks received by different agents are modeled as samples from some unknown but fixed distributions. Though this formulation of heterogeneous feedback allows elegant statistical analysis of the regret, it may not be adequate for dynamic (non-stationary) environments. For example, consider the task of finding popular news articles within a region mentioned above. The popularity of news articles on different topics can be time-varying, e.g. the news on football may become most popular during the FIFA World Cup but may be less popular afterwards.

In contrast with the prior work, we introduce a new *non-stochastic* federated multi-armed bandit problem in which the heterogeneous feedback received by different agents are chosen by an oblivious adversary. We consider a federated bandit problem with K arms and N agents. The agents can share their information via a communication network. At each time step, each agent will choose one arm, receive the feedback and exchange their information with their neighbors in the network. The problem is *doubly adversarial*, i.e. the losses are determined by an oblivious adversary which specifies the loss of each arm not only for each time step but also for each agent. As a result, the agents which choose the same arm at the same time step may observe different losses. The goal is to find the *globally* best arm in hindsight, whose cumulative loss averaged over all agents is lowest, without exchanging raw information consisting of arm identity or loss value sequences among agents. As standard in online learning problems, we focus on regret minimization over an arbitrary time horizon.

4.1.1 Related work

The doubly adversarial federated bandit problem is related to several lines of research, namely that on federated bandits, multi-agent cooperative adversarial bandits, and distributed optimization. Here we briefly discuss these related works.

Federated bandits Solving bandit problems in the federated learning setting has gained attention in recent years. Dubey and Pentland (2020) and Huang et al. (2021) considered the linear contextual bandit problem and extended the LinUCB algorithm (Li et al., 2010) to the federated learning setting. Zhu et al. (2021) and Shi et al. (2021) studied a federated multi-armed bandit problem where the losses

observed by different agents are i.i.d. samples from some common unknown distribution. Réda et al. (2022) considered the problem of identifying a globally best arm for multi-armed bandit problems in a centralized federated learning setting. All these works focus on the stochastic setting, i.e. the reward or loss of an arm is sampled from some unknown but fixed distribution. Our work considers the non-stochastic setting, i.e. losses are chosen by an oblivious adversary, which is a more appropriate assumption for non-stationary environments.

Multi-agent cooperative adversarial bandit Cesa-Bianchi et al. (2016); Bar-On and Mansour (2019); Yi and Vojnovic (2022) studied the adversarial case where agents receive the same loss for the same action chosen at the same time step, whose algorithms require the agents to exchange their raw data with neighbors. Cesa-Bianchi et al. (2020) discussed the cooperative online learning setting where the agents have access to the full-information feedback and the communication is asynchronous. In these works, the agents that choose the same action at the same time step receive the same reward or loss value and agents aggregate messages received from their neighbors. Our work relaxes this assumption by allowing agents to receive different losses even for the same action in a time step. Besides, we propose a new algorithm that uses a different aggregation of messages than in the aforementioned papers, which is based on distributed dual averaging method in Nesterov (2009); Xiao (2009); Duchi et al. (2011).

Distributed optimization Duchi et al. (2011) proposed the dual averaging algorithm for distributed convex optimization via a gossip communication mechanism. Subsequently, Hosseini et al. (2013) extended this algorithm to the online optimization setting. Scaman et al. (2019) found optimal distributed algorithms for distributed convex optimization and a lower bound which applies to strongly convex functions. The doubly adversarial federated bandit problem with full-information feedback is a special case of distributed online linear optimization problems. Our work complements these existing studies by providing a lower bound for the distributed online linear optimization problems. Moreover, our work proposes a near-optimal algorithm for the more challenging bandit feedback setting.

4.1.2 Organization of this chapter

We first formally formulate the doubly adversarial federated bandit problem and the federated bandit algorithms we study in Section 4.2. Then, in Section 4.3, we provide two regret lower bounds for any federated bandit algorithm under the full-information and bandit feedback setting, respectively. In Section 4.4, we present a federated bandit algorithm adapted from the celebrated Exp3 algorithm for the bandit-feedback setting, together with its regret upper bound. Finally, we show results of our numerical experiments in Section 4.5. All the proofs are available in Section 4.6

4.2 Problem setting

Consider a communication network defined by an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of N agents and $(u, v) \in \mathcal{E}$ if agent u and agent v can directly exchange messages. We assume that \mathcal{G} is simple, i.e. it contains no self loops nor multiple edges. The agents in the communication network collaboratively aim to solve a non-stochastic multi-armed bandit problem. In this problem, there is a fixed set \mathcal{A} of K arms and a fixed time horizon T . Each instance of the problem is parameterized by a tensor $L = (\ell_t^v(i)) \in [0, 1]^{T \times N \times K}$ where $\ell_t^v(i)$ is the loss associated with agent $v \in \mathcal{V}$ if it chooses arm $i \in \mathcal{A}$ at time step t .

At each time step t , each agent v will choose its action $a_t^v = i$, observe the feedback I_t^v and incur a loss defined as the average of losses of arm i over all agents, i.e.,

$$\bar{\ell}_t(i) = \frac{1}{N} \sum_{v \in \mathcal{V}} \ell_t^v(i). \quad (4.2.1)$$

At the end of each time step, each agent $v \in \mathcal{V}$ can communicate with their neighbors $\mathcal{N}(v) = \{u \in \mathcal{V} : (u, v) \in \mathcal{E}\}$. We assume a *non-stochastic* setting, i.e. the loss tensor L is determined by an oblivious adversary. In this setting, the adversary has the knowledge of the description of the algorithm running by the agents but the losses in L do not depend on the specific arms selected by the agents.

The performance of each agent $v \in \mathcal{V}$ is measured by its *regret*, defined as the difference of the expected cumulative loss incurred and the cumulative loss of a *globally* best fixed arm in hindsight, i.e.

$$R_T^v(\pi, L) = \mathbb{E} \left[\sum_{t=1}^T \bar{\ell}_t(a_t^v) - \min_{i \in \mathcal{A}} \left\{ \sum_{t=1}^T \bar{\ell}_t(i) \right\} \right] \quad (4.2.2)$$

where the expectation is taken over the action of all agents under algorithm π on instance L . We will abbreviate $R_T^v(\pi, L)$ as R_T^v when the algorithm π and instance L have no ambiguity in the context. We aim to characterize $\max_L R_T^v(\pi, L)$ for each agent $v \in \mathcal{V}$ under two feedback settings,

- full-information feedback: $I_t^v = \ell_t^v$, and
- bandit feedback: $I_t^v = \ell_t^v(a_t^v)$.

Let \mathcal{F}_t^v be the sequence of agent v 's actions and feedback up to time step t , i.e., $\mathcal{F}_t^v = \bigcup_{s=1}^t \{a_s^v, I_s^v\}$. For a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, we denote as $d(u, v)$ the number of edges of a shortest path connecting nodes u and v in \mathcal{V} and $d(v, v) = 0$.

We focus on the case when π is a *federated bandit algorithm* in which each agent $v \in \mathcal{V}$ can only communicate with their neighbors within a time step.

Definition 4.2.1 (federated bandit algorithm). *A federated bandit algorithm π is a multi-agent learning algorithm such that for each round t and each agent $v \in \mathcal{V}$, the action selection distribution p_t^v only depends on $\bigcup_{u \in \mathcal{V}} \mathcal{F}_{t-d(u,v)-1}^u$.*

From Definition 4.2.1, the communication between any two agents u and v in \mathcal{G} comes with a delay equal to $d(u, v) + 1$. Here we give some examples of π in different settings:

- when $|\mathcal{V}| = 1$ and $I_t^v = \ell_t^v$, π is an online learning algorithm for learning with expert advice problems (Cesa-Bianchi et al., 1997),
- when $|\mathcal{V}| = 1$ and $I_t^v = \ell_t^v(a_t^v)$, π is a sequential algorithm for a multi-armed bandit problem (Auer et al., 2002a),
- when $I_t^v \in \partial f_i(x_t^v)$ for some convex function $f(x)$, π belongs to the black-box procedure for distributed convex optimization over a simplex (Scaman et al., 2019), and
- when \mathcal{G} is a star graph, π is a centralized federated bandit algorithm discussed in Réda et al. (2022).

4.3 Lower bounds

In this section, we show two lower bounds on the cumulative regret of any federated bandit algorithm π in which all agents exchange their messages through the communication network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, for full-information and bandit feedback

setting. Both lower bounds highlight how the cumulative regret of the federated algorithm π is related to the minimum time it takes for all agents in \mathcal{G} to reach an agreement on a globally best arm.

Agents reaching an agreement about a globally best arm in hindsight is to find $i^* \in \arg \min_{i \in \mathcal{A}} \sum_{t=1}^T \bar{\ell}_t(i)$ by each agent v exchanging their private information about $\{\sum_{t=1}^T \ell_t^v(i) : i \in \mathcal{A}\}$ with their neighbors. This is known as a distributed consensus averaging problem (Boyd et al., 2004). Let $d_v = |\mathcal{N}(v)|$ and $d_{\max} = \max_{v \in \mathcal{V}} d_v$ and $d_{\min} = \min_{v \in \mathcal{V}} d_v$. The dynamics of a consensus averaging procedure is usually characterized by spectrum of the Laplacian matrix M of graph \mathcal{G} defined as

$$M_{u,v} := \begin{cases} d_u & \text{if } u = v \\ -1 & \text{if } u \neq v \text{ and } (u, v) \in \mathcal{E} \\ 0 & \text{otherwise.} \end{cases}$$

Let $\lambda_1(M) \geq \dots \geq \lambda_N(M) = 0$ be the eigenvalues of the Laplacian matrix M . The second smallest eigenvalue $\lambda_{N-1}(M)$ is the *algebraic connectivity* which approximates the sparsest-cut of graph \mathcal{G} (Arora et al., 2009). In the following two theorems, we show that for any federated bandit algorithm π , there always exists a problem instance and an agent whose worst-case cumulative regret is $\Omega(\lambda_{N-1}(M)^{-1/4} \sqrt{T})$.

Theorem 4.3.1 (Full-information feedback). *For any federated bandit algorithm π , there exists a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with Laplacian matrix M and a full-information feedback instance $L \in [0, 1]^{T \times N \times K}$ such that for some $v_1 \in \mathcal{V}$,*

$$R_T^{v_1}(\pi, L) = \Omega \left(\sqrt[4]{\frac{1 + d_{\max}}{\lambda_{N-1}(M)}} \sqrt{T \log K} \right). \quad (4.3.1)$$

The proof, in Section 4.6.1, relies on the existence of a graph in which there exist two clusters of agents, A and B , with distance $d(A, B) = \min_{u \in A, v \in B} d(u, v) = \Omega \left(\sqrt{(d_{\max} + 1) / \lambda_{N-1}(M)} \right)$. Then, we consider an instance where only agents in cluster A receive non-zero losses. Based on a reduction argument, the cumulative regrets for agents in cluster B are the same as (up to a constant factor) the cumulative regret in a single-agent adversarial bandit problem with feedback of delay $d(A, B)$ (see Lemma 4.6.4 in Section 4.6.1). Hence, one can show that the cumulative regret of agents in cluster B is $\Omega \left(\sqrt{d(A, B)} \sqrt{T \log K} \right)$.

Note that the doubly adversarial federated bandit with full-information feedback is a special case of distributed online linear optimization, with the decision set being a $K - 1$ -dimensional simplex. Hence, Theorem 4.3.1 immediately implies a regret lower bound for the distributed online linear optimization problem. To the best of our knowledge, this is the first non-trivial lower bound that relates the hardness of distributed online linear optimization problem to the algebraic connectivity of the communication network.

Leveraging the lower bound for the full-information setting, we show a lower bound for the bandit feedback setting.

Theorem 4.3.2 (Bandit feedback). *For any federated bandit algorithm π , there exists a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with Laplacian matrix M and a bandit feedback instance $L \in [0, 1]^{T \times N \times K}$ such that for some $v_1 \in \mathcal{V}$,*

$$R_T^{v_1}(\pi, L) = \Omega \left(\max \left\{ \sqrt{\frac{1}{1 + d_{v_1}}} \sqrt{KT}, \sqrt[4]{\frac{1 + d_{\max}}{\lambda_{N-1}(M)}} \sqrt{T \log K} \right\} \right). \quad (4.3.2)$$

The proof is provided in Section 4.6.2. The lower bound contains two parts. The first part, derived from the information-theoretic argument in Shamir (2014), captures the effect from bandit feedback. The second part is inherited from Theorem 4.3.1 by the fact that the regret of an agent in bandit feedback setting cannot be smaller than its regret in full-information setting.

4.4 FEDEXP3: a federated regret-minimization algorithm

Inspired by the fact that the cumulative regret is related to the time need to reach consensus about a globally best arm, we introduce a new federated bandit algorithm based on the gossip communication mechanism, called FEDEXP3. The details of FEDEXP3 are described in Algorithm 5. We shall also show that FEDEXP3 has a sub-linear cumulative regret upper bound which holds for all agents simultaneously.

The FEDEXP3 algorithm is adapted from the Exp3 algorithm, in which each agent v maintains an estimator $z_t^v \in \mathbb{R}^K$ of the cumulative losses for all arms and a tentative action selection distribution $x_t^v \in [0, 1]^K$. At the beginning of each time step t , each agent follows the action selection distribution x_t^v with probability $1 - \gamma_t$, or performs a uniform random exploration with probability γ_t . Once the action a_t^v is sampled, the agent observes the associated loss $\ell_t^v(a_t^v)$ and then computes an importance-weighted loss estimator $g_t^v \in \mathbb{R}^K$ using the sampling

probability $p_t^v(a_t^v)$. Before g_t^v is integrated into the cumulative loss estimator z_{t+1}^v , the agent communicates with its neighbors to average its cumulative loss estimator z_t^v .

Algorithm 5: FEDEXP3

Input : Non-increasing sequence of learning rates $\{\eta_t > 0\}$,
non-increasing sequence of exploration ratios $\{\gamma_t > 0\}$,
and a gossip matrix $W \in [0, 1]^{N \times N}$.

Initialization: $z_1^v(i) = 0, x_1^v(i) = 1/K$ for all $i \in \mathcal{A}$ and $v \in \mathcal{V}$.

- 1 **for** each time step $t = 1, 2, \dots, T$ **do**
- 2 **for** each agent $v \in \mathcal{V}$ **do**
- 3 compute the action distribution $p_t^v(i) = (1 - \gamma_t)x_t^v(i) + \gamma_t/K$;
- 4 choose the action $a_t^v \sim p_t^v$;
- 5 compute the loss estimators $g_t^v(i) = \ell_t^v(i)\mathbb{I}\{a_t^v = i\} / p_t^v(i)$ for all
 $i \in \mathcal{A}$;
- 6 update the gossip accumulative loss $z_{t+1}^v = \sum_{u:(u,v) \in \mathcal{E}} W_{u,v}z_t^u + g_t^v$;
// communication step
- 7 ;
- 8 update the exploitation distribution $x_{t+1}^v = \frac{\exp(-\eta_t z_{t+1}^v(i))}{\sum_{j \in \mathcal{A}} \exp(-\eta_t z_{t+1}^v(j))}$;
- 9 **end**
- 10 **end**

The communication step is characterized by the *gossip* matrix which is a doubly stochastic matrix $W \in [0, 1]^{N \times N}$ satisfying the following constraints

$$\sum_{v \in \mathcal{V}} W_{u,v} = \sum_{u \in \mathcal{V}} W_{u,v} = 1$$

and $W_{u,v} \geq 0$ where equality holds when $(u, v) \notin \mathcal{E}$. This gossip communication step facilitates the agents to reach a consensus on the estimators of the cumulative losses of all arms, and hence allows the agents to identify a globally best arm in hindsight. We present below an upper bound on the cumulative regret of each agent in FEDEXP3.

Theorem 4.4.1. *Assume that the network runs Algorithm 5 with*

$$\gamma_t = \sqrt[3]{\frac{\left(C_W + \frac{1}{2}\right) K^2 \log K}{t}}$$

and

$$\eta_t = \frac{\log K}{T\gamma_T} = \sqrt[3]{\frac{(\log K)^2}{\left(C_W + \frac{1}{2}\right) K^2 T^2}}$$

with $C_W = \min\{2\log T + \log N, \sqrt{N}\}/(1 - \sigma_2(W)) + 3$. Then, the expected regret of each agent $v \in \mathcal{V}$ is bounded as

$$R_T^v = \tilde{O}\left(\frac{1}{\sqrt[3]{1 - \sigma_2(W)}} K^{2/3} T^{2/3}\right)$$

where $\sigma_2(W)$ is the second largest singular value of W .

Proof sketch Let $\hat{\ell}_t$ and \bar{z}_t be the average instant loss estimator and average cumulative loss estimator,

$$f_t = \frac{1}{N} \sum_{v \in \mathcal{V}} g_t^v \quad \text{and} \quad \bar{z}_t = \frac{1}{N} \sum_{v \in \mathcal{V}} z_t^v,$$

and let y_t be action distribution that minimizes the regularized average cumulative loss estimator

$$y_t(i) = \frac{\exp(-\eta_{t-1} \bar{z}_t(i))}{\sum_{j \in \mathcal{A}} \exp(-\eta_{t-1} \bar{z}_t(j))}.$$

The cumulative regret can be bounded by the sum of three terms

$$\begin{aligned} R_t^v \leq & \underbrace{\mathbb{E} \left[\sum_{t=1}^T (\langle f_t, y_t^v \rangle - f_t(i^*)) \right]}_{\text{FTRL}} \\ & + K \underbrace{\sum_{t=1}^T \eta_{t-1} \mathbb{E} \|z_t^v - \bar{z}_t\|_*}_{\text{CONSENSUS}} + \underbrace{\sum_{t=1}^T \gamma_t}_{\text{EXPLORATION}} \end{aligned}$$

where $i^* \in \arg \min_{i \in \mathcal{A}} \sum_{t=1}^T \bar{\ell}_t(i)$ is a globally best arm in hindsight.

The FTRL term is a typical regret term from the classic analysis for the Follow-The-Regularized-Leader algorithm (Lattimore and Szepesvári, 2020). The CONSENSUS term measures the cumulative approximation error generated during the consensus reaching process, which can be bounded using the convergence analysis of distributed averaging algorithm based on doubly stochastic matrices (Duchi et al., 2011; Hosseini et al., 2013). The last EXPLORATION term can be bounded by specifying the time-decaying exploration ratio γ_t . The full proof of

Theorem 4.4.1 is provided in Section 4.6.4.

The FEDEXP3 algorithm is also a valid algorithm for the multi-agent adversarial bandit problem (Cesa-Bianchi et al., 2016) which is a special case of the doubly adversarial federated bandit problem when $\ell_t^v(i) = \ell_t(i)$ for all $v \in \mathcal{V}$. According to the distributed consensus process of FEDEXP3, each agent $v \in \mathcal{V}$ only communicates cumulative loss estimator values z_t^v , instead of the actual loss values $\ell_t^v(a_t^v)$, and the selection distribution p_t^v . FEDEXP3 can guarantee a sub-linear regret without the exchange of sequences of selected arm identities or loss sequences of agents, which resolves an open question raised in Cesa-Bianchi et al. (2016).

Choice of the gossip matrix The gossip matrix W can be constructed using the *max-degree* trick in Duchi et al. (2011), i.e.,

$$W = I - \frac{D - A}{2(1 + d_{\max})}$$

where $D = \text{diag}(d_1, \dots, d_N)$ and A is the adjacency matrix of \mathcal{G} . This construction of W requires that all agents have knowledge of the maximum degree d_{\max} , which can indeed be easily computed in a distributed system by nodes exchanging messages and updating their states using the maximum reduce operator.

Another choice of W comes from the effort to minimize the cumulative regret. The leading factor $1/\sqrt[3]{1 - \sigma_2(W)}$ in the regret upper bound of FEDEXP3 can be minimized by choosing a gossip matrix W with smallest $\sigma_2(W)$. Suppose that the agents have knowledge of the topology structure of the communication graph \mathcal{G} , then the agents can choose the gossip matrix to minimize their regret by solving the following convex optimization problem:

$$\begin{aligned} & \text{minimize} && \|W - (1/n)\mathbf{1}\mathbf{1}^T\|_2 \\ & \text{subject to} && W \geq 0, W\mathbf{1} = \mathbf{1}, W = W^T, \\ & && W_{ij} = 0, \text{ for } (i, j) \notin \mathcal{E} \end{aligned}$$

which has an equivalent semi-definite programming formulation as noted in Boyd et al. (2004).

Gap between upper and lower bounds The regret upper bound of FEDEXP3 algorithm in Theorem 4.4.1 grows sublinearly in the number of arms K and horizon time T . There is only a small polynomial gap between the regret upper bound

and the lower bound in Theorem 4.3.2 with respect to these two parameters. The regret upper bound depends also on the second largest singular value $\sigma_2(W)$ of W . The related term in the lower bound in Theorem 4.3.2 is the second smallest eigenvalue $\lambda_{N-1}(M)$ of the Laplacian matrix M . To compare these two terms, we point that when the gossip matrix is constructed using the max-degree method, as discussed in Corollary 1 in Duchi et al. (2011),

$$\frac{1}{\sqrt[3]{1 - \sigma_2(W)}} \leq \sqrt[3]{2 \frac{d_{\max} + 1}{\lambda_{N-1}(M)}}.$$

With respect to $\sqrt[4]{(d_{\max} + 1)/\lambda_{N-1}(M)}$ in Theorem 4.3.2, there is only a small polynomial gap between the regret upper bound and the lower bound. We note that a similar dependence on $\sigma_2(W)$ is present in the analysis of distributed optimization algorithms (Duchi et al., 2011; Hosseini et al., 2013).

4.5 Numerical experiments

We present experimental results for the FEDEXP3 algorithm (W constructed by the max-degree method) using both synthetic and real-world datasets. We aim to validate our theoretical analysis and demonstrate the effectiveness of FEDEXP3 on finding a globally best arm in non-stationary environments. All the experiments are performed with 10 independent runs. The code for producing our experimental results is available online in the Github repository: [link].

4.5.1 Synthetic datasets

We validate our theoretical analysis of the FEDEXP3 algorithm on synthetic datasets. The objective is two-fold. First, we demonstrate that the cumulative regret of FEDEXP3 grows sub-linearly with time. Second, we examine the dependence of the regret on the second largest singular value of the gossip matrix.

A motivation for finding a globally best arm in recommender systems is to provide recommendations for those users whose feedback is sparse. In this setting, we construct a federated bandit setting in which a subset of agents will be activated at each time step and only activated agents may receive non-zero loss. Specifically, we set $T = 3,000$ with $N = 36$ and $K = 20$. At each time step t , a subset U_t of $N/2$ agents are selected from \mathcal{V} with replacement. For all activated agents U_t , the loss for arm i is sampled independently from Bernoulli distribution




Figure 5: (Top) The average cumulative regret, i.e. $\sum_{v \in \mathcal{V}} R_T^v / N$, versus T , for FEDEXP3 and Exp3 on the grid graph. (Down) The average cumulative regret versus $(1 - \sigma_2(W))^{-1/3}$ for FEDEXP3 on different networks at $T = 3000$.

with mean $\mu_i = (i - 1) / (K - 1)$. All non-activated agents receive a loss of 0 for any arm they choose at time step t .

We evaluate the performance of FEDEXP3 on different networks, i.e. for a complete graph, a \sqrt{N} by \sqrt{N} grid network, and random geometric graphs. The random geometric graph RGG(d) is constructed by uniform random placement of each node in $[0, 1]^2$ and connecting any two nodes whose distance is less or equal to d (Penrose, 2003). Random geometric graphs are commonly used for modeling spatial networks.

In our experiments, we set $d \in \{0.3, \dots, 0.9\}$. The results in Figure 5 confirm that the cumulative regret of the FEDEXP3 algorithm grows sub-linearly with respect to time and suggest that the cumulative regret of FEDEXP3 is proportional to $(1 - \sigma_2(W))^{-1/3}$. This is compatible with the regret upper bound in Theorem 4.4.1.

4.5.2 MovieLens dataset: recommending popular movie genres

We compare FEDEXP3 with a UCB-based federated bandit algorithm in a movie recommendation scenario using a real-world dataset. In movie recommendation settings, users' preferences over different genres of movies can change over time. In such non-stationary environments, we demonstrate that a significant performance improvement can be achieved by FEDEXP3 against the GossipUCB algorithm (we refer to as GUCB) proposed in Zhu et al. (2021), which is defined for stationary settings.

We evaluate the two algorithms using the MovieLens-Latest-full dataset which contains 58,000 movies, classified into 20 genres, with 27,000,000 ratings




Figure 6: The average cumulative regret versus horizon time for FEDEXP3 and GUCB in the movie recommendation setting with the communication networks: (top) complete graph, (mid) the grid network, and (down) RGG(0.5).

(rating scores in $\{0.5, 1, \dots, 5\}$) from 280,000 users. Among all the users, there are 3,364 users who rated at least one movie for every genre. We select these users as our agents, i.e. $N = 3,364$, and the 20 genres as the arms to be recommended, i.e. $K = 20$.

We create a federated bandit environment for movie recommendation based on this dataset. Let $m^v(i)$ be the number of ratings that agent v has for genre i . We set the horizon $T = \max_{v \in \mathcal{N}} m^v(i) = 12,800$. To reflect the changes in agents' preferences over genres as time evolves, we sort the ratings in an increasing order by their Unix timestamps and construct the loss tensor in the following way. Let $r_j^v(i)$ be the j -th rating of agent v on genre i , the loss of recommending an movie to agent v of genre i at time step t is defined as

$$\ell_t^v(i) = \frac{5.5 - r_j^v(i)}{5.5}$$

for $t \in \left[(j-1) \left\lfloor \frac{T}{m^v(i)} \right\rfloor, j \left\lfloor \frac{T}{m^v(i)} \right\rfloor \right)$. The performance of FEDEXP3 and GUCB is shown in Figure 6. The results demonstrate that FEDEXP3 can outperform GUCB by a significant margin for different communication networks.

4.6 Proofs

Notation For a vector x , we use $x(i)$ to denote the i -th coordinate of x . We define $\mathcal{F}_t = \bigcup_{v \in \mathcal{V}} \mathcal{F}_t^v$ where \mathcal{F}_t^v is the sequence of agent v 's actions and feedback up to time step t , i.e., $\mathcal{F}_t^v = \bigcup_{s=1}^t \{a_s^v, I_s^v\}$.

4.6.1 Proof of Theorem 4.3.1

We first define a new class of cluster-based distributed online learning procedure, referred to as *cluster-based federated algorithms*, in which the delay only occurs when the communication is between different clusters. The regret lower bound for federated bandit algorithms will be no less than the regret lower bound for cluster-based federated algorithms, as shown in Lemma 4.6.2. Then we show in Lemma 4.6.3 that there exists a special graph in which there exist two clusters of agents A and B with distance $d(A, B) = \min_{u \in A, v \in B} d(u, v) = \Omega\left(\sqrt{(d_{\max} + 1)/\lambda_{N-1}(M)}\right)$. Then, we consider an instance where only agents in cluster A receive non-zero losses. Based on a reduction argument, the cumulative regrets of agents in cluster B are the same as (up to a constant factor) the cumulative regrets in a single-agent adversarial bandit setting with feedback delay $d(A, B)$ (see Lemma 4.6.4 in Section 4.6.1). Hence, one can show that the cumulative regret of agents in cluster B is $\Omega\left(\sqrt{d(A, B)}\sqrt{T \log K}\right)$.

We denote with $d(\mathcal{U}, \mathcal{U}')$ the smallest distance between any two nodes in $\mathcal{U}, \mathcal{U}' \subset \mathcal{V}$, i.e.

$$d(\mathcal{U}, \mathcal{U}') = \min_{u \in \mathcal{U}, u' \in \mathcal{U}'} d(u, u')$$

where $d(u, v)$ is the length of a shortest path connecting u and v .

Definition 4.6.1 (Cluster-based federated algorithms). *A cluster-based federated algorithm is a multi-agent learning algorithm defined by a partition of graph $\bigcup_r \mathcal{U}_r = \mathcal{V}$ where \mathcal{U}_r is called cluster. In the cluster-based federated algorithm, at each round t , the action selection probability p_t^v of agent $v \in \mathcal{U}_r$ depends on the history information up to round $t - d(\mathcal{U}_r, \mathcal{U}_{r'}) - 1$ of all agents $u' \in \mathcal{U}_{r'}$.*

Note that when all agents are in the same cluster \mathcal{V} , the centralized federated algorithm in Réda et al. (2022) is an instance of a cluster-based federated algorithm.

Lemma 4.6.2 (Monotonicity). *Let Π and Π' be two sets of all cluster-based federated algorithms with two partitions $\bigcup_r \mathcal{U}_r$ and $\bigcup_s \mathcal{U}'_s$, respectively. Suppose for any cluster \mathcal{U}'_s of π' , there exists a cluster \mathcal{U}_r of π such that $\mathcal{U}'_s \subset \mathcal{U}_r$, then*

$$\Pi' \subset \Pi \quad \text{and} \quad \min_{\pi \in \Pi} R_T^v(\pi, L) \leq \min_{\pi' \in \Pi'} R_T^v(\pi', L)$$

for any $L \in [0, 1]^{T \times N \times K}$ and any $v \in \mathcal{V}$.

Proof. It suffices to show $\Pi' \subset \Pi$. Consider a cluster-based federated algorithm $\pi' \in \Pi'$. For any agent $v \in \mathcal{V}$, let \mathcal{U}'_s be the cluster of v in Π' . By definition of cluster-based procedure, agent v 's action selection distribution probability p_t^v depends on the history information up to round $t - d(\mathcal{U}'_s, \mathcal{U}'_h) - 1$ of all agents $u' \in \mathcal{U}'_h$.

By the assumption, there exists two subset $\mathcal{U}_{r_1}, \mathcal{U}_{r_2} \subset \mathcal{V}$ such that $\mathcal{U}'_s \subset \mathcal{U}_{r_1}$ and $\mathcal{U}'_h \subset \mathcal{U}_{r_2}$. Hence $d(\mathcal{U}_{r_1}, \mathcal{U}_{r_2}) \leq d(\mathcal{U}'_s, \mathcal{U}'_h)$, from which it follows $t - d(\mathcal{U}'_s, \mathcal{U}'_h) - 1 \leq t - d(\mathcal{U}_{r_1}, \mathcal{U}_{r_2}) - 1$. Hence $\pi' \in \Pi$ which completes the proof. \square

Lemma 4.6.3. *There exists a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with N nodes and a matrix $M \in \mathcal{M}_{\mathcal{G}}$, together with two subsets of nodes $I_0, I_1 \subset \mathcal{V}$ of size $|I_0| = |I_1| \geq N/4$ and such that*

$$d(I_0, I_1) \geq \tilde{\Delta},$$

where $d(I_0, I_1)$ is the shortest-path distance in \mathcal{G} between the two sets and

$$\tilde{\Delta} = \frac{\sqrt{2}}{3} \sqrt{\frac{1 + d_{\max}}{\lambda_{N-1}(M)}}.$$

Proof. From Lemma 24 in Scaman et al. (2019), there exists exists a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with N nodes and a matrix $M \in \mathcal{M}_{\mathcal{G}}$, together with two subsets of nodes $I_0, I_1 \subset \mathcal{V}$ of size $|I_0| = |I_1| \geq N/4$ and such that

$$d(I_0, I_1) \geq \frac{\sqrt{2}}{3} \sqrt{\frac{\lambda_1(M)}{\lambda_{N-1}(M)}}.$$

We show that $\lambda_1(M) \geq 1 + d_{\max}$. To see this, note that $\lambda_1(M)$ is the Rayleigh quotient $\max_{x \neq 0} \frac{x^T M x}{x^T x}$. By the definition of the Laplacian matrix,

$$x^T M x = \sum_{(v,u) \in \mathcal{E}} (x_v - x_u)^2.$$

Let v be a vertex whose degree is d_{\max} and

$$x_u := \begin{cases} \sqrt{\frac{d_{\max}}{1+d_{\max}}} & \text{if } u = v \\ -\frac{1}{\sqrt{d_{\max}}\sqrt{1+d_{\max}}} & \text{if } u \neq v \text{ and } v_i \text{ is adjacent to } v_j \\ 0 & \text{otherwise} \end{cases}$$

then

$$\begin{aligned}
\sum_{(v,u) \in \mathcal{E}} (x_v - x_u)^2 &= d_{\max} \left(\sqrt{\frac{d_{\max}}{1+d_{\max}}} + \frac{1}{\sqrt{d_{\max}}\sqrt{1+d_{\max}}} \right)^2 \\
&= d_{\max} \left(\frac{1}{\sqrt{d_{\max}}\sqrt{1+d_{\max}}} \right)^2 (d_{\max} + 1)^2 \\
&= 1 + d_{\max}
\end{aligned}$$

and

$$\sum_{u \in \mathcal{V}} x_u^2 = \frac{d_{\max}}{1+d_{\max}} + d_{\max} \frac{1}{d_{\max}(1+d_{\max})} = 1.$$

Hence, $\lambda_1(M) \geq 1 + d_{\max}$. □

Let I_0, I_1 be two subsets of nodes satisfying

$$d(I_0, I_1) \geq \tilde{\Delta} \quad \text{and} \quad |I_0| = |I_1| = N/4.$$

The number of rounds needed to communicate between any node in I_0 and any node I_1 is at least $\tilde{\Delta}$.

Lemma 4.6.4. *Let $v_0 \in I_0$ and $v_1 \in \mathcal{V} \setminus I_0$. Consider a cluster-based federated algorithm with clusters I_0 and $V \setminus I_0$. Then, any distributed online learning algorithm σ for full information feedback setting has an expected regret*

$$R_T^{v_1} \geq \frac{1 - o(1)}{4} \sqrt{\frac{(\tilde{\Delta} + 1)}{2}} T \log K$$

as $T \rightarrow \infty$.

Proof. Consider an online learning with expert advice problem with the action set \mathcal{A} over B rounds (Cesa-Bianchi et al., 1997). Let ℓ'_1, \dots, ℓ'_B be an arbitrary sequence of losses and p'_b be the action selection distribution at round b . We show that σ can be used to design an algorithm for this online learning with expert advice problem, adapted from Cesa-Bianchi et al. (2016).

Consider the loss sequences $\{\ell_t^v\}_{t=1}^T$ for each $v \in \mathcal{V}$ with $T = (\tilde{\Delta} + 1)B$ such that

$$\ell_t^v = \begin{cases} \ell'_{\lceil t/(\tilde{\Delta}+1) \rceil}, & v \in I_0 \\ 0 & \text{otherwise.} \end{cases}$$

Let p_t^v be the action select distribution of agent $v \in \mathcal{V}$ running the algorithm σ . Define the algorithm for the online learning with expert advice problem as follows:

$$p'_b = \frac{1}{\tilde{\Delta} + 1} \sum_{s=1}^{\tilde{\Delta}+1} p_{(\tilde{\Delta}+1)(b-1)+s}^{v_1}$$

where $p_t^v = (1/k, \dots, 1/k)$ for all $t \leq 1$ and $v \in \mathcal{V}$.

Note that p'_b is defined by $p_{(\tilde{\Delta}+1)(b-1)+1}^{v_1}, \dots, p_{(\tilde{\Delta}+1)b}^{v_1}$. These are in turn defined by $\ell_1^{v_0}, \dots, \ell_{(\tilde{\Delta}+1)(b-1)}^{v_0}$ by the definition of cluster-based communication protocol. Also note that $\lceil t/(\tilde{\Delta} + 1) \rceil \leq b - 1$ for $t \leq (\tilde{\Delta} + 1)b$, hence p'_b is determined by $\ell'_1, \dots, \ell'_{b-1}$. Therefore p'_1, \dots, p'_B are generated by a legitimate algorithm for online learning with expert advice problem.

Note that the cumulative regret of agent v_1 is

$$\begin{aligned} \sum_{t=1}^T \langle p_t^{v_1}, \bar{\ell}_t \rangle &= \frac{1}{N} \sum_{t=1}^T \left[\sum_{v \in I_0} \langle p_t^{v_1}, \ell_t^v \rangle + \sum_{v \notin I_0} \langle p_t^{v_1}, \ell_t^v \rangle \right] \\ &= \frac{1}{4} \sum_{t=1}^T \langle p_t^{v_1}, \ell'_{\lceil t/(\tilde{\Delta}+1) \rceil} \rangle \\ &= \frac{1}{4} \sum_{b=1}^B \sum_{s=1}^{\tilde{\Delta}+1} \langle p_{(\tilde{\Delta}+1)(b-1)+s}^{v_1}, \ell'_b \rangle \\ &= \frac{\tilde{\Delta} + 1}{4} \sum_{b=1}^B \langle p'_b, \ell'_b \rangle \end{aligned} \quad (4.6.1)$$

where the second equality comes from the definition of ℓ_t^v and the fourth equality comes from the definition of p'_b .

Also note that

$$\begin{aligned} \min_{i \in \mathcal{A}} \sum_{t=1}^T \bar{\ell}_t(i) &= \frac{1}{4} \min_{i \in \mathcal{A}} \sum_{t=1}^T \ell'_{\lceil t/(\tilde{\Delta}+1) \rceil}(i) \\ &= \frac{\tilde{\Delta} + 1}{4} \min_{i \in \mathcal{A}} \sum_{b=1}^B \ell'_b(i). \end{aligned} \quad (4.6.2)$$

From (4.6.1) and (4.6.2), it follows that

$$\sum_{t=1}^T \langle p_t^{v_1}, \bar{\ell}_t \rangle - \min_{i \in \mathcal{A}} \sum_{t=1}^T \bar{\ell}_t(i) = \frac{\tilde{\Delta} + 1}{4} \left[\sum_{b=1}^B \langle p'_b, \ell'_b \rangle - \min_{i \in \mathcal{A}} \sum_{b=1}^B \ell'_b(i) \right].$$

There exists a sequence of losses ℓ'_1, \dots, ℓ'_B such that for any algorithm for online

learning with expert advice problem, the expected regret satisfies (Cesa-Bianchi and Lugosi, 2006, Theorem 3.7),

$$\sum_{b=1}^B \langle p'_b, \ell'_b \rangle - \min_{i \in \mathcal{A}} \sum_{b=1}^B \ell'_b(i) \geq (1 - o(1)) \sqrt{\frac{B}{2} \ln K}.$$

Hence, we have

$$\sum_{t=1}^T \langle p_t^{v_1}, \bar{\ell}_t \rangle - \min_{i \in \mathcal{A}} \sum_{t=1}^T \bar{\ell}_t(i) \geq \frac{1 - o(1)}{4} \sqrt{(\tilde{\Delta} + 1) \frac{T}{2} \ln K}.$$

□

4.6.2 Proof of Theorem 4.3.2

The lower bound contains two parts. The first part is derived by using information-theoretic arguments in Shamir (2014) and it captures the effect of bandit feedback. The second part is inherited from the full-information feedback lower bound in Theorem 4.3.1 by the fact that the regret of an agent in the bandit feedback setting cannot be smaller than the regret in the full-information setting.

Consider a centralized federated algorithm with all the agents in the same cluster \mathcal{V} , denoted as Π^C . Note that by Lemma 4.6.2, for a federated bandit algorithm Π^G ,

$$\Pi^G \subset \Pi^C \quad \text{and} \quad \min_{\pi' \in \Pi^C} R_T^v(\pi', L) \leq \min_{\pi \in \Pi^G} R_T^v(\pi, L)$$

for any $L \in [0, 1]^{T \times N \times K}$ and any $v \in \mathcal{V}$.

For any $\pi' \in \Pi^C$, at each round t , every agent $v \in \mathcal{V}$ receives $O(|\mathcal{N}(v)|)$ bits since its neighboring agents can choose at most $|\mathcal{N}(v)|$ distinct actions. By Theorem 4 in Shamir (2014), there exists some distribution \mathcal{D} over $[0, 1]^K$ such that loss vectors $\bar{\ell}_t \stackrel{i.i.d.}{\sim} \mathcal{D}$ for all $t = 1, 2, \dots, T$ and $\min_{i \in \mathcal{A}} \mathbb{E} \left[\sum_{t=1}^T \bar{\ell}_t(a_t(v)) - \sum_{t=1}^T \bar{\ell}_t(i) \right] = \Omega \left(\min \{ T, \sqrt{KT / (1 + |\mathcal{N}(v)|)} \} \right)$.

Hence, it follows that

$$\begin{aligned}
\min_L R_T^v(\pi, L) &\geq \min_L R_T^v(\pi', L) \\
&\geq \mathbb{E}_{\bar{\ell}_t \sim \mathcal{D}} \left[\sum_{t=1}^T \bar{\ell}_t(a_t(v)) - \min_{i \in \mathcal{A}} \sum_{t=1}^T \bar{\ell}_t(i) \right] \\
&\geq \max_{i \in \mathcal{A}} \mathbb{E}_{\bar{\ell}_t \sim \mathcal{D}} \left[\sum_{t=1}^T \bar{\ell}_t(a_t(v)) - \sum_{t=1}^T \bar{\ell}_t(i) \right] \\
&\geq \min_{i \in \mathcal{A}} \mathbb{E}_{\bar{\ell}_t \sim \mathcal{D}} \left[\sum_{t=1}^T \bar{\ell}_t(a_t(v)) - \sum_{t=1}^T \bar{\ell}_t(i) \right] \\
&= \Omega \left(\min \{ T, \sqrt{KT/(1 + |\mathcal{N}(v)|)} \} \right)
\end{aligned}$$

where the third inequality comes from Jensen's inequality.

Also note that any federated bandit algorithm for bandit feedback setting is also a federated bandit algorithm for full-information setting, from which it follows

$$\begin{aligned}
\min_L R_T^v(\pi, L) &\geq \max \left\{ \Omega \left(\min \{ T, \sqrt{KT/(1 + |\mathcal{N}(v)|)} \} \right), \Omega \left(\sqrt[4]{\frac{1 + d_{\max}}{\lambda_{N-1}(M)}} \sqrt{T \log K} \right) \right\} \\
&= \Omega \left(\min \left\{ T, \max \left\{ \sqrt{K/(1 + |\mathcal{N}(v)|)}, \sqrt[4]{\frac{1 + d_{\max}}{\lambda_{N-1}(M)}} \sqrt{\log K} \right\} \sqrt{T} \right\} \right).
\end{aligned}$$

4.6.3 Auxiliary lemmas

Here we present some auxiliary lemmas which are used in the proof of Theorem 4.4.1. Recall that $\hat{\ell}_t$ and \bar{z}_t are the average instant loss estimator and average cumulative loss,

$$f_t = \frac{1}{N} \sum_{v \in \mathcal{V}} g_t^v \quad \text{and} \quad \bar{z}_t = \frac{1}{N} \sum_{v \in \mathcal{V}} z_t^v$$

and y_t is action distribution to minimize the regularized average cumulative loss

$$y_t(i) = \frac{\exp(-\eta_{t-1} \bar{z}_t(i))}{\sum_{j \in \mathcal{A}} \exp(-\eta_{t-1} \bar{z}_t(j))}.$$

Lemma 4.6.5. For each time step $t = 1, \dots, T$,

$$\bar{z}_{t+1} = \bar{z}_t + f_t$$

and

$$\max\{\|g_t^v\|_*, \|f_t\|_*\} \leq \frac{K}{\gamma_t}.$$

Proof.

$$\begin{aligned} \bar{z}_{t+1} &= \frac{1}{N} \sum_{v \in \mathcal{V}} z_{t+1}^v \\ &= \frac{1}{N} \sum_{v \in \mathcal{V}} \sum_{u: (u,v) \in \mathcal{E}} W_{u,v} z_t^u + \frac{1}{N} \sum_{v \in \mathcal{V}} g_t^v \\ &= \frac{1}{N} \sum_{v \in \mathcal{V}} z_t^v + \frac{1}{N} \sum_{v \in \mathcal{V}} g_t^v \\ &= \bar{z}_t + f_t \end{aligned}$$

where the second equality comes from Line 7 in Algorithm 5 and the third equality comes from the double-stochasticity of W .

Noting that $p_t^v(i) \geq \gamma/K$ for all $v \in \mathcal{V}, i \in \mathcal{A}$ and $t \in \{1, \dots, T\}$, it follows that

$$\|g_t^v\|_* = \frac{\ell_t^v(a_t^v)}{p_t^v(a_t^v)} \leq \frac{K}{\gamma_t} \quad \text{and} \quad \|f_t\|_* \leq \frac{1}{N} \sum_{v \in \mathcal{V}} \|g_t^v\|_* \leq \frac{K}{\gamma_t}.$$

□

Lemma 4.6.6. *For any $v \in \mathcal{V}$ and $t \geq 1$, it holds that*

$$\mathbb{E}[g_t^v \mid \mathcal{F}_{t-1}] = \ell_t^v \quad \text{and} \quad \mathbb{E}[f_t \mid \mathcal{F}_{t-1}] = \bar{\ell}_t$$

with

$$\mathbb{E}[\|f_t\|_*] \leq K \quad \text{and} \quad \mathbb{E}[\|f_t\|_*^2] \leq \frac{K^2}{\gamma_t}.$$

Proof. Note that p_t^v is determined by \mathcal{F}_{t-1} , hence

$$\mathbb{E}[g_t^v(i) \mid \mathcal{F}_{t-1}] = \frac{\ell_t^v(i)}{p_t^v(i)} \mathbb{E}[\mathbb{I}\{a_t^v = i\} \mid \mathcal{F}_{t-1}] = \frac{\ell_t^v(i)}{p_t^v(i)} p_t^v(i) = \ell_t^v(i)$$

and

$$\mathbb{E}[\|g_t^v\|_*] = \mathbb{E}\left[\frac{\ell_t^v(a_t^v)}{p_t^v(a_t^v)}\right] = \mathbb{E}\left[\mathbb{E}\left[\frac{\ell_t^v(a_t^v)}{p_t^v(a_t^v)} \mid \mathcal{F}_{t-1}\right]\right] = \mathbb{E}\left[\sum_{i \in \mathcal{A}} p_t^v(i) \frac{\ell_t^v(i)}{p_t^v(i)}\right] = \sum_{i \in \mathcal{A}} \ell_t^v(i) \leq K$$

where the last inequality comes from $\ell_t^v(i) \leq 1$. Since $f_t(i) = \frac{1}{N} \sum_{v \in \mathcal{V}} g_t^v$, it follows

that

$$\mathbb{E} [f_t(i) \mid \mathcal{F}_{t-1}] = \frac{1}{N} \sum_{v \in \mathcal{V}} \ell_t^v(i) = \bar{\ell}_t(i)$$

and

$$\mathbb{E} [\|f_t\|_*] \leq \frac{1}{N} \sum_{v \in \mathcal{V}} \mathbb{E} [\|g_t^v\|_*] \leq K$$

which comes from Jensen's inequality. Notice that

$$\begin{aligned} \mathbb{E} [\|g_t^v\|_*^2] &= \mathbb{E} \left[\frac{\ell_t^v(a_t^v)^2}{p_t^v(a_t^v)^2} \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\frac{\ell_t^v(a_t^v)^2}{p_t^v(a_t^v)^2} \mid \mathcal{F}_{t-1} \right] \right] \\ &= \mathbb{E} \left[\sum_{i \in \mathcal{A}} p_t^v(i) \frac{\ell_t^v(i)^2}{p_t^v(i)^2} \right] \leq \mathbb{E} \left[\sum_{i \in \mathcal{A}} \frac{1}{p_t^v(i)} \right] \leq \frac{K^2}{\gamma_t} \end{aligned}$$

where the last inequality comes from $p_t^v(i) \geq \gamma_t/K$. Again, from Jensen's inequality, it follows

$$\mathbb{E} [\|f_t\|_*^2] \leq \frac{1}{N} \sum_{v \in \mathcal{V}} \mathbb{E} [\|g_t^v\|_*^2] \leq \frac{K^2}{\gamma_t}.$$

□

Before presenting the next lemma, we recall the definition of strongly-convex functions and Fenchel duality. A function ϕ is said to be α -strongly convex function on a convex set \mathcal{X} if

$$\phi(x') \geq \phi(x) + \langle \nabla \phi(x), x' - x \rangle + \frac{1}{2} \alpha \|x' - x\|^2$$

for all $x', x \in \mathcal{X}$, for some $\alpha \geq 0$.

Let ϕ^* denote the *Fenchel conjugate* of ϕ , i.e.,

$$\phi^*(y) = \max_{x \in \mathcal{X}} \{\langle x, y \rangle - \phi(x)\}$$

with the projection,

$$\nabla \phi^*(y) = \arg \max_{x \in \mathcal{X}} \{\langle x, y \rangle - \phi(x)\}.$$

Lemma 4.6.7. *Let ψ the normalized negative entropy function Lattimore and Szepesvári*

(2020) on $\mathcal{P}_{K-1} = \{x \in [0, 1]^K : \sum_{i=1}^K x(i) = 1\}$,

$$\psi_\eta(x) = \frac{1}{\eta} \sum_{i=1}^k x(i) (\log(x(i)) - 1).$$

For all $t = 1, \dots, T$, it holds that

$$x_t^v = \operatorname{argmin}_{x \in \mathcal{P}_{K-1}} \{\langle x, z_t^v \rangle + \psi_{\eta_t}(x)\} = \nabla \psi_{\eta_{t-1}}^*(-z_t^v)$$

with $\mathcal{X} = \mathcal{P}_{K-1}$ and

$$y_t = \operatorname{argmin}_{x \in \mathcal{P}_{K-1}} \{\langle x, \bar{z}_t \rangle + \psi_{\eta_{t-1}}(x)\} = \nabla \psi_{\eta_{t-1}}^*(-\bar{z}_t).$$

Furthermore, it holds

$$\|x_t^v - y_t\| \leq \eta_{t-1} \|z_t^v - \bar{z}_t\|_*.$$

Proof. We prove for $y_t = \operatorname{argmin}_{x \in \mathcal{P}_{K-1}} \{\langle x, \bar{z}_t \rangle + \psi_{\eta_{t-1}}(x)\}$ whose argument also applies to x_t^v .

Notice it suffices to consider the minimization problem

$$\begin{aligned} \min_{x \in \mathcal{P}_{K-1}} \quad & \eta_{t-1} \sum_{k=1}^K x(i) \bar{z}_t(i) + \sum_{i=1}^k x(i) \log(x(i)) \\ \text{subject to} \quad & \sum_{k=1}^K x(i) = 1. \end{aligned}$$

It suffices to consider the Lagrangian,

$$\mathcal{L} = -\eta_{t-1} \sum_{k=1}^K x(i) \bar{z}_t(i) - \sum_{i=1}^k x(i) \log(x(i)) - \lambda \left(\sum_{k=1}^K x(i) - 1 \right).$$

Consider the first-order conditions for all $i = 1, \dots, K$

$$\frac{\partial \mathcal{L}}{\partial x(i)} = -\eta_{t-1} \bar{z}_t(i) - \log(x(i)) - 1 - \lambda = 0$$

which gives $x(i) = \exp(-\eta_{t-1} \bar{z}_t(i)) / \exp(1 + \lambda)$ for all $i = 1, \dots, K$. Plugging into the constraint $\sum_{k=1}^K x(i) = 1$ together with the definition of Fenchel duality Hiriart-Urruty and Lemarechal (2010) completes the proof for y_t .

Note that the normalized negative entropy $\psi(x)$ is 1-strongly convex,

$$\psi(x') \geq \psi(x) + \langle \nabla \psi(x), x' - x \rangle + \frac{1}{2} \|x' - x\|^2.$$

Multiplying $1/\eta_{t-1}$ both sides of the inequality yields that $\psi_{\eta_{t-1}}(x)$ is $1/\eta_{t-1}$ -strongly convex. By Theorem 4.2.1 in Hiriart-Urruty and Lemarechal (2010), we have that $\nabla \psi_{\eta_{t-1}}^*(z)$ is η_{t-1} -Lipschitz.

It follows that

$$\|p_t^v - \bar{p}_t\| = \|\nabla \psi_{\eta}^*(-z_t^v) - \nabla \psi_{\eta}^*(-\bar{z}_t)\| \leq \eta_{t-1} \|\bar{z}_t - z_t^v\|_*.$$

□

We state an upper bound on the network disagreement on the cumulative loss estimators from Duchi et al. (2011) and Hosseini et al. (2013).

Lemma 4.6.8. *For any $v \in \mathcal{V}$ and $t = 1, 2, \dots, T$,*

$$\|\bar{z}_t - z_t^v\|_* \leq \frac{K}{\gamma_T} \left(\frac{\min\{2 \log T + \log n, \sqrt{n}\}}{1 - \sigma_2(W)} + 3 \right) = \frac{K}{\gamma_T} C_W$$

where $\sigma_2(W)$ is the second largest singular value of W .

Proof. From Lemma 4.6.5, it follows that $\|g_t^v\|_* \leq K/\gamma_t$. Since $\{\gamma_t\}$ is non-increasing, let $L = K/\gamma_T$ in Eq. (29) in Duchi et al. (2011) and Lemma 6 in Hosseini et al. (2013) completes the proof. □

4.6.4 Proof of Theorem 4.4.1

Let $i^* = \arg \min_{i \in \mathcal{A}} \sum_{t=1}^T \bar{\ell}_t(i)$. Note that p_t^v is determined by \mathcal{F}_{t-1} and $\mathbb{E}[f_t | \mathcal{F}_{t-1}] = \bar{\ell}_t$ from Lemma 4.6.6. It follows that for each agent $v \in \mathcal{V}$

$$\begin{aligned} R_T^v &= \mathbb{E} \left[\sum_{t=1}^T \langle \bar{\ell}_t, p_t^v \rangle - \sum_{t=1}^T \bar{\ell}_t(i^*) \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \langle \mathbb{E}[f_t | \mathcal{F}_{t-1}], p_t^v \rangle - \sum_{t=1}^T \mathbb{E}[f_t(i^*) | \mathcal{F}_{t-1}] \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \langle f_t, p_t^v \rangle - \sum_{t=1}^T f_t(i^*) \right]. \end{aligned} \tag{4.6.3}$$

By the definition of p_t^v , it follows

$$\begin{aligned}
R_T^v &= \mathbb{E} \left[\sum_{t=1}^T (\langle f_t, (1-\gamma)x_t^v + \gamma x_1^v \rangle - f_t(i^*)) \right] \\
&= \mathbb{E} \left[\sum_{t=1}^T (1-\gamma_t) (\langle f_t, x_t^v \rangle - f_t(i^*)) \right] + \sum_{t=1}^T \gamma_t \mathbb{E} [\langle f_t, x_1^v \rangle - f_t(i^*)] \\
&= \mathbb{E} \left[\sum_{t=1}^T (1-\gamma_t) (\langle f_t, x_t^v \rangle - f_t(i^*)) \right] + \sum_{t=1}^T \gamma_t (\langle \bar{\ell}_t, x_1^v \rangle - \bar{\ell}_t(i^*)) \\
&\leq \mathbb{E} \left[\sum_{t=1}^T (\langle f_t, x_t^v \rangle - f_t(i^*)) \right] + \sum_{t=1}^T \gamma_t \\
&= \underbrace{\mathbb{E} \left[\sum_{t=1}^T (\langle f_t, y_t^v \rangle - f_t(i^*)) \right]}_{\text{(I)}} + \underbrace{\mathbb{E} \left[\sum_{t=1}^T \langle f_t, x_t^v - y_t^v \rangle \right]}_{\text{(II)}} + \sum_{t=1}^T \gamma_t
\end{aligned}$$

where the first inequality comes from the fact that $\gamma_t > 0$ and the fact that $\|\bar{\ell}_t\|_* \leq \sum_{v \in \mathcal{V}} 1/N \|\bar{\ell}_t^v\|_* \leq 1$.

From Lemma 4.6.5, it follows $\bar{z}_t = \sum_{s=1}^{t-1} f_s$. Hence, it follows from Lemma 4.6.7 that

$$y_t = \arg \min_{x \in \mathcal{P}_{K-1}} \left\{ \sum_{s=1}^t \langle f_s, x \rangle + \frac{1}{\eta_{t-1}} \psi(x) \right\}.$$

From Lemma 3 in Duchi et al. (2011) and Corollary 28.8 in Lattimore and Szepesvári (2020), we have

$$\text{(I)} \leq \frac{1}{2} \sum_{t=1}^T \eta_{t-1} \mathbb{E} [\|f_t\|_*^2] + \frac{1}{\eta_T} \log(K) \quad (4.6.4)$$

which is because $\{\eta_t\}$ is a non-increasing sequence. Note that

$$\|x_t^v - y_t\| \leq \eta_{t-1} \|z_t^v - \bar{z}_t\|_*$$

by Lemma 4.6.7. This yields that

$$\text{(II)} \leq \sum_{t=1}^T \eta_{t-1} \mathbb{E} [\|f_t\|_* \|z_t^v - \bar{z}_t\|_*]. \quad (4.6.5)$$

Plugging Equations (4.6.4) and (4.6.5) into (I) yields that

$$\begin{aligned}
R_T^v &\leq \frac{1}{2} \sum_{t=1}^T \eta_{t-1} \mathbb{E} \left[\|f_t\|_*^2 \right] + \frac{1}{\eta_T} \log(K) + \sum_{t=1}^T \eta_{t-1} \mathbb{E} \left[\|f_t\|_* \|z_t^v - \bar{z}_t\|_* \right] + \sum_{t=1}^T \gamma_t \\
&\leq \frac{1}{2} \sum_{t=1}^T \eta_{t-1} \mathbb{E} \left[\|f_t\|_*^2 \right] + \frac{K}{\gamma_T} C_W \sum_{t=1}^T \eta_{t-1} \mathbb{E} \left[\|f_t\|_* \right] + \sum_{t=1}^T \gamma_t + \frac{1}{\eta_T} \log(K) \\
&\leq \frac{K^2}{2} \sum_{t=1}^T \frac{\eta_{t-1}}{\gamma_t} + \frac{K^2}{\gamma_T} C_W \sum_{t=1}^T \eta_{t-1} + \sum_{t=1}^T \gamma_t + \frac{1}{\eta_T} \log(K)
\end{aligned}$$

where the second inequality comes from Lemma 4.6.8 and the third inequality comes from Lemma 4.6.6.

Let

$$\gamma_t = \sqrt[3]{\frac{\left(C_W + \frac{1}{2}\right) K^2 \log K}{t}} \quad \text{and} \quad \eta_t = \frac{\log K}{T \gamma_T} = \sqrt[3]{\frac{(\log K)^2}{\left(C_W + \frac{1}{2}\right) K^2 T^2}}.$$

Then, for every $v \in \mathcal{V}$, we have

$$\begin{aligned}
R_T^v &\leq \frac{3}{8} \sqrt[3]{\frac{K^2 \log K}{\left(C_W + \frac{1}{2}\right)^2} T^{\frac{2}{3}}} + \sqrt[3]{K^2 \log K \frac{C_W^3}{\left(C_W + \frac{1}{2}\right)^2} T^{\frac{2}{3}}} \tag{4.6.6} \\
&\quad + \frac{3}{2} \sqrt[3]{\left(C_W + \frac{1}{2}\right) K^2 \log K T^{\frac{2}{3}}} + \sqrt[3]{\left(C_W + \frac{1}{2}\right) K^2 \log K T^{\frac{2}{3}}} \\
&\leq \frac{3}{4} \sqrt[3]{K^2 \log K T^{\frac{2}{3}}} + \sqrt[3]{C_W K^2 \log K T^{\frac{2}{3}}} + \frac{5\sqrt[3]{2}}{2} \sqrt[3]{C_W K^2 \log K T^{\frac{2}{3}}} \\
&\leq 5 \sqrt[3]{C_W K^2 \log K T^{\frac{2}{3}}}.
\end{aligned}$$

Chapter 5

Communication-Efficient Federated Online Convex Optimization

In a cooperative learning system, multiple agents exchange information to make predictions and update their models, and the amount and frequency of communication can greatly impact the accuracy, speed, and scalability of the system. If the communication complexity is high, it can lead to increased communication overhead, which can result in slower and less efficient decision-making. This is especially important in real-world applications, where communication resources are limited and the volume of data generated is vast. The challenge here is to design cooperative learning algorithms which can achieve the desired performance while minimizing the communication cost.

5.1 Introduction

Large-scale federated learning systems have garnered significant research attention, due to their advantages in terms of privacy, security, and scalability (Kairouz et al., 2021). Prominent technology companies have implemented federated learning systems for various applications, such as learning from wearable devices, sentiment analysis, and location-based services Bonawitz et al. (2019). However, a major challenge in the deployment of high-performance large-scale federated learning systems is the high communication overhead associated with the training process. Designing a communication-efficient federated learning method that balances the trade-off between communication costs and convergence rate remains a challenging task.

Removing the communication bottleneck for federated learning methods has been the focus of much research (Hamer et al., 2020; Reisizadeh et al., 2020; Rothchild et al., 2020). A number of quantization and sketching methods have been proposed to reduce the *per-message communication costs*, that is, the coding length of each message transmitted within the federated learning systems. In this chapter, we aim to decrease the *communication complexity*, which is measured as the number of messages transmitted during the training process. Prior work has

relied on deterministic communication protocols, which facilitate elegant analysis of communication complexity. However, the system lacks the flexibility to balance the trade-off between communication complexity and convergence rate.

In this chapter, we investigate the problem of online convex optimization in the federated learning setting. The agents in the network can exchange information with their neighbors. At each time step, each agent makes a decision, receives feedback, and communicates with its neighbors. We consider a heterogeneous feedback scenario, where different agents receive feedback from different convex functions at each time step. The goal is to identify the best global solution retrospectively, with the lowest cumulative loss in terms of the averaged convex function over all agents in the network.

5.1.1 Related work

In this chapter, we examine the field of communication-efficient algorithms for Distributed Online Convex Optimization (DOCO) (Duchi et al., 2011; Hosseini et al., 2013; Wan et al., 2020; Cesa-Bianchi et al., 2020). Our study builds upon the previous research in this area and provides a comprehensive review of the recent literature in this field. The objective of this review is to showcase the unique contributions of our work and differentiate it from the prior research.

The focus of communication-efficient algorithms for DOCO is to simultaneously minimize the amount of information transmitted between nodes in the distributed system and the cumulative regret over a given horizon. Previous works in this area have explored two main directions. The first direction is to minimize the coding length of the transmitted message, while the second direction is to reduce the rounds of communication between nodes.

One popular method to reduce the coding length of each message is the use of quantization methods, which perform lossy compression on the original information contained in the message. Simple quantization methods, such as 1BitSGD (Seide et al., 2014; Strom, 2015), were initially proposed for the distributed training of deep neural networks with Stochastic Gradient Descent (SGD). Subsequently, more advanced methods, such as QSGD (Alistarh et al., 2017) and the random-dropping strategy (Wangni et al., 2018), have been developed with convergence guarantees. Quantization methods have also been studied in the context of federated learning (Hamer et al., 2020; Reisizadeh et al., 2020; Rothchild et al., 2020).

Reducing the rounds of communication, referred to as *communication complex-*

ity, has become a crucial area of research in communication-efficient DOCO in recent years (Zhang et al., 2017; Wan et al., 2020; Reisizadeh et al., 2020). Zhang et al. (2017) proposed the first conditional gradient descent algorithm for the distributed online setting, which reduces the cost of expensive projection operations by using simpler linear optimization steps. Subsequently, Wan et al. (2020) presented the D-OCG method, which divides the total rounds into \sqrt{T} equal-sized blocks. Reisizadeh et al. (2020) proposed the periodic communication strategy to minimize the communication complexity. Both methods enjoy a $O(\sqrt{T})$ communication complexity. It remains unknown before this thesis whether the communication complexity can be further reduced to $o(\sqrt{T})$.

In this chapter, we focus on reducing the communication complexity for distributed online convex optimization problems. Specifically, our algorithm builds on the Dual Average algorithm (Duchi et al., 2011; Nesterov, 2009; Hosseini et al., 2013). We propose a stochastic communication protocol instead of a deterministic communication protocol. By setting the communication probability as $T^{-\alpha}$ for any $\alpha \in [0, 1)$, we show that the resulting communication complexity scales as $\Theta(T^{1-\alpha})$ and the regret upper bound scales as $\tilde{O}(\sqrt{T^{1+\alpha}})$ with high probability. Therefore, our algorithm balances a trade-off between the communication complexity and the regret.

The minimization of communication complexity has also been discussed in the context of Distributed Offline Convex Optimization and Stochastic Multi-Agent Multi-Armed Bandits. Interested readers are referred Smith et al. (2018) and Agarwal et al. (2022) for a comprehensive review of these topics.

5.1.2 Organization of this chapter

In Section 5.2, the problem setting is elucidated and the metric of communication complexity is rigorously defined. Subsequently, Section 5.3 presents a communication-efficient federated online convex optimization algorithm, equipped with both regret and communication complexity bounds. The validity of our results is demonstrated in Section 5.4 through a rigorous mathematical proof, which can be found in Section 5.5.

5.2 Federated online convex optimization

Consider a communication network defined by an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of N agents and $(u, v) \in \mathcal{E}$ if agent u and agent v can directly

exchange messages. We assume that \mathcal{G} is simple, i.e. it contains no self loops nor multiple edges. The agents in the communication network collaboratively aim to solve a distributed online convex optimization problem. In this problem, there is a convex decision set $\mathcal{X} \subset \mathbb{R}^K$ and a fixed time horizon T .

At each time step t , each agent v will choose its action $x_t^v \in \mathcal{X}$, observe the feedback g_t^v . At the end of each time step, each agent $v \in \mathcal{V}$ can communicate with their neighbors $\mathcal{N}(v) = \{u \in \mathcal{V} : (u, v) \in \mathcal{E}\}$.

Let $f_t^v(\cdot)$ be a convex function defined on \mathcal{X} for agent v at time step t , and f_t be the network loss function defined as

$$f_t(x) = \frac{1}{N} \sum_{v \in \mathcal{V}} f_t^v(x) \quad \text{for } x \in \mathcal{X}.$$

Each agent v receives a loss of $f_t(x_t^v)$ at the time step t .

The performance of each agent $v \in \mathcal{V}$ is measured by its *regret*, defined as the difference of the expected cumulative loss incurred and the cumulative loss of a *globally* best fixed solution in hindsight, i.e.

$$R_t^v = \sum_{t=1}^v f_t(x_t^v) - \min_{x \in \mathcal{X}} \sum_{t=1}^T f_t(x).$$

Here we add two assumptions on the problem we consider, which are standard in online convex optimization:

Assumption 1 (Gradient-based optimization). $g_t^v \in \partial f_t^v(x_t^v)$

Assumption 2 (Lipschitz-continuity). *There exists some $L > 0$ such that $\|g_t^v\|_* \leq L$ for all $v \in \mathcal{V}$ and $t = 1, 2, \dots, T$.*

Except for the regret of each agent, we also want to minimize the *communication complexity* of the system, defined as

$$Q_T = \sum_{t=1}^T m_t$$

where m_t is the number of messages transmitted in the communication network at time step t .

The ODDA algorithm proposed by Hosseini et al. (2013); Duchi et al. (2011), which is designed to minimize the regret, enjoys a regret upper bound of $O(T^{1/2})$ with communication complexity of $\Theta(T)$. We show in the next section a stochas-

tic algorithm which has sub-linear regret upper bounds and communication complexity in high probability simultaneously.

5.3 FEDOCO: a communication-efficient stochastic regret-minimization algorithm

Algorithm 6: FEDOCO

Input : Non-increasing sequence of learning rates $\{\eta_t > 0\}$, a regularizer ψ and a skipping parameter $\alpha \in [0, 1)$.

Initialization: $z_1^v = 0$ and $x_1^v = \max_{x \in \mathcal{X}} \psi(x)$ for all $v \in \mathcal{V}$.

- 1 **for** each time step $t = 1, 2, \dots, T$ **do**
- 2 sample $X_t \sim \text{Bernoulli}(T^{-\alpha})$;
- 3 **if** $X_t = 1$ **then**
- 4 uniformly sample an edge $(u, u') \in \mathcal{E}$;
- 5 **end**
- 6 **for** each agent $v \in \mathcal{V}$ **do**
- 7 select x_t^v ;
- 8 receive feedback g_t^v ;
- 9 update the accumulative loss ; // communication step
- 10
$$z_{t+1}^v = \begin{cases} \frac{1}{2} (z_t^u + z_t^{u'}) + g_t^v, & v \in \{u, u'\} \quad \text{and} \quad X_t = 1 \\ z_t^v + \ell_t^v, & \text{otherwise} \end{cases}$$
- 11 update $x_{t+1}^v = \Pi_{\mathcal{X}}^{\psi} (z_{t+1}^v, \eta_t)$;
- 12 **end**

In this section, we present a novel communication-efficient stochastic algorithm, named FEDOCO. The proposed algorithm enables the communication complexity to grow sub-linearly with the horizon T , while ensuring that the cumulative regret of each agent remains sub-linear in the same horizon. This innovative approach offers a compelling solution to the challenge of limiting communication while maintaining the performance of the system.

The FEDOCO algorithm is parameterized by α , called *skipping parameter*, which lies in the interval $[0, 1)$. The selection of α enables the FEDOCO algorithm to strike a balance between the cumulative regret of each agent and the communication complexity of the system.

FEDOCO takes a *regularizer* function ψ as an input argument. The regularizer ψ is a strongly convex function defined on \mathcal{X} with $\max_{x \in \mathcal{X}} \psi(x) < \infty$. And the

decision variable chosen by agent v is computed by solving the minimization algorithm

$$x_t^v = \Pi_{\mathcal{X}}^\psi(z_t^v, \eta_{t-1}) = \arg \min_{x \in \mathcal{X}} \left\{ \langle z_t^v, x \rangle + \frac{1}{\eta_{t-1}} \psi(x) \right\}.$$

In the communication step of FEDOCO, each agent has the opportunity to either be matched with a single neighbor to communicate with, or skip the communication step. Specifically, at the start of time step t , the FEDOCO will decide whether this time step is a communication step or skipping step. When it is a communication step, the system will select one edge from the network at random, and the agents in the edge will average their estimates before integrating the new observation g_t^v . In a skipping step, the agents will not communicate with others.

The probability of a communication step is inversely proportional to a polynomial function in the horizon T . Consequently, the communication complexity scales sub-linearly with T with high probability. The details of FEDOCO are described in Algorithm 6.

5.4 Trade-off between communication complexity and regret

In this section, a theoretical examination of the cumulative regret and communication complexity of each agent in the FEDOCO algorithm is performed. The objective is to derive upper bounds for these metrics.

Recall that $\delta_{\min} = \min\{\delta_v : v \in \mathcal{V}\}$. Let $D = \text{diag}\{\delta_v : v \in \mathcal{V}\}$, A be the adjacent matrix of \mathcal{G} and M be its Laplacian matrix.

Theorem 5.4.1. *Let $R > 0$ such that $\psi(x^*) \leq R^2$. Assume that the network runs Algorithm 6 with $\alpha \in [0, 1)$ and*

$$\eta_t = \frac{R\sqrt{1 - \lambda_2(W)}}{L\sqrt{t}}.$$

Then, for any $\delta \in (0, 1)$, with probability at least

$$1 - \frac{1}{T} - 2 \exp\left(-\frac{\delta^2}{3} T^{1-\alpha}\right)$$

the communication complexity

$$Q_T \in \left[2(1 - \delta)T^{1-\alpha}, \quad 2(1 + \delta)T^{1-\alpha} \right]$$

and the cumulative regret of each agent $v \in \mathcal{V}$ is bounded as

$$R_T^v \leq \tilde{O} \left(\sqrt{\frac{|\mathcal{E}|}{\delta_{\min} \lambda_{N-1}(M)}} \sqrt{T^{1+\alpha}} \right)$$

where $\lambda_{N-1}(M)$ is the second smallest eigenvalue value of the Laplacian matrix.

Proof sketch The proof of Theorem 5.4.1 relies on a decomposition used in the proof of Theorem 4.4.1, replacing the expectation argument with a high probability argument.

Let $\bar{z}_t = \frac{1}{N} \sum_{v \in \mathcal{V}} z_t^v$ and $y_{t+1} = \Pi_{\mathcal{X}}^\psi(\bar{z}_{t+1}, \eta_t)$. The cumulative regret of agent v can be bounded by the sum of two terms

$$R_T^v \leq \underbrace{\sum_{t=1}^T \frac{1}{N} \sum_{v \in \mathcal{V}} \langle g_t^v, y_t - x^* \rangle}_{\text{FTRL}} + \underbrace{\sum_{t=1}^T \left(\frac{2L}{n} \sum_{i=1}^n \|x_t^v - y_t\| + L \|x_t^v - y_t\| \right)}_{\text{CONSENSUS}}.$$

With the probabilistic peer-to-peer communication mechanism in Algorithm 6, the CONSENSUS term is upper bounded by $\tilde{O}(1/(1 - \lambda_2(W)))$ where

$$W = I - \frac{1}{2T^\alpha |\mathcal{E}|} (D - A).$$

Note that $D^{1/2} M D^{1/2} = D - A$, and from Lemma 4 in Duchi et al. (2011),

$$\lambda_2(W) \leq 1 - \frac{\delta_{\min}}{|\mathcal{E}| T^\alpha} \lambda_{N-1}(M).$$

Choosing the adequate learning rate η_t completes the proof. The details of the proof is presented in Section 5.5.

Communication complexity and regret trade-off Note that both the communication complexity and regret of FEDOCO scale sub-linearly with respect to the horizon. As the skipping parameter α approaches 1, the number of communication steps decreases. As a result, with high probability, the communication complexity decreases as $T^{1-\alpha}$ and the cumulative regret upper bound increases

as $\sqrt{T^{(1+\alpha)}}$ (ignoring logarithmic terms). Note that the case when $\alpha = 1/2$ recovers the communication complexity of $O(\sqrt{T})$ and the regret upper bound of $O(T^{3/4})$ presented in the analysis of online conditional gradient algorithms in the distributed setting (Wan et al., 2020).

5.5 Proofs

5.5.1 Proof of Theorem 5.4.1

Communication complexity Note that at each time step, with probability $T^{-\alpha}$, there are two agents exchanging their messages; otherwise there is no message exchanged. Hence the cumulative number of messages exchanged

$$Q_T = \sum_{t=1}^T m_t = 2 \sum_{t=1}^T X_t.$$

Note that $\{X_t : t = 1, 2, \dots, T\}$ are i.i.d. Bernoulli random variables with mean $T^{-\alpha}$. From the Chernoff bound for binomial distribution (Mitzenmacher and Upfal, 2017), we have for any $\delta \in (0, 1)$

$$\mathbb{P}(Q_T \geq 2(1 + \delta)T^{1-\alpha}) \leq \exp\left(-\frac{\delta^2 T^{1-\alpha}}{3}\right)$$

and

$$\mathbb{P}(Q_T \leq 2(1 - \delta)T^{1-\alpha}) \leq \exp\left(-\frac{\delta^2 T^{1-\alpha}}{2}\right).$$

Cumulative regret The proof follows the general approach in Duchi et al. (2011); Hosseini et al. (2013) adapted to the stochastic gossip matrix setting. It follows that , it follows that

$$\begin{aligned} R_T^v &\leq \sum_{t=1}^T (f_t(y_t) - f_t(x^*) + L \|x_t^v - y_t\|) \\ &\leq \sum_{t=1}^T \left(\frac{1}{N} \sum_{v \in \mathcal{V}} \underbrace{\langle g_t^v, x_t^v - x^* \rangle}_{(*)} + \frac{L}{N} \sum_{v \in \mathcal{V}} \|x_t^v - y_t\| + L \|x_t^v - y_t\| \right) \end{aligned}$$

where the first inequality comes from the fact that f_t^v is L -Lipschitz continuous, the second inequality comes from the fact that f_t^v is convex function.

For term (*), we can expand as

$$\begin{aligned}\langle g_t^v, x_t^v - x^* \rangle &= \langle g_t^v, y_t - x^* \rangle + \langle g_t^v, x_t^v - y_t \rangle \\ &\leq \langle g_t^v, y_t - x^* \rangle + L \|x_t^v - y_t\|\end{aligned}$$

the inequality comes from the fact that $\|g_t^v\|_* \leq L$.

It follows that

$$R_T^v \leq \underbrace{\sum_{t=1}^T \frac{1}{N} \sum_{v \in \mathcal{V}} \langle g_t^v, y_t - x^* \rangle}_{\text{(I)}} + \underbrace{\sum_{t=1}^T \left(\frac{2L}{N} \sum_{v \in \mathcal{V}} \|x_t^v - y_t\| + L \|x_t^v - y_t\| \right)}_{\text{(II)}} \quad (5.5.1)$$

Note that $\bar{z}_t = \sum_{s=1}^{t-1} \frac{1}{N} \sum_{v \in \mathcal{V}} g_s^v$. Hence, it follows from Lemma 4.6.7 that

$$y_t = \arg \min_{x \in \mathcal{X}} \left\{ \sum_{s=1}^t \left\langle \frac{1}{N} \sum_{v \in \mathcal{V}} g_s^v, x \right\rangle + \frac{1}{\eta_{t-1}} \psi(x) \right\}.$$

From Lemma 3 in Duchi et al. (2011) and Corollary 28.8 in Lattimore and Szepesvári (2020), we have

$$\text{(I)} \leq \frac{L^2}{2} \sum_{t=1}^T \eta_{t-1} + \frac{1}{\eta_T} \psi(x^*) \quad (5.5.2)$$

which is because $\{\eta_t\}$ is a non-increasing sequence. Note that

$$\|x_t^v - y_t\| \leq \eta_{t-1} \|z_t^v - \bar{z}_t\|_*$$

by Lemma 4.6.7. This yields that

$$\text{(II)} \leq L \sum_{t=1}^T \eta_t \left(\|\bar{z}_t - z_t^v\|_* + \frac{2}{N} \sum_{v \in \mathcal{V}} \|\bar{z}_t - z_t^v\|_* \right). \quad (5.5.3)$$

Plugging Equations (5.5.2) and (5.5.3) into Equation (5.5.1) yields that

$$R_T^v \leq \frac{L^2}{2} \sum_{t=1}^T \eta_{t-1} + \frac{1}{\eta_T} \psi(x^*) + L \sum_{t=1}^T \eta_t \left(\|\bar{z}_t - z_t^v\|_* + \frac{2}{N} \sum_{v \in \mathcal{V}} \|\bar{z}_t - z_t^v\|_* \right).$$

Let $e_v \in \{0, 1\}^K$ be the base vector whose coordinates are all 0 except that the v -th coordinate is 1, and (u_t, u_t') be the edge selected at time step t when $X_t = 1$.

Define the gossip matrix at time step t

$$W_t = I - \frac{1}{2}(e_{u_t} - e_{u'_t})(e_{u_t} - e_{u'_t})^\top \mathbb{I}\{X_t = 1\}$$

then

$$z_{t+1}^v = \sum_{u:(u,v) \in \mathcal{E}} [W_t]_{u,v} z_t^u + \ell_t^v$$

and

$$\mathbb{E}[W_t] = W = I - \frac{1}{T^\alpha |\mathcal{E}|} (D - A)$$

which comes from the fact that $\sum_{u:(u,v) \in \mathcal{E}} (e_{u_t} - e_{u'_t})(e_{u_t} - e_{u'_t})^\top = 2(D - A)$.

From Theorem 3 in Duchi et al. (2011), we have

$$\mathbb{P} \left[\max_{t \leq T, v \in \mathcal{V}} \|\bar{z}_t - z_t^v\|_* / L > \frac{6 \log(T^2 N)}{1 - \lambda_2(W)} + \frac{1}{T\sqrt{N}} + 2 \right] \leq \frac{1}{T}.$$

Let

$$\eta_t = \frac{R\sqrt{1 - \lambda_2(W)}}{L\sqrt{t}}.$$

Then, with probability at least $1 - 1/T$, for every $v \in \mathcal{V}$,

$$R_T^v \leq cRL \cdot \frac{\log(T\sqrt{N})}{\sqrt{1 - \lambda_2(W)}} \sqrt{T}$$

for some universal constant $c > 0$.

Note that $D - A = D^{1/2} \mathcal{L} D^{1/2}$, hence

$$\lambda_2(W) = \lambda_2 \left(I - \frac{1}{T^\alpha |\mathcal{E}|} D^{1/2} \mathcal{L} D^{1/2} \right) \leq 1 - \frac{\delta_{\min}}{T^\alpha |\mathcal{E}|} \lambda_{N-1}(M).$$

It follows that with probability at least $1 - 1/T$, for every agent v ,

$$R_T^v \leq cRL \cdot \log(T\sqrt{N}) \sqrt{\frac{|\mathcal{E}|}{\delta_{\min} \lambda_{N-1}(M)}} \sqrt{T^{1+\alpha}}$$

for some universal constant $c > 0$.

Applying the union bound on Q_T and R_T^v completes the proof.

Chapter 6

Greedy Bayes Incremental Matching Strategies

Efficiently finding a maximum value matching for a given set of nodes based on observations of values of matches over a sequence of rounds is a fundamental sequential learning problem. Here a matching refers to a set of matches of node pairs (or matches in short) such that each node is matched to at most one other node. This problem arises in many applications including team formation, online social networks, online gaming, online labor, and other online platforms. We may think of the nodes as agents in MACL systems and the value of a matched node pair to represent the reward or welfare generated from interaction or joint production of the two agents, contributing to the total reward or welfare. For example, in team production, the value of a matched pair of nodes may represent the quality of the work produced. In online social networks, the value of a matched pair of nodes may represent user satisfaction or social capital gained through interaction.

6.1 Introduction

We consider the case when the value of each matched pair of nodes is according to a function – we refer to as a value function – of latent node features in some domain. We focus on the case of a finite domain such that each node is either of low or high type, which can be represented by each node having value 0 or 1 – we refer to this as latent binary types. For value functions, we consider two Boolean functions, AND and OR. The types of nodes are a priori unknown and partial information about node types is acquired as values of matches are observed. The underlying value function plays a key role in how much information is revealed about individual node types from observed values of matches. For example, for Boolean OR function, having observed value of a match to be 0, we can immediately deduce that both nodes are of low type, and, otherwise, the only conclusion that can be drawn is that at least one node is of high type. The decision maker gains partial information about individual node types as outcomes of matches are observed over a sequence of rounds. We consider this decision problem under the constraint that sequence of matchings must be incremental, meaning that

in each round the selected matching is derived by a limited change of the matching in the previous round. Specifically, we consider the case when in each round the decision maker can only choose a matching that is identical to the matching selected in the previous round, except for possibly re-matching any two previously matched pairs. Learning how to match through an incremental sequence of matchings limits how much information can be gained about individual node types in each round. This constraint is motivated by applications in practice when matched node pairs remain matched over several rounds, allowing only a small change of matching to be made in each round.

We consider the underlying incremental matching problem for a set of n nodes in a Bayesian setting where latent binary types of nodes are according to a prior product-form Bernoulli distribution with parameter p . Here p is the probability that a node is of high type according to the prior distribution. We study greedy Bayes incremental matching strategies that in each round select a matching that has highest marginal value from the set of feasible matchings (obeying incremental matching constraint), according to posterior distribution, given the information observed up to the current round. We evaluate performance of greedy Bayes incremental matching strategies by considering expected Bayesian regret, defined as the difference between cumulative expected value of an optimal matching and expected value of the sequence of matchings chosen by the algorithm asymptotically for large number of rounds. Our goals are to identify greedy Bayes incremental matching strategies for both AND and OR value functions, and characterize expected Bayesian regret under these two functions.

6.1.1 Related work

Team formation problems were studied under various assumptions in different scientific fields and application domains. Kleinberg and Raghu (2018), Sekar et al. (2021), Lee et al. (2022) and Mehta et al. (2020) studied individual skill scoring methods for the best team selection, or assigning individuals to teams, for various functions determining team performance from individual performances. Neuman et al. (1999) studied the relationship between work team effectiveness and two aspects of the personality traits of team members, namely, the average level of a given trait within a team, and the variability in traits within a team. Carlier and Ekeland (2010) studied the problem of formations of teams through a pricing mechanism, where individuals are partitioned into groups of equal size and each team requires exactly one member from each group. When there are two

groups, the problem corresponds to the classic marriage problem. Liemhetcharat and Veloso (2012) studied the team formation problem where each individual has its own capabilities and the performance of a team depends on capabilities of team members as well their pairwise synergies. Rajkumar et al. (2017) studied the problem of partitioning individuals into disjoint groups for different objectives that account for pairwise compatibility of individuals. Flores-Szwagrzak and Treibich (2020) studied the problem of inferring individual from team productivity using a scoring method leveraging varying membership and level of success of all teams. Popular skill rating systems used in online gaming platforms, such as TrueSkill, assume team performances to be sum of individual performances with some added noise, and matching strategies are used aiming to match players of similar skills in order to make interesting competitions Graepel et al. (2006).

The most closely related work to ours is Johari et al. (2018) which introduced the problem of incremental matching with match outcomes according to a function of binary latent types. They considered regret for maximum value and minimum value functions. For maximum value (Boolean OR) function, their 1-chain algorithm is a greedy Bayes incremental matching algorithm, for which we show a tight regret characterization. For minimum value (Boolean AND), they studied an algorithm and established some comparison results. Our work identifies the greedy Bayes incremental matching algorithm and establishes a regret bound. Eichhorn et al. (2022) studied the same problem, but without the constraint that the sequence of matchings must be according to an incremental matching sequence.

The problem we consider has some connections with combinatorial bandits that have been studied under various assumptions, e.g. Cesa-Bianchi and Lugosi (2012), Combes et al. (2015), Kveton et al. (2015), Chen et al. (2016), Rejwan and Mansour (2020), Agarwal et al. (2021). For combinatorial bandits, in each round, the decision maker selects an arm that corresponds to a feasible subset of base arms. For example, given a set of N base arms, an arm is an arbitrary subset of base arms of cardinality K . Under full bandit feedback, the reward of the selected arm is observed, which is according to an underlying function of the corresponding base arms' values. For example, this may be sum of values Cesa-Bianchi and Lugosi (2012); Combes et al. (2015); Rejwan and Mansour (2020), or some non-linear function of values (such as maximum value) Chen et al. (2016); Agarwal et al. (2021). Under semi-bandit feedback, individual values of all selected base

arms are observed. In our setting, we may think of node pairs to be base arms, matchings to be arms, and we have semi-bandit feedback. Alternatively, we may regard nodes to be base arms, matchings to be arms, and we have full-bandit feedback for each match in a matching. Note that requiring for sequence of matching selections to be according to an incremental matching sequence is not studied in the line of work on combinatorial bandits.

Simchowitz et al. (2016) studied the best-of- K bandit problem in which at each time step, the decision maker chooses a subset of K arms from a set of N arms and observers reward equal to maximum of individual arm's values drawn from a joint distribution. Here the objective is to identify the subset that achieves the highest expected reward using as few queries as possible.

Regret minimization for maximum weight matching problems with stochastic edge weights whose mean values obey a rank-1 structure was studied in Sentenac et al. (2021). Here the mean weight $w_{u,v}$ for two nodes u and v is of the form $w_{u,v} = \theta_u \theta_v$, where $\theta_1, \dots, \theta_n$ are hidden node parameters with values in $[0, 1]$. The stochastic weight of an edge is defined as the sum of mean edge weight and a sub-Gaussian random noise. An adaptive matching algorithm is shown to have regret of order $n \log(T)$ up to a poly-logarithmic factor. Note that the rank-1 assumption along with assuming that node parameters are binary and zero noise boil, down to weights according to Boolean OR function. Our work is different in that we consider noiseless weights, binary node types, and we admit the incremental matchings constraint.

6.1.2 Organization of this chapter

In Section 6.2 we present problem formulation and notation used in this chapter. Section 6.3 contain our results, starting with a regret lower bound, and then showing greedy Bayes incremental matching algorithm and its analysis for each value function studied separately in different subsections. Section 6.4 contains proofs for theorems stated in this chapter.

6.2 Problem formulation

Let $V = \{1, \dots, n\}$ be a set of n nodes, where n is an even, positive integer. Each node $v \in V$ has a latent variable θ_v that takes value in $\Theta = \{0, 1\}$. We refer to the latent variable of a node as the *type* of node. Type 0 is referred to as *low type* and type 1 is referred to as *high type*.

The latent variables $\theta = (\theta_1, \dots, \theta_n)$ are assumed to have a prior distribution of product-form with identical Bernoulli marginal distributions, i.e. $\mathbb{P}[\theta_v = 1] = 1 - \mathbb{P}[\theta_v = 0] = p$, for all $v \in V$, for some $p \in [0, 1]$. Intuitively, parameter p characterises the population of nodes with respect to their types—we may think of as the portion of high-type nodes in the population.

Each match of a pair of nodes (u, v) , has *weight* $w_{u,v} = w(\theta_u, \theta_v)$ where w is a symmetric function $w : \Theta^2 \rightarrow \tilde{\Theta}$. A function is said to be symmetric if its output is invariant to permutations of its inputs. Set $\tilde{\Theta}$ is assumed to be fixed—its elements do not depend on either n or p . We refer to w as a *value function*, quantifying the value of a pair of nodes. When $\tilde{\Theta} = \{0, 1\}$, then w is a Boolean function. We consider two Boolean functions with dual input, which are defined as follows:

- AND: 1 if both inputs are 1, otherwise 0
- OR: 0 if both inputs are 0, otherwise 1.

We refer to Boolean OR as *maximum* value function, Boolean AND as *minimum* value function.

These value functions are natural choices and have been studied in the context of production systems. Boolean OR value function requires at least one node of a pair of nodes to be of high type for the value of the pair to be high. In other words, the value of a pair is the value of the highest type node of this pair. This function is sometimes referred to as *strongest link* function. In the other extreme, Boolean AND value function requires both nodes of a pair of nodes to be of high type for the value of the pair to be high. This value function is sometimes referred to as the *weakest link* function, indicating that the value of a pair of nodes is the value of the lowest type node of this pair.

The value $w_{u,v} = w(\theta_u, \theta_v)$ of a pair of nodes (u, v) may only reveal a *partial information* about values of node types θ_u and θ_v . We discuss this for functions w that we consider as follows. For Boolean OR value function, if $w_{u,v} = 0$, then we can deduce that *both* u and v are of low type, and, otherwise, if $w_{u,v} = 1$, then we can deduce that *either* u or v is of high type. For minimum (Boolean AND) value function, if $w_{u,v} = 0$, then we can deduce that *either* u or v is of low type, and, otherwise, if $w_{u,v} = 1$, then we can deduce that *both* u and v are of high type.

We consider the problem of finding a maximum value matching of nodes in a sequential decision making setting. An adaptive matching algorithm \mathcal{A} selects a matching $M_t^{\mathcal{A}}$ in each round $t = 1, \dots, T$ from a set of *permissible matchings* \mathcal{M}_t , observes the value of each pair in this matching, i.e. $Y_t = \{w(\theta_u, \theta_v) : (u, v) \in$




Figure 7: An example of an incremental matching sequence: for example, the second matching is derived from the first matching by re-matching node pairs (2,3) and (4,5).

$M_t^A\}$, and receives an instantaneous reward $\mu_w(M_t^A) = \sum_{(u,v) \in M_t^A} w(\theta_u, \theta_v)$. The set of permissible matchings has to satisfy the *incremental matching constraint*: each $M \in \mathcal{M}_t$ is such that M has at most two new matched pairs with respect to M_{t-1}^A . Formally, $\mathcal{M}_t = \{M \in \mathcal{M} : |M \Delta M_{t-1}^A| \leq 4\}$ where \mathcal{M} is the set of all matchings, and Δ is the symmetric set difference operator $A \Delta B = (A \setminus B) \cup (B \setminus A)$. The algorithm \mathcal{A} specifies a mapping from the history \mathcal{H}_t to M_t^A , where $\mathcal{H}_1 = \emptyset$ and $\mathcal{H}_t = \{M_1^A, Y_1, \dots, M_{t-1}^A, Y_{t-1}\}$, for $t > 1$. See Figure 7 for an illustration.

An adaptive matching obeying the incremental matching constraint is referred to as an incremental matching algorithm. An adaptive incremental matching algorithm that in each round re-matches pairs of nodes greedily with respect to the conditional expected instantaneous regret in the next round according to the posterior belief about node states is referred to as a *greedy Bayes algorithm*. For a given value function w , we let \mathcal{G} denote a greedy Bayes algorithm. Note that we omit explicit referral to w in the notation \mathcal{G} for simplicity.

For any given value function w , let M^* be a maximum value matching of nodes in V , i.e. $M^* \in \arg \max_M \mu_w(M)$. Then, the *regret* of an adaptive algorithm \mathcal{A} is defined as

$$R_w^A(T) = T\mathbb{E}[\mu_w(M^*)] - \mathbb{E} \left[\sum_{t=1}^T \mu_w(M_t^A) \right]$$

where in the right-hand side of the last equation, the expectations are taken with respect to the prior distribution of the node types and the prior distribution of node types and any randomization of algorithm \mathcal{A} , respectively. Note that $R_w^A(T)$ is the *expected Bayesian regret*. We aim to characterise the asymptotic regret of

Table 1: Regret lower bounds for different value functions.

w	Regret lower bound
OR	$\Omega(\min\{p, 1 - p\}^4 n^2)$
AND	$\Omega((p(1 - p))^2 n^2)$

greedy Bayes algorithm \mathcal{G} for some value functions w as the time horizon T goes to infinity and the number of nodes n goes to infinity.

Note that for any value function w and any adaptive algorithm \mathcal{A} , the initial matching is a random matching of nodes in V because the prior distribution of node types is non-informative for maximum value matching.

6.3 Algorithms and regret bounds

We first show a regret lower bound that holds for any adaptive matching algorithm under mild assumptions. Let \tilde{R}_w denote the expected instantaneous regret incurred for a random matching of nodes and value function w . The following lemma provides a regret lower bound that holds for any adaptive incremental matching algorithm.

Lemma 6.3.1 (lower bound). *For any adaptive incremental matching algorithm for value function w , there exists a constant $c > 0$ such that the expected regret is at least $c\tilde{R}_w^2$, provided that \tilde{R}_w is larger than some constant value that depends only on c and w .*

The proof of the lemma is provided in Section 6.4.1. The proof follows by the fact that any adaptive incremental matching algorithm reduces the instantaneous regret in each round for at most a constant amount. From this and some simple arguments, it follows that the cumulative regret is lower bounded by the quadratic function of the expected instantaneous regret of initial random matching.

Regret lower bounds for two value functions are shown in Table 1. These lower bounds follow from Lemma 6.3.1 and computations given in Section 6.4.1.

We will compare the regret upper bounds (which in some cases are tight) developed for greedy Bayes algorithms in the next sections and compare then with given lower bounds. Any additional factor in a regret upper bound that depends on the population parameter p to the one that is present in the corresponding lower bound quantifies the search cost of a greedy Bayes strategy.

We will see that different value functions w can be separated into two classes depending on whether or not they satisfy the constant improvement property. A

value function is said to satisfy *a constant improvement property* if in each round of the execution of the greedy Bayes algorithm, the instantaneous regret is guaranteed to be decreased for at least a constant amount L over a constant number of next M rounds, where L and M are constants that do not depend on either n or p . For the class of value functions that satisfy the constant improvement property, the regret of a greedy Bayes algorithm is essentially determined by the instantaneous regret incurred in the initial random matching and the regret differs from the corresponding lower bound only for a constant factor. For value functions that do not satisfy the constant improvement property, the regret has an additional factor depending on the population parameter p that is due to the search cost of the greedy Bayes algorithm.

6.3.1 Maximum value function (OR)

In this section, we consider the maximum value function, defined as follows

$$w(x, y) = x \vee y.$$

In this case, any optimal matching has the maximum number of matches of nodes of different types. Hence, an adaptive algorithm has to efficiently identify pairs of nodes of different types. Note that based on the observed value of a node pair, this node pair belongs to one of the following two sets $S_0 = \{(u, v) : w(\theta_u, \theta_v) = 0\}$ or $S_1 = \{(u, v) : w(\theta_u, \theta_v) = 1\}$. Here, S_0 is the set of all 0-0 type pairs, and S_1 is the set of all 0-1 and 1-1 type pairs. Under the incremental matching constraint, the conditional expected reward is maximized by choosing one pair from S_0 and one pair from S_1 and rematching them. This strategy underlies the greedy Bayes algorithm.

The main idea behind the algorithm is to exploit the *known* low type nodes to explore the *unknown* high type nodes. The algorithm maintains a set of node pairs known to be all of low type (thus each pair in this set has two nodes of low type) and a set of node pairs whose values were observed to be of high type (thus each pair in this set has at least one node of high type). In each round, the algorithm chooses a pair from the known set and a pair from the unknown set and then changes pairing of nodes of these two pairs. From observed pair values, the algorithm can identify the type of all nodes of the two chosen pairs. The algorithm terminates as soon as the types of all nodes in the unknown set are identified. The greedy Bayes algorithm corresponds to the 1-chain algorithm in

Kamble et al. (2018).

The asymptotic regret of the greedy Bayes algorithm is characterized in the following theorem.

Theorem 6.3.2. *For maximum value function w , the regret of greedy Bayes algorithm satisfies*

$$\lim_{T \rightarrow \infty} R_w^G(T) \sim \frac{1 - (1 - p)^2}{8p^2} \min\{p, 1 - p\}^4 n^2, \text{ for large } n.$$

The proof of Theorem 6.3.2 is provided in Section 6.4.2. Here we provide the proof sketch. The key observation in the proof is that the instantaneous regret in each round is a linear function of the number of matches of 1-1 type in the next round. We can show that the number of matches of 1-1 type evolves according to a time-inhomogeneous Markov chain. The regret analysis boils down to computing the expected values of the states of this Markov chain until a stopping time at which either the number of matches known to be of 0-0 type or the number of matches known to be either of 0-0 or 0-1 type is equal to zero. The analysis uses the Wormald's convergence theorem for bounding the expected value of some nonlinear functions of the Markov chain state.

Under maximum value function, matching nodes of different type is preferred and the difficulty of finding the optimal matching is in removing all pairs of nodes of same type whenever possible. If the nodes are of predominantly of one type, i.e. when p or $1 - p$ is small, a minority-type node is matched with a node of the same type with a small probability in the initial random matching. This means that the initial random matching would typically be close to an optimal matching and, hence, the regret of the algorithm is small. In the other extreme, when there is a balanced number of nodes of different types, i.e. when p is close to $1/2$, any node type is equally likely to be matched with a node of either type in the initial random matching. The initial random matching can thus be far from an optimal matching, and hence, the regret will be high.

Theorem 6.3.2 gives us a precise characterization of how the regret depends on the population parameter p . It indeed confirms the intuition that regret would be small when the population is predominantly of one type, and that the worst-case regret is achieved when there is a balanced number of node types. Observe that the dependency of the asymptotic regret on the population parameter p is not symmetric around $1/2$. It can be readily checked that for any $0 < \gamma < 1/2$, the asymptotic regret when $p = \gamma$ is larger than the asymptotic regret when $p = 1 -$




Figure 8: The asymptotic normalized regret for maximum value function, in the limit of large number of nodes n , versus population parameter p .

γ . In the limit of small p , the asymptotic regret is approximately $(1/4)p^3n^2$, while in the limit of small $1 - p$, the asymptotic regret is approximately $(1/8)(1 - p)^4n^2$. This shows that it is harder to find an optimal matching by greedy search when p is small (many low-type nodes) than when $1 - p$ is small (many high-type nodes).

Compared with the lower bound shown in Table 1, the regret upper bound has the extra factor $(1 - (1 - p)^2)/p^2 = (2 - p)/p$. This extra factor quantifies the regret due to the search cost of the greedy Bayes algorithm. This extra factor scales as $1/p$ for small value of p .

6.3.2 Minimum value function (AND)

In this section we consider the case of minimum value function defined as

$$w(x, y) = x \wedge y.$$

Under minimum value function, a greedy Bayes algorithm aims at searching for high-type nodes in order to match them together, as only in this case a positive reward is accrued.

The posterior distribution of node types can be represented by a collection of sets, each of which we refer to as a *matching set*, and the identities of nodes whose types are determined with certainty from observed pair values. Each matching set contains a set of matches such that there is *at most one node* of high type in this set. The greedy Bayes algorithm explores pairs of nodes that reside in different matching sets, and, as we will show, prioritizes the search to matching sets of smallest size.

For any two matching sets, evaluation of all pairs of nodes belonging to these

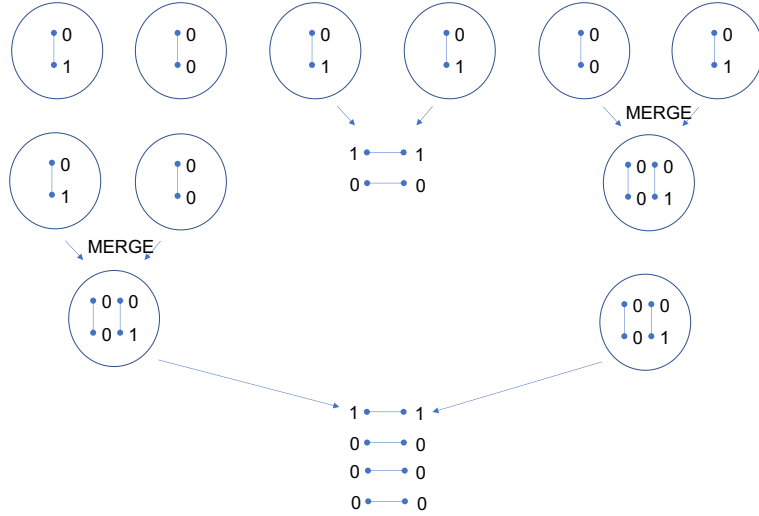


Figure 9: An illustration of evolution of matching sets of greedy Bayes matching strategy for AND value function.

two matching sets results in two possible outcomes: either (a) a 1-1 match is found, in which case the two matching sets are removed from further consideration, or (b) a 1-1 match is not found, in which case the two matching sets are merged into a new matching set. In the former case, each matching set has exactly one high-type node, and having found a match between the two high-type nodes, all other nodes in the two matching sets are identified to be of low-type with certainty. On the other hand, in the latter case, having not found a 1-1 match of nodes spanning the two matching sets, the only information deduced from observed pair evaluations is that there is at most one high-type node in the union of the two matching sets. Hence, in this case we merge the two matching sets into a new matching set.

This exploration procedure ends when matching sets coalesce into one set or no matching set remains. In the former case, the remaining matching set contains at most one high-type node, which cannot be matched with any other high-type node as all of them are already matched to a high-type node.

The following is a key lemma to show that greedy Bayes strategy explores matching sets of smallest sizes first.

Lemma 6.3.3. *For every round of greedy Bayes algorithm and every set of matches A of cardinality a for which the available information is that A contains at most one high-type*

node, the posterior probability of A containing a high-type node is

$$\pi(a) = \frac{pa}{1 - p + pa}.$$

The proof of the lemma is provided in Section 6.4.3. The proof is by induction over cardinality of matching sets.

Note that the posterior probability of a matching set containing a high-type node depends only on its cardinality. Intuitively, the larger the cardinality of a matching set, the larger the posterior probability of this matching set containing a high-type node. Note that any specific node in a matching set A of cardinality a is a high-type node with probability $\pi(a)/(2a) = p/(2(1 - p + pa))$, which is decreasing in the cardinality of A . This implies that the greedy Bayes strategy is to explore pairs of nodes residing in two different matching sets of smallest size.

We refer to greedy Bayes algorithm for minimum value function as LEAST-SIZE-MERGE algorithm, which is formally defined in Algorithm 7. This algorithm uses a procedure IS-MERGE defined in Algorithm 8. This procedure uses as input two matching sets and constructs matchings of nodes by exploring pairs of nodes residing in the two matching sets.

Algorithm 7: Least-Size-Merge

Input : nodes set N .
Initialization: iteration $t = 1$, $M_1 \leftarrow$ a random matching on nodes N and
 $S \leftarrow \{\{(i, j)\} \in M_1 : Y_1(i, j) = 0\}$.

```

1 while  $|S| > 1$  do
2    $A \leftarrow \arg \min_{S' \in S} |S'|$ ;
3    $B \leftarrow \arg \min_{S' \in S \setminus \{A\}} |S'|$ ;
4    $S \leftarrow S \setminus \{A, B\}$ ;
5   if IS-MERGE( $A, B$ ) == True then
6      $S \leftarrow S \cup \{A \cup B\}$ 
7   end
8 end

```

From the above discussion, it is clear that greedy Bayes algorithm results in a stochastic process that describes evolution of a collection of matching sets as pairs of matching sets are examined, which results in either merging two matching sets or removing them from the collection of matching sets.

The initial matching is a random matching. From this initial matching, the algorithm identifies high-type nodes that belong to matches yielding value 1, and a set of matching sets, each consisting of a pair of nodes for which the evaluation is observed to be of value 0. Indeed, all pairs in these matching sets are either of 0-0 or 0-1 type. Let X_1 be the number of matching sets at step 1. It is readily observed that

$$X_1 \sim \text{Bin} \left(\frac{n}{2}, 1 - p^2 \right).$$

We divide rounds into epochs. In each epoch r , let S_r denote the set of matching sets at the beginning of epoch r . Initially, as we discussed, S_1 consists of pairs of type 0-0 or 0-1 defined by the initial random matching and the cardinality of this set is X_1 . At epoch r , a pair of matching sets A and B of the smallest size is removed from S_r and pairs of nodes spanning A and B are evaluated until either a match 1-1 is found, or not such a match is found. In the latter case, we deduce that nodes in $A \cup B$ contain at most one high-type node. The set of matching sets in the next epoch, S_{r+1} , consists of all unexplored matching sets from S_r , and $A \cup B$ if the outcome of examining A and B resulted in not finding a 1-1 match.

In each epoch, the posterior distribution of node types is fully defined by the identities of nodes inferred to be of high-type, and the set of matching sets. The posterior probability of a matching set containing a high-type node is given in Lemma 6.3.3.

It is readily observed that at the beginning of each epoch, all matching sets are of the same cardinality except for at most one set of larger cardinality. We call this latter matching set a *special matching set*, if one exists, and all other as *regular matching sets*.

As the greedy strategy prioritizes evaluating nodes residing in smallest matching sets, we can divide epochs in super-epochs. In each super-epoch, disjoint pairs of regular matching sets are created and are then examined in arbitrary order, and if there is a special matching set, it is paired at the end of the super-epoch with a regular set for examination. For super-epoch s , let X_s denote the number of regular matching sets and let Y_s denote the cardinality of the special matching set; if there is no special matching set, we define $Y_s = 0$.

Stochastic process $\{(X_s, Y_s)\}_{s \geq 0}$ is a Markov chain with initial state $(X_0, Y_0) = (n/2, 0)$, and for each super-epoch $s \geq 0$,

$$X_{s+1} \mid (X_s, Y_s) \sim \text{Bin} \left(\left\lfloor \frac{X_s}{2} \right\rfloor, 1 - \pi(2^s)^2 \right) \quad (6.3.1)$$

Algorithm 8: IS-MERGE

Input : two matching sets A, B .
Initialization: create a candidate set $P_i \leftarrow B$ for all $i \in A$

```

1 for all node  $u \in A$  do
2   while  $P_u \neq \emptyset$  do
3      $t \leftarrow t + 1$ ;
4     choose an node  $v \in P_u$  and let  $(u, u'), (v, v') \in M_{t-1}$ ;
5      $M_t \leftarrow (M_{t-1} \setminus \{(u, u'), (v, v')\}) \cup \{(u, v), (u', v')\}$ ;
6     if  $Y_t(u, v) + Y_t(u', v') = 1$  then
7       | return False;
8     end
9      $P_u \leftarrow P_u \setminus \{v\}$ ;
10    if  $u' \in A$  then
11      |  $P_{u'} \leftarrow P_{u'} \setminus \{v'\}$ ;
12    end
13    if  $v' \in A$  then
14      |  $P_{v'} \leftarrow P_{v'} \setminus \{u'\}$ ;
15    end
16  end
17  return True;
18 end

```

and

$$Y_{s+1} \mid (X_s, Y_s, X_{s+1}) \sim \begin{cases} (Y_s + \xi_s) \text{Ber}(1 - \pi(\xi_s)\pi(Y_s)) & \text{if } Y_s > 0 \\ (2^s + 2^{s+1}) \text{Ber}(1 - \pi(2^s)\pi(2^{s+1})) & \text{if } Y_s = 0, X_s \text{ is odd, } X_{s+1} > 0 \\ 0, & \text{otherwise} \end{cases} \quad (6.3.2)$$

where

$$\xi_s = \begin{cases} 2^s & \text{if } X_s \text{ is odd} \\ 2^{s+1} & \text{if } X_s \text{ is even, and } X_{s+1} > 0 \\ 0 & \text{otherwise} \end{cases} .$$

The transition probabilities of the Markov chain $\{(X_s, Y_s)\}_{s \geq 0}$ can be interpreted as follows. In each super-epoch, matchings are first created by swapping matches for nodes belonging to different regular matchings sets, by examining one pair of regular matching sets after other. There are $\lfloor X_s/2 \rfloor$ pairs of regular matching sets that are examined in super-epoch s . Each such examination of a pair of regular matching sets A and B results in either A and B being merged, which occurs with probability $1 - \pi(|A|)\pi(|B|)$, or A and B being removed from further consideration which occurs with probability $\pi(|A|)\pi(|B|)$. Observe that




Figure 10: Two possible outcomes: merge or remove two matching sets.

$\pi(|A|)\pi(|B|)$ is the probability that both A and B contain a high-type node. The outcomes of this procedure are illustrated in Figure 10. All regular matching sets in super-epoch s are of cardinality 2^s . Hence, in super-epoch s , two regular matching sets are merged with probability $1 - \pi(2^s)^2$, and are removed from further consideration with probability $\pi(2^s)^2$. This explains the conditional distribution of the number of regular matching sets in super-epoch $s + 1$, X_{s+1} , given in (6.3.1).

We next explain the conditional distribution of the cardinality of the special set in super-epoch $s + 1$, Y_{s+1} , conditional on the values of X_s, Y_s and X_{s+1} given by (6.3.2). There are three possible cases. First, if $Y_s > 0$, then at the beginning of super-epoch s there is a special matching set of cardinality Y_s . This set is examined with a matching set A of cardinality ζ_s at the end of super-epoch s after regular matching sets are examined, if there is at least one matching set left for consideration besides the special matching set. The set A is either (a) one of regular matching sets of cardinality 2^s if at the beginning of super-epoch s the number of regular matching sets is odd, or (b) a matching set of cardinality 2^{s+1} that resulted from merging two regular matching sets in super-epoch s if the number of regular matching sets at the beginning of super-epoch s is even and at least one merging resulted from examinations of regular matching sets. Second, if $Y_s = 0$, then all matching sets at the beginning of super-epoch s are regular matching sets. In this case, a special matching set can exist in super-epoch $s + 1$, only if the number of regular matching sets at the beginning of super-epoch s is odd and at least one merging resulted from examination of regular matching sets in super-epoch s . The cardinality Y_{s+1} of this special set is clearly $2^s + 2^{s+1}$. Finally, if $Y_s = 0$ and either number of regular matching sets at the beginning of super-epoch s is even or no examination of regular matching sets in super-epoch s resulted in merging,

there is no special matching set in super-epoch $s + 1$. Note that cardinality of the special matching set in super-epoch s is of cardinality $Y_1 = 0$ for $s = 1$ and $Y_s \leq 2^1 + \dots + 2^s \leq 2^{s+1}$ for $s > 1$.

The Markov chain $\{(X_s, Y_s)\}_{s \geq 0}$ has absorbing states $(0, 0)$, $(1, 0)$, $(0, z)$, for $0 < z \leq n$. In state $(0, 0)$, all high-type nodes are identified. In states $(1, 0)$ or $(0, z)$, there is at most one unidentified high-type node. It is not needed to identify the type of this node, as all other high-type nodes are already matched with high-type nodes.

The number of rounds for examination of a pair of matching sets is discussed as follows. For any two matching sets A and B , with cardinalities $|A| = a$ and $|B| = b$, let $N(a, b)$ be the smallest number of adjacent matchings that cover all pairs $(u, v) \in A \times B$. It can be readily noted that $ab/2 \leq N(a, b) \leq ab$. When exploring pairs in matching sets A and B such that either A or B contains only low-type nodes, the number of exploration steps is equal to $N(a, b)$. On the other hand, when exploring pairs in matching sets A and B such that both A and B contain a high-type node, then the number of exploration steps is a random variable with expected value $N(a, b)/2$.

The number of steps until LEAST-SIZE-MERGE algorithm terminates is characterised in the following proposition.

Proposition 6.3.4. *The expected number of steps τ until termination of algorithm LEAST-SIZE-MERGE is bounded as follows*

$$\mathbb{E}[\tau] = O\left(\min\left\{\frac{n}{p}, n^2\right\}\right).$$

Proposition 6.3.4 characterizes the expected number of steps used by LEAST-SIZE-MERGE algorithm, which leads to an upper bound on the regret of LEAST-SIZE-MERGE algorithm. The details of the proof of Proposition 6.3.4 is in Section 6.4.4.

Theorem 6.3.5. *For the minimum value function w , the regret of greedy Bayes algorithm LEAST-SIZE-MERGE satisfies*

$$\lim_{T \rightarrow \infty} R_w^{\mathcal{G}}(T) \leq (a_{n,p} + b_{n,p} + c_{n,p}) n^2 + O(n \log(n))$$




Figure 11: The number of rounds until greedy Bayes incremental matching algorithm terminates for minimum (Boolean AND) value function versus population parameter p , for different numbers of nodes n .

where

$$a_{n,p} = \frac{1}{4} \sum_{s=1}^{\log_2(n)} \left(1 - \frac{1}{2}p_s\right) \left(1 - \frac{1}{2}p_s^2\right) p_s \left(\prod_{i=0}^{s-1} q_i^2 + \frac{1}{n} \sum_{i=0}^{s-1} 2^{i+1} \left(\prod_{j=0}^i q_j \right) \left(\prod_{j=i+1}^{s-1} q_j^2 \right) \right),$$

$$b_{n,p} = \frac{1}{n} \sum_{s=1}^{\log_2(n)} 2^{s+1} p_s \prod_{i=0}^{s-1} q_i,$$

and

$$c_{n,p} = \min \left\{ \frac{c}{np}, \frac{4}{3} \right\}$$

with $p_s = p2^s / (1 - p + p2^s)$, $q_s = 1 - p_s^2$, and c is a constant such that $c > 16/(\sqrt{2} - 1)$. Moreover, for any $n \geq 1, p \in (0, 1]$, there exists some constant $M > 0$ such that $\max\{a_{n,p}, b_{n,p}, c_{n,p}\} < M$.

Proof of the theorem is provided in Section 6.4.5. The proof separately considers contributions to regret over rounds at which regular matching sets are exam-

ined and contributions to regret over rounds at which a special set is examined. In the regret bound, $a_{n,p}$ is due to regret accumulated while examining regular matching sets while $b_{n,p}$ and $c_{n,p}$ are due to regret accumulated while examining a special matching set. Theorem 6.3.5 shows that the asymptotic regret bound of greedy Bayes algorithm scales LEAST-SIZE-MERGE in $O(n^2)$.

Here we give more detailed characterizations of $a_{n,p}$ and $b_{n,p}$ when the population parameter p depends on the size of population n . First, when $p = 0$ or $p = 1$, any algorithm algorithm has a zero regret. On the one hand, when the high-type nodes are scarce, i.e. $p \in (0, 1/2]$, it holds that $a_{n,p} = O(1)$ and $b_{n,p} = O(1)$. The search cost of greedy Bayes algorithm LEAST-SIZE-MERGE remains limited and independent of the scarcity of the high-type nodes. On the other hand, when the low-type nodes are scarce, i.e. $p \in (1/2, 1]$, it holds that $a_{n,p} = O((1-p)^2)$ and $b_{n,p} = O((1-p)1/n)$. The search cost of greedy Bayes algorithm LEAST-SIZE-MERGE increases quadratically in the scarcity of the low-type nodes. The detailed explicit bounds of $a_{n,p}$ and $b_{n,p}$ for various scenarios are provided in Section 6.4.6.

In Figure 12, we show the theoretical regret bound in Theorem 6.3.5 along with the regret estimated through simulations versus the value of population parameter p for several different values of the number of nodes n . We observe that the theoretical bound is tight for large values of n . Observe that the regret decreases with the value of the population parameter p for any sufficiently large value of this parameter. This is intuitive as for pair values according to the MIN function, the goal is to maximize the number of 1-1 type matches, which easier for greedy search when there is a larger number of high-type nodes.

6.4 Proofs

6.4.1 Proof of Lemma 6.3.1

Let w be an arbitrary value function and \mathcal{A} be an arbitrary local switching adaptive matching algorithm for value function w . Algorithm \mathcal{A} reduces the instantaneous regret in each round for at most some constant value $L > 0$. Recall that any adaptive matching algorithm selects random matching for initial matching because the prior distribution is non-informative for maximum value matching.

Let $R_t^{\mathcal{A}}$ denote the instantaneous regret in round t , for $t \geq 1$. Note that it holds

$$R_t^{\mathcal{A}} \geq \max\{R_1 - L(t-1), 0\}, \text{ for all } t \geq 1$$




Figure 12: Regret of the greedy Bayes algorithm for the MIN value function versus the population parameter p , for different values of the number of nodes n .

where R_1 denotes the instantaneous regret of initial random matching (we omit the subscript \mathcal{A} in R_1 because this instantaneous regret does not depend on \mathcal{A}).

Hence, for any $T \geq \lfloor \mathbb{E}[R_1]/L \rfloor$,

$$\begin{aligned}
\mathbb{E} \left[\sum_{t=1}^T R_t^{\mathcal{A}} \right] &\geq \sum_{t=1}^T \mathbb{E}[\max\{R_1 - L(t-1), 0\}] \\
&\geq \sum_{t=1}^T \max\{\mathbb{E}[R_1] - L(t-1), 0\} \\
&\geq \sum_{t=0}^{\lfloor \mathbb{E}[R_1]/L \rfloor - 1} (\mathbb{E}[R_1] - Lt) \\
&= \mathbb{E}[R_1] \left\lfloor \frac{\mathbb{E}[R_1]}{L} \right\rfloor - L \frac{1}{2} \left(\left\lfloor \frac{\mathbb{E}[R_1]}{L} \right\rfloor - 1 \right) \left\lfloor \frac{\mathbb{E}[R_1]}{L} \right\rfloor \\
&\geq \frac{L}{2} \left\lfloor \frac{\mathbb{E}[R_1]}{L} \right\rfloor^2
\end{aligned}$$

where the second inequality is by Jensen's inequality.

For any fixed $c \in (0, 1/(2L))$, assume that $\mathbb{E}[R_1] \geq L/(1 - \sqrt{2cL})$. Then, $\mathbb{E}[R_1]/L - 1 \geq \sqrt{2cL}(\mathbb{E}[R_1]/L)$ and, thus, it follows

$$\frac{L}{2} \left[\frac{\mathbb{E}[R_1]}{L} \right]^2 \geq \frac{L}{2} \left(\frac{\mathbb{E}[R_1]}{L} - 1 \right)^2 \geq c\mathbb{E}[R_1]^2.$$

Regret for random matching Let M be a random matching and let N^{00} , N^{01} and N^{11} denote the number of matches in M of type 0-0, 0-1 and 1-1, respectively. Let w be an arbitrary value function. Then, we have

$$\mathbb{E}[\mu_w(M)] = w(0,0)\mathbb{E}[N^{00}] + w(0,1)\mathbb{E}[N^{01}] + w(1,1)\mathbb{E}[N^{11}].$$

It follows that

$$\mathbb{E}[\mu_w(M)] = \frac{1}{2} \left(w(0,0)(1-p)^2 + w(0,1)2p(1-p) + w(1,1)p^2 \right) n.$$

In the following we consider the instantaneous regret of random matching for different value functions w .

Maximum value function In this case, we have

$$\mathbb{E}[\mu_w(M)] = \frac{1}{2}(1 - (1-p)^2)n.$$

Any optimal matching M^* has a maximum possible number of 0-1 type pairs and any remaining high-type nodes matched with high-type nodes. Let N^1 denote the number of high-type nodes in the given set of nodes. Then, conditional on $\{N_1 = n_1\}$, we have

$$\mathbb{E}[\mu_w(M^*) \mid N_1 = n_1] = \frac{1}{2} \min \{2n_1, n\}.$$

This is because if $n_1 \leq n/2$, then every high-type node is matched with a low-type node yielding the total value of n_1 , and, otherwise, if $n_1 > n/2$, then $n - n_1$ high-type nodes are matched with low-type nodes and the remaining $n_1 - (n - n_1)$ high-type nodes are matched with high-type nodes, yielding the total value $(n - n_1) + (1/2)(n_1 - (n - n_1)) = n/2$.

It follows that the expected weight of an optimal matching is

$$\mathbb{E}[\mu_w(M^*)] = \frac{1}{2}\mathbb{E}[\min\{2N_1, n\}]$$

Hence, we have

$$\tilde{R}_w = \mathbb{E}[\mu_w(M^*)] - \mathbb{E}[\mu_w(M)] = \mathbb{E}\left[\min\left\{\frac{N_1}{n} - p + \frac{1}{2}p^2, \frac{1}{2}(1-p)^2\right\}\right]n.$$

Since, for any real number a , $x \mapsto \min\{x, a\}$ is a concave function, by Jensen's inequality and $\mathbb{E}[N_1/n] - p = 0$, we have

$$\tilde{R}_w \leq \frac{1}{2}\min\{p, 1-p\}^2n.$$

We next show that it holds

$$\tilde{R}_w \geq \frac{1}{2}\min\{p, 1-p\}^2n(1-o(1))$$

provided that

$$\frac{p}{\log(1/p)^{1/4}} = \omega\left(\frac{1}{(c_n^2n)^{1/4}}\right),$$

and

$$1-p = \omega\left(e^{-\frac{1}{26}c_n^2n}\right)$$

where c_n is any sequence such that $c_n = o(1)$. Hence, under these conditions,

$$\tilde{R}_w \sim \frac{1}{2}\min\{p, 1-p\}^2n, \text{ for large } n.$$

By Hoeffding's inequality,

$$\Pr\left[\frac{N_1}{n} - p < -c_n\frac{1}{2}p^2\right] \leq e^{-\frac{1}{2}c_n^2p^4n}.$$

It follows

$$\begin{aligned} \tilde{R}_w &= \mathbb{E}\left[\min\left\{\frac{N_1}{n} - p + \frac{1}{2}p^2, \frac{1}{2}(1-p)^2\right\}\right]n \\ &\geq \min\left\{(1-c_n)\frac{1}{2}p^2, \frac{1}{2}(1-p)^2\right\}n\left(1 - e^{-\frac{1}{2}c_n^2p^4n}\right) - pe^{-\frac{1}{2}c_n^2p^4n} \\ &\geq \frac{1}{2}\min\{p, 1-p\}^2n\left((1-c_n)\left(1 - e^{-\frac{1}{2}c_n^2p^4n}\right) - 2\max\left\{\frac{1}{p}, \frac{p}{(1-p)^2}\right\}e^{-\frac{1}{2}c_n^2p^4n}\right). \end{aligned}$$

Under conditions $c_n = o(1)$ and $p/\log(1/p)^{1/4} = \omega(1/(c_n^2 n)^{1/4})$, it clearly holds

$$(1 - c_n) \left(1 - e^{-\frac{1}{2}c_n^2 p^4 n}\right) \geq 1 - o(1).$$

We next show that

$$2 \max \left\{ \frac{1}{p}, \frac{p}{(1-p)^2} \right\} e^{-\frac{1}{2}c_n^2 p^4 n} = o(1).$$

For the case when $p \leq 1/2$, we have

$$\begin{aligned} 2 \max \left\{ \frac{1}{p}, \frac{p}{(1-p)^2} \right\} e^{-\frac{1}{2}c_n^2 p^4 n} &= 2 \frac{1}{p} e^{-\frac{1}{2}c_n^2 p^4 n} \\ &= 2e^{-\frac{1}{2}c_n^2 p^4 n(1 - \log(1/p)/((1/2)c_n^2 p^4 n))} \\ &= 2e^{-\frac{1}{2}c_n^2 p^4 n(1 - o(1))} \\ &= o(1) \end{aligned}$$

where the two last relations hold by the condition $p/\log(1/p)^{1/4} = \omega(1/(c_n^2 n)^{1/4})$.

For the case when $p > 1/2$,

$$\begin{aligned} 2 \max \left\{ \frac{1}{p}, \frac{p}{(1-p)^2} \right\} e^{-\frac{1}{2}c_n^2 p^4 n} &= 2 \frac{p}{(1-p)^2} e^{-\frac{1}{2}c_n^2 p^4 n} \\ &\leq 2e^{-\frac{1}{25}c_n^2 n + 2\log\left(\frac{1}{1-p}\right)}. \end{aligned}$$

The right-hand side in the last inequality is $o(1)$ if, and only if,

$$\frac{1}{25}c_n^2 n + \log\left(\frac{1}{1-p}\right) \rightarrow \infty \text{ as } n \rightarrow \infty,$$

i.e.

$$1 - p = \omega\left(e^{-\frac{1}{25}c_n^2 n}\right).$$

Minimum value function In this case, we have

$$\mathbb{E}[\mu_w(M)] = \frac{1}{2}p^2 n.$$

Any optimal matching has maximum number of 1-1 type pairs. Hence, we

have

$$\mathbb{E}[\mu_w(M^*)] = \mathbb{E} \left[\left\lfloor \frac{N_1}{2} \right\rfloor \right].$$

We next show that the regret of random matching satisfies

$$\tilde{R}_w = \mathbb{E}[\mu_w(M^*)] - \mathbb{E}[\mu_w(M)] = \frac{1}{2}p(1-p)n(1-o(1))$$

provided that $p = \omega(1/n)$ and $1-p = \omega(1/n)$.

The regret of random matching has the following upper bound,

$$\tilde{R}_w = \mathbb{E}[\mu_w(M^*)] - \mathbb{E}[\mu_w(M)] \leq \frac{1}{2}pn - \frac{1}{2}p^2n = \frac{1}{2}p(1-p)n$$

and the following lower bound,

$$\tilde{R}_w = \mathbb{E}[\mu_w(M^*)] - \mathbb{E}[\mu_w(M)] \geq \frac{1}{2}pn - 1 - \frac{1}{2}p^2n = \frac{1}{2}p(1-p)n - 1 = \frac{1}{2}p(1-p)n(1-o(1))$$

where the last relation holds under conditions $p = \omega(1/n)$ and $1-p = \omega(1/n)$.

6.4.2 Proof of Theorem 6.3.2

For each round $t \geq 0$, let S_t be the number of matches known to be of 0-0 type, U_t be the number of matches known to be of either 0-1 or 1-1 type, and K_t be the number of matches known to be of 0-1 type.

Let N_1 be the number of type 1 items in round 1. Let U_t^{01} be the number of matches of 0-1 type and $U_t^{11} := U_t - U_t^{01}$. The following equations hold, for every $t \geq 0$,

$$U_t^{01} + U_t^{11} = U_t \tag{6.4.1}$$

$$S_t + U_t + K_t = \frac{n}{2} \tag{6.4.2}$$

$$U_t^{11} + U_t + K_t = N_1. \tag{6.4.3}$$

By the law of large numbers, the following limits hold almost surely,

$$\lim_{n \rightarrow \infty} \frac{1}{n} N_1 = p \quad (6.4.4)$$

$$\lim_{n \rightarrow \infty} \frac{2}{n} U_1^{11} = p^2 \quad (6.4.5)$$

$$\lim_{n \rightarrow \infty} \frac{2}{n} S_1 = (1 - p)^2 \quad (6.4.6)$$

$$\lim_{n \rightarrow \infty} \frac{2}{n} U_1 = 1 - (1 - p)^2. \quad (6.4.7)$$

For any $p \in (0, 1)$, $S_1 \wedge U_1 > 0$ almost surely. Hence, in what follows we condition on $\{S_1 > 0\} \cap \{U_1 > 0\}$.

The system dynamics is according to the Markov chain $\{X_t\}_{t \geq 0}$ where $X_t = (S_t, U_t^{01}, U_t^{11}, K_t)$ with transition probabilities given as follows

$$S_{t+1} = S_t - \xi_t \mathbf{1}_{\{S_t \wedge U_t > 0\}} \quad (6.4.8)$$

$$U_{t+1}^{11} = U_t^{11} - \xi_t \mathbf{1}_{\{S_t \wedge U_t > 0\}} \quad (6.4.9)$$

$$U_{t+1}^{01} = U_t^{01} - (1 - \xi_t) \mathbf{1}_{\{S_t \wedge U_t > 0\}} \quad (6.4.10)$$

$$K_{t+1} = K_t + (1 + \xi_t) \mathbf{1}_{\{S_t \wedge U_t > 0\}} \quad (6.4.11)$$

where ξ_t is a Bernoulli random variable with mean U_t^{11}/U_t under condition $U_t > 0$.

Conditional on $\{N_1 = n_1\}$, the expected reward of an optimal matching is

$$\mu^* = \frac{n_1 + \min\{n_1, n - n_1\}}{2}.$$

The instantaneous reward at round t is

$$\mu_t = U_t + K_t.$$

Hence, the instantaneous regret R_t at round t is equal to

$$R_t = U_t^{11} + \min\left\{\frac{n}{2} - N_1, 0\right\}. \quad (6.4.12)$$

The algorithm terminates at the stopping time τ defined as

$$\tau = \min\{t > 1 \mid S_t = 0 \text{ or } U_t = 0\}.$$

Note that τ satisfies, with probability 1,

$$S_1 \wedge U_1 + 1 \leq \tau \leq U_1 + 1.$$

To see this, note that from (6.4.8), $S_t \geq S_1 - (t - 1)$. Hence, $S_t = 0$ implies $t \geq S_1 + 1$. Similarly, from (6.4.9) and (6.4.10), $U_t \geq U_1 - (t - 1)$, hence $U_t = 0$ implies $t \geq U_1 + 1$. From this, it follows that $\tau \geq S_1 \wedge U_1 + 1$. Now, for the upper bound, suppose first that $S_\tau \geq U_\tau = 0$. Since in this case $U_\tau = U_1 - (\tau - 1)$, it follows that $\tau = U_1 + 1$. Otherwise, if $0 = S_\tau < U_\tau$, it follows that $\tau \leq U_1 + 1$, which establishes the upper bound.

From (6.4.1), (6.4.9) and (6.4.10), for all $t \geq 1$, $U_{t+1} = U_t - \mathbf{1}_{\{S_t \wedge U_t > 0\}}$. Hence, for all $t \geq 1$,

$$U_t = U_1 - (t - 1)\mathbf{1}_{\{S_t \wedge U_t > 0\}} - \tau\mathbf{1}_{\{S_t \wedge U_t = 0\}}.$$

From (6.4.8) and (6.4.9), for all $t \geq 1$, $S_t - U_t^{11} = S_1 - U_1^{11}$. Hence, $S_t > 0$ is equivalent to $U_t^{11} > U_1^{11} - S_1$.

It follows that the system evolves according to a one-dimensional time-inhomogenous Markov chain $\{U_t^{11}\}_{t \geq 1}$ with transition probabilities given as, for $t \leq u_1$,

$$\mathbb{P}_1[U_{t+1}^{11} = u - 1 \mid U_t^{11} = u] = \frac{u}{u_1 - (t - 1)} \mathbf{1}_{\{u > u_1^{11} - s_1\}}$$

and

$$\mathbb{P}_1[U_{t+1}^{11} = u \mid U_t^{11} = u] = 1 - \mathbb{P}_1[U_{t+1}^{11} = u - 1 \mid U_t^{11} = u]$$

and, for $t > u_1$,

$$\mathbb{P}_1[U_{t+1}^{11} = u \mid U_t^{11} = u] = 1.$$

For all $1 \leq t \leq u_1$, we have

$$\mathbb{E}_1[U_{t+1}^{11}] = \mathbb{E}_1[U_t^{11}] - \mathbb{E}_1 \left[\frac{U_t^{11}}{U_t} \mathbf{1}_{\{U_t^{11} > u_1^{11} - s_1\}} \right].$$

and, for $t > u_1$,

$$\mathbb{E}_1[U_{t+1}^{11}] = \mathbb{E}_1[U_t^{11}].$$

In the following we consider two cases with respect to the sign of $p - 1/2$.

Case 1: $p \leq 1/2$. From (6.4.4), (6.4.5) and (6.4.4), it holds, almost surely

$$U_1^{11} \leq S_1, \quad N_1 \leq \frac{n}{2}.$$

By (6.4.12), it holds, almost surely,

$$R_t = U_t^{11}.$$

In this case, we have

$$\mathbb{E}_1 \left[\frac{U_t^{11}}{U_t} \mathbf{1}_{\{U_t^{11} > u_1^{11} - s_1\}} \right] = \mathbb{E}_1 \left[\frac{U_t^{11}}{U_t} \right] = \frac{1}{u_t} \mathbb{E}_1[U_t^{11}]$$

and $u_t = u_1 - (t - 1)$. Hence, for $1 \leq t \leq u_1$, $\mathbb{E}_1[U_{t+1}^{11}] = \frac{u_{t+1}}{u_t} \mathbb{E}_1[U_t^{11}]$, and, thus, expanding this recurrence, we obtain

$$\mathbb{E}_1[U_t^{11}] = \frac{u_t}{u_1} u_1^{11}.$$

It follows that for all $t \geq 1$,

$$\mathbb{E}_1[U_t^{11}] = \left(1 - \frac{t-1}{u_1}\right)^+ u_1^{11} \tag{6.4.13}$$

where $(\cdot)^+ = \max\{\cdot, 0\}$.

We have, for $t \geq 1$,

$$\mathbb{E}_1[R_t] = \mathbb{E}_1[U_t^{11}] = \left(1 - \frac{t-1}{u_1}\right)^+ u_1^{11}.$$

Hence, we have

$$\begin{aligned}
\mathbb{E} \left[\sum_{t=1}^{U_1+1} R_t \right] &= \mathbb{E} \left[\sum_{t=1}^{U_1+1} \left(1 - \frac{t-1}{U_1} \right)^+ U_1^{11} \right] \\
&= \frac{1}{2} \mathbb{E} \left[U_1^{11} (U_1 + 1) \right] \\
&= \frac{1}{2} \mathbb{E}[U_1^{11} U_1] + \frac{1}{2} \mathbb{E}[U_1^{11}] \\
&= \frac{1}{2} \mathbb{E}[\mathbb{E}[U_1^{11} | U_1] U_1] + \frac{1}{2} \mathbb{E}[U_1^{11}] \\
&= \frac{1}{2} \frac{p^2}{1 - (1-p)^2} \mathbb{E}[U_1^2] + \frac{1}{4} p^2 n \\
&= \frac{1}{2} \frac{p^2}{1 - (1-p)^2} \left((1 - (1-p)^2)^2 \frac{n^2}{4} + \frac{n}{2} (1 - (1-p)^2) (1-p)^2 \right) + \frac{1}{4} p^2 n \\
&= \frac{1}{8} (1 - (1-p)^2) p^2 n^2 + \frac{1}{4} p^2 (1-p)^2 n + \frac{1}{4} p^2 n \\
&= \frac{1}{8} (1 - (1-p)^2) p^2 n^2 + \frac{1}{4} p^2 (1 + (1-p)^2) n \\
&= \frac{1}{8} (1 - (1-p)^2) p^2 n^2 (1 + o(1))
\end{aligned}$$

where the last relation holds under condition $p = \omega(1/n)$.

Case 2: $p > 1/2$ From (6.4.4), (6.4.5) and (6.4.4), we have, almost surely,

$$U_1^{11} > S_1, \quad N_1 > \frac{n}{2}.$$

By (6.4.12)

$$R_t = U_t^{11} + \frac{n}{2} - N_1$$

We next establish a limit value of U_t^{11} as n goes to infinity, for $p > 1/2$ by applying Theorem 2.3.2.

Condition (C1) holds because for all $t \geq 1$, it always holds,

$$|U_{t+1}^{11} - U_t^{11}| \leq 1.$$

Condition (C2) holds as follows

$$\mathbb{E}[U_{t+1}^{11} - U_t^{11} | \mathcal{F}_t] = f(t/n, U_t^{11}/n) \mathbf{1}_{\{U_t^{11}/n > \bar{z}\}}$$

where $U_1^{11} - S_1 = \tilde{z}n + o(n)$, and

$$f(t, x) = -\frac{x}{\tilde{u}_1 - t'}$$

for $x > \tilde{z}$, where $u_1 = \tilde{u}_1 n + o(n)$ and $\tilde{u}_1 = \frac{1}{2}p(2-p)$.

Hence, when $p \in (1/2, 1)$ is fixed,

$$\mathbb{E}[U_{t+1}^{11} - U_t^{11} \mid \mathcal{F}_t] = -\frac{U_t^{11}/n}{\frac{1}{2}p(2-p) - t/n} + o(1),$$

and, when $p = 1 - o(1)$,

$$\mathbb{E}[U_{t+1}^{11} - U_t^{11} \mid \mathcal{F}_t] = -\frac{U_t^{11}/n}{\frac{1}{2} - t/n} + o(1).$$

For condition (C3), we need to show that f satisfies a Lipschitz condition on $D = \{(t, x) : 0 \leq t < t^*, 0 \leq x \leq 1\}$. Note that

$$\begin{aligned} |f(t, x) - f(t', x')| &= |f(t, x) - f(t', x) + f(t', x) - f(t', x')| \\ &\leq \frac{|f(t, x) - f(t', x)|}{|t - t'|} |t - t'| + \frac{|f(t', x) - f(t', x')|}{|x - x'|} |x - x'| \\ &\leq L(|t - t'| + |x - x'|) \end{aligned}$$

where

$$L = \max_{(t,x) \in D} \left| \frac{\partial}{\partial t} f(t, x) \right| \vee \max_{(t,x) \in D} \left| \frac{\partial}{\partial x} f(t, x) \right|.$$

Since $\partial f(x, t)/\partial t = x/(\tilde{u}_1 - t)^2$ and $\partial f(x, t)/\partial x = -1/(\tilde{u}_1 - t)$, we have

$$L = \frac{\tilde{u}_1^{11}}{(\tilde{u}_1 - t^*)^2} \vee \frac{1}{\tilde{u}_1 - t^*}$$

where t^* is the unique positive value such that $u^{11}(t^*) = \tilde{z}$. Note that

$$t^* = \left(1 - \frac{\tilde{z}}{\tilde{u}_1^{11}}\right) \tilde{u}_1 = \frac{(1 - (1-p)^2)(1-p)^2}{2p^2}$$

and

$$\tilde{u}_1 - t^* = \frac{1 - (1-p)^2}{p^2} \left(p - \frac{1}{2}\right).$$

Next, consider the differential equation,

$$\frac{d}{dt}u^{11}(t) = f(t, u^{11}(t)), \text{ for } 0 \leq t < t^*.$$

This differential equation has a unique solution which is given by

$$u^{11}(t) = \left(1 - \frac{t}{\tilde{u}_1}\right) \tilde{u}_1^{11}$$

where $u_1^{11} = \tilde{u}_1^{11}n + o(n)$.

By Theorem 2.3.2, we have, almost surely, for every $1 \leq t < t^*n$,

$$U_t^{11} = \left(1 - \frac{t}{u_1}\right) u_1^{11} + o(n).$$

and $o(n) \leq 2n^{4/15} \log(n)$.

For the regret, we have

$$\begin{aligned} \sum_{t=1}^{t^*n} R_t &= \sum_{t=1}^{t^*n} \left(1 - \frac{t}{U_1}\right) U_1^{11} + t^*n \left(\frac{n}{2} - N_1\right) + t^*o(n^2) \\ &= t^*n \left(U_1^{11} - N_1 + \frac{n}{2}\right) - \frac{1}{2}(t^*n)^2 \frac{U_1^{11}}{U_1} + t^*o(n^2) \\ &= \frac{1}{2}(1-p)^2 t^*n^2 - \frac{1}{2} \frac{p^2}{1 - (1-p)^2} (t^*)^2 n^2 + t^*o(n^2) \\ &= \frac{(1 - (1-p)^2)(1-p)^4}{8p^2} n^2 + \frac{(2-p)(1-p)^2}{2p} o(n^2). \end{aligned}$$

6.4.3 Proof of Lemma 6.3.3

We need to show that for any set A , the posterior probability that A contains a type 1 node is equal to

$$\pi(|A|) = \frac{|A|p}{1-p+|A|p}. \quad (6.4.14)$$

We will show this by induction over the cardinality of sets. We first consider the base case for sets of cardinality 2. It suffices to consider (X_1, X_2) that are two independent Bernoulli (p) random variables, conditional on the event $(X_1, X_2) \neq$

(1, 1). Then,

$$\begin{aligned}
\mathbb{P}[(X_1, X_2) \neq (0, 0) \mid (X_1, X_2) \neq (1, 1)] &= \frac{\mathbb{P}[(X_1, X_2) \in \{(0, 1), (1, 0)\}]}{\mathbb{P}[(X_1, X_2) \neq (1, 1)]} \\
&= \frac{2p(1-p)}{1-p^2} \\
&= \frac{2p}{1+p}.
\end{aligned}$$

This shows that (6.4.14) holds for the base case.

For the induction step, suppose that (6.4.14) holds for sets of cardinality $\leq k$. We then show that it holds for sets of cardinality $\leq k+1$. To this end, let $A \cup B$ be a set of cardinality $k+1$, with A and B being two disjoint sets of cardinality $\leq k$. Let $\tau(S)$ denote the event that set S has at least one type 1 item. Then, we have

$$\begin{aligned}
\mathbb{P}[\tau(A \cup B) \mid \tau(A)^c \vee \tau(B)^c] &= \mathbb{P}[\tau(A) \vee \tau(B) \mid \tau(A)^c \vee \tau(B)^c] \\
&= \frac{\mathbb{P}[(\tau(A)^c \wedge \tau(B)) \vee (\tau(A) \wedge \tau(B)^c)]}{\mathbb{P}[\tau(A)^c \vee \tau(B)^c]} \\
&= \frac{(1 - \mathbb{P}[\tau(A)])\mathbb{P}[\tau(B)] + \mathbb{P}[\tau(A)](1 - \mathbb{P}[\tau(B)])}{1 - \mathbb{P}[\tau(A)]\mathbb{P}[\tau(B)]}.
\end{aligned}$$

Now, note

$$\begin{aligned}
(1 - \mathbb{P}[\tau(A)])\mathbb{P}[\tau(B)] &= \frac{1}{C}p(1-p)|B| \\
\mathbb{P}[\tau(A)](1 - \mathbb{P}[\tau(B)]) &= \frac{1}{C}p(1-p)|A|
\end{aligned}$$

and

$$1 - \mathbb{P}[\tau(A)]\mathbb{P}[\tau(B)] = \frac{1}{C}(1-p)(1-p - (|A| + |B|)p)$$

where $C := (1-p + |A|p)(1-p + |B|p)$.

Hence,

$$\mathbb{P}[\tau(A \cup B) \mid \tau(A)^c \vee \tau(B)^c] = \frac{p(|A| + |B|)}{1-p + (|A| + |B|)p}$$

which shows that (6.4.14) holds for $A \cup B$.

6.4.4 Proof of Proposition 6.3.4

We break τ into two parts: (a) τ_1 denoting the number of steps spent examining pairs of regular matching sets, and (b) τ_2 denoting the number of steps spent for examining a pair of matching sets involving a special set.

We first upper bound the expected value of τ_1 . Conditional on X_s , X_{s+1} is a binomial random variable with parameters $\lfloor X_s/2 \rfloor$ and $1 - \pi(2^s)^2$. Hence, for $s \geq 1$,

$$\mathbb{E}[X_{s+1}] = \mathbb{E} \left[\left\lfloor \frac{X_s}{2} \right\rfloor \right] (1 - \pi(2^s)^2) \quad (6.4.15)$$

and

$$\mathbb{E}[X_1] = \frac{n}{2}(1 - \pi(1)^2). \quad (6.4.16)$$

From (6.4.15), we have

$$\mathbb{E}[X_{s+1}] \leq \mathbb{E}[X_s] \frac{1 - \pi(2^s)^2}{2}, \text{ for } s \geq 1$$

which together with (6.4.16) yields

$$\mathbb{E}[X_s] \leq n \frac{1}{2^s} \prod_{i=0}^{s-1} (1 - \pi(2^i)^2). \quad (6.4.17)$$

Since in each super-epoch s exploration of pairs from any two regular matching sets takes at most $2^s 2^s = 4^s$ steps and there are at most $X_s/2$ pairs of matching sets, and there are at most $\log_2(n)$ super-epochs, we have

$$\mathbb{E}[\tau_1] \leq \frac{1}{2} \sum_{s=1}^{\log_2(n)} 4^s \mathbb{E}[X_s]. \quad (6.4.18)$$

From (6.4.15), $\mathbb{E}[X_s] \leq n2^{-s}$, hence it follows

$$\mathbb{E}[\tau_1] \leq n^2.$$

Let us define $s^* = \lfloor \log_2(c/p) \rfloor$, where c is some constant such that $c > 2/(\sqrt{2} - 1)$.

From (6.4.15) and (6.4.18), we have

$$\mathbb{E}[\tau_1] \leq n(I(1, s^*) + I(s^*, \infty))$$

where

$$I(x, y) := \sum_{s=x}^{y-1} 2^s \prod_{i=0}^{s-1} (1 - \pi(2^i)^2).$$

We first bound the term $I(1, s^*)$ as follows

$$I(1, s^*) \leq \sum_{s=0}^{s^*-1} 2^s \leq 2^{s^*} \leq \frac{c}{p}.$$

We next bound the term $I(s^*, \infty)$. By definition of s^* , note that $s^* \geq \log_2(c/p) - 1$, hence $p2^{s^*} \geq c/2$. It is readily observed that this with $c > 2/(\sqrt{2} - 1)$ implies $2(1 - \pi(2^{s^*})^2) < 1$.

We can write

$$\begin{aligned} I(s^*, \infty) &\leq \sum_{s=s^*}^{\infty} (2(1 - \pi(2^{s^*})^2))^s \\ &= \frac{1}{2\pi(2^{s^*})^2 - 1} \left(2(1 - \pi(2^{s^*})^2)\right)^{s^*} \\ &\leq \frac{1}{2\pi(2^{s^*})^2 - 1} \\ &\leq \frac{1}{2(c/(c+2))^2 - 1} \\ &= O(1) \end{aligned}$$

where in the last inequality we use the fact that $\pi(a)$ is increasing in a , and $2^{s^*} \geq c/(2p)$, hence $\pi(2^{s^*}) \geq (c/2)/(1 - p + c/2) \geq c/(c+2)$.

Therefore, we have

$$\frac{1}{2} \sum_{s=1}^{\log_2(n)} 4^s \mathbb{E}[X_s] \leq \left(\frac{1}{2(c/(c+2))^2 - 1} + \frac{c}{p} \right) n. \quad (6.4.19)$$

Hence, it follows

$$\mathbb{E}[\tau_1] = O\left(\min\left\{\frac{n}{p}, n^2\right\}\right).$$

We next consider the number of steps τ_2 spent on exploring special matching sets. We have already noted that in each super-epoch s , the cardinality of the special set is at most 2^{s+1} . The special matching set can be paired with a set of cardinality at most 2^{s+1} . Hence, in super-epoch s , the expected number of steps spent on examining matching sets involving the special matching set is at most

$2^{s+1}2^{s+1} = 4^{s+1}$. Now, note

$$\begin{aligned}\mathbb{E}[\tau_2] &\leq 4 \sum_{s=1}^{\log_2(n)} 4^s \Pr[X_s \geq 1] \\ &\leq 4 \sum_{s=1}^{\log_2(n)} 4^s \mathbb{E}[X_s]\end{aligned}$$

where the second inequality is by Markov's inequality. We have already shown that

$$\sum_{s=1}^{\log_2(n)} 4^s \mathbb{E}[X_s] = O(\min\{n/p, n^2\}) \quad (6.4.20)$$

, hence, the result follows.

6.4.5 Proof of Theorem 6.3.5

Let X_s be the number of regular matching sets and let Y_s be the size of the special matching set at the beginning of super-epoch s . Let Z_r be a random variable taking value 1 if both of the two matching sets examined in epoch r contain a type 1 node, and $Z_r = 0$, otherwise.

Note that conditional on that in epoch r of super-epoch s two regular matching sets are examined, we have $Z_r \sim \text{Ber}(\pi(2^s)^2)$, and if the examination involves the special matching set, we have $Z_r \sim \text{Ber}(\pi(Y_s)\pi(\zeta_s))$, where $\zeta_s = 2^s$ if X_s is odd, and $\zeta_s = 2^{s+1}$ if X_s is even and $X_{s+1} > 0$.

Let S_r be the number of steps in epoch r , T_{s-1} be the first epoch of super-epoch s with $T_0 = 1$. Let \tilde{R}_s denote the regret accumulated in super-epoch s . To simplify notation, we define $K_s := \lfloor X_s/2 \rfloor$, which corresponds to the number of epochs in which two regular matching sets are explored in super-epoch s . We also use the following notation $p_s := \pi(2^s)$ and $q_s := 1 - p_s^2$. We show a graphic illustration of various definitions in Figure 13.

Note that there can be at most $\log_2(n)$ super-epochs.

Let X_s^1 be the number of regular matching sets at the beginning of super-epoch s that contain a type 1 node. Let Y_s^1 be a random variable taking value 1 if the special matching set has a type 1 node, and otherwise $Y_s^1 = 0$. If $Y_s = 0$ (indicating that at the beginning of super-epoch s there is no special matching set), we define $Y_s^1 = 0$. Note that $X_s^1 | X_s \sim \text{Bin}(X_s, p_s)$ and $Y_s^1 | Y_s \sim \text{Ber}(\pi(Y_s))$.

Let α_s (resp. β_s) be the conditional expected number of steps in epoch r , conditional on $Z_r = 0$ (resp. $Z_r = 1$), where s is the super-epoch of epoch r , i.e.

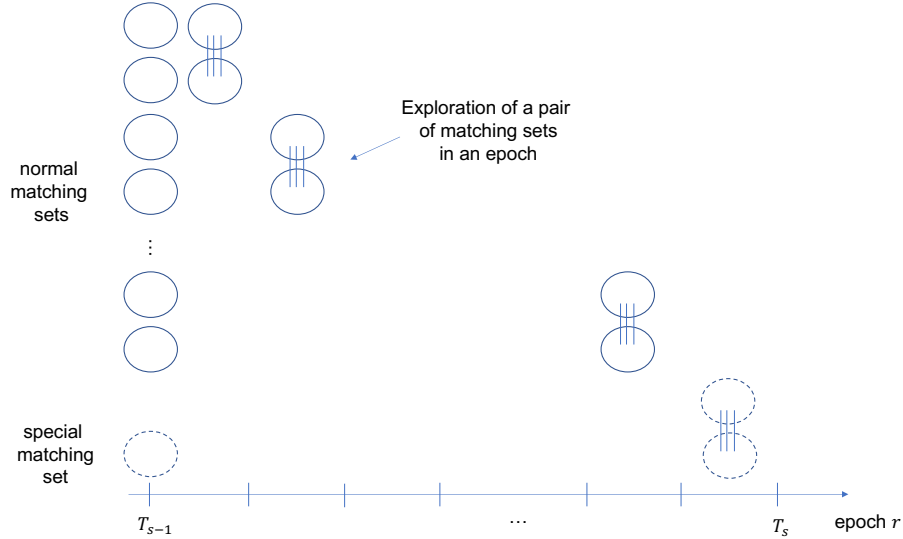


Figure 13: Matchings created in super-epoch s .

$\alpha_s = \mathbb{E}[S_r | Z_r = 0]$ and $\beta_s = \mathbb{E}[S_r | Z_r = 1]$. Note that $\beta_s = \alpha_s/2$.

In each step, the instantaneous regret is equal to the number of pairs of type 1 nodes that can be formed from the number of unidentified type 1 nodes. Since each regular or special matching set contains at most one type 1 node, we can express the instantaneous regret at the beginning of super-epoch s as follows:

$$R_{T_{s-1}} = \left\lfloor \frac{X_s^1}{2} \right\rfloor + Y_s^1 \mathbf{1}_{\{X_s^1 \text{ is odd}\}}. \quad (6.4.21)$$

The instantaneous regret in epoch r , such that $T_{s-1} < r \leq T_s$, satisfies the following relation

$$R_r = R_{T_{s-1}} - \sum_{i=T_{s-1}+1}^r Z_i. \quad (6.4.22)$$

We will evaluate the cumulative regret of the algorithm by summing up the regret accumulated in each super-epoch. The regret accumulated in super-epoch s can be expressed as follows:

$$\tilde{R}_s = A_s + B_s$$

where

$$A_s := \sum_{r=T_{s-1}+1}^{T_{s-1}+K_s} R_{r-1} S_r \quad (6.4.23)$$

and

$$B_s := R_{T_{s-1}+K_s} S_{T_{s-1}+K_s+1} \left(\mathbf{1}_{\{X_s \text{ is odd}\}} + \mathbf{1}_{\{Y_s > 0\}} \mathbf{1}_{\{X_s \text{ is even}\}} + \mathbf{1}_{\{X_{s+1} > 0\}} \right) \mathbf{1}_{\{X_s > 0\}}. \quad (6.4.24)$$

Here A_s is the cumulative regret incurred in epochs of super-epoch s in which regular matching sets are examined, and B_s is the regret incurred in the last epoch of super-epoch s that involves examination of the special matching set, or a regular matching set of size 2^{s+1} and a remaining regular matching set of size 2^s if there is odd number of regular matching sets of size 2^s at the initial epoch of the super-epoch s .

We next note some basic properties of binomial distribution, which will be used in the rest of the proof.

Lemma 6.4.1. *Assume that X is a binomial random variable with parameters n and p . Then,*

1. $\mathbb{E}[X^2] = n(n-1)p^2 + np$
2. $\mathbb{E}[X^3] = n(n-1)(n-2)p^3 + 3n(n-1)p^2 + np$.

In the following we evaluate the expected values of A_s and B_s defined in (6.4.23) and (6.4.24), respectively.

Lemma 6.4.2. *For every super-epoch $1 \leq s \leq \log_2(n)$, we have*

$$\mathbb{E}[A_s] = \frac{1}{4} \left(1 - \frac{1}{2}p_s\right) \left(1 - \frac{1}{2}p_s^2\right) p_s \alpha_s \mathbb{E}[X_s^2] + O(n).$$

Proof. From (6.4.22) and (6.4.23), $A_s = A_s^{(1)} - A_s^{(2)}$ where

$$A_s^{(1)} = R_{T_{s-1}} \sum_{r=T_{s-1}+1}^{T_{s-1}+K_s} S_r \quad (6.4.25)$$

and

$$A_s^{(2)} := \sum_{r=T_{s-1}+1}^{T_{s-1}+K_s} S_r \sum_{i=T_{s-1}+1}^r Z_i. \quad (6.4.26)$$

We separately upper bound the expected values of $A_s^{(1)}$ and $A_s^{(2)}$. First, we upper bound the expected value of $A_s^{(1)}$. Note that

$$\mathbb{E} \left[A_s^{(1)} \mid T_{s-1}, X_s, X_s^1, Y_s^1 \right] = \mathbb{E} \left[R_{T_{s-1}} \mid X_s, X_s^1, Y_s^1 \right] N_s$$

where

$$N_s := \mathbb{E} \left[\sum_{r=T_{s-1}+1}^{T_{s-1}+K_s} S_r \mid T_{s-1}, X_s, X_s^1 \right].$$

Combining with (6.4.21), we have

$$\mathbb{E} \left[A_s^{(1)} \mid T_{s-1} X_s, X_s^1, Y_s^1 \right] = \left(\left\lfloor \frac{X_s^1}{2} \right\rfloor + Y_s^1 \mathbf{1}_{\{X_s^1 \text{ is odd}\}} \right) N_s. \quad (6.4.27)$$

We next evaluate N_s . Let us define K_s^{11} to be the number of epochs in super-epoch s in which two regular matching sets are examined such that both have a value 1 node, i.e. $K_s^{11} := \sum_{r=T_{s-1}+1}^{T_{s-1}+K_s} Z_r$. Then, note

$$N_s = \beta_s \mathbb{E} \left[K_s^{11} \mid X_s, X_s^1 \right] + \alpha_s \mathbb{E} \left[K_s - K_s^{11} \mid X_s, X_s^1 \right] = \alpha_s \left(K_s - \frac{1}{2} \mathbb{E}[K_s^{11} \mid X_s, X_s^1] \right).$$

Now, note $\mathbb{E} [K_s^{11} \mid X_s, X_s^1] = K_s (X_s^1 / X_s)^2$ and, hence,

$$N_s = \alpha_s K_s \left(1 - \frac{1}{2} \left(\frac{X_s^1}{X_s} \right)^2 \right) = \alpha_s \left\lfloor \frac{X_s}{2} \right\rfloor \left(1 - \frac{1}{2} \left(\frac{X_s^1}{X_s} \right)^2 \right).$$

Note that

$$\begin{aligned} \mathbb{E} \left[\left\lfloor \frac{X_s^1}{2} \right\rfloor \left\lfloor \frac{X_s}{2} \right\rfloor \left(1 - \frac{1}{2} \left(\frac{X_s^1}{X_s} \right)^2 \right) \mid X_s \right] &= \frac{X_s}{4} \mathbb{E} \left[X_s^1 - \frac{(X_s^1)^3}{2X_s^2} \mid X_s \right] \\ &= \frac{X_s \mathbb{E}[X_s^1 \mid X_s]}{4} - \frac{\mathbb{E}[(X_s^1)^3 \mid X_s]}{8X_s} \\ &= \frac{1}{4} p_s X_s^2 - \frac{1}{8} p_s^3 X_s^2 + O(n) \\ &= \frac{1}{4} p_s \left(1 - \frac{1}{2} p_s^2 \right) X_s^2 + O(n) \end{aligned}$$

where we used Lemma 6.4.1.

Combining with (6.4.27) we have

$$\mathbb{E} \left[A_s^{(1)} \right] = \frac{1}{4} p_s \left(1 - \frac{1}{2} p_s^2 \right) \alpha_s \mathbb{E}[X_s^2] + O(n).$$

Next, we upper bound the expected value of $A_s^{(2)}$. Note that

$$\begin{aligned}
\mathbb{E} \left[A_s^{(2)} \right] &= \mathbb{E} \left[\sum_{r=T_{s-1}+1}^{T_{s-1}+K_s} S_r \sum_{i=T_{s-1}+1}^r Z_i \right] \\
&= \mathbb{E} \left[\mathbb{E} \left[\sum_{r=T_{s-1}+1}^{T_{s-1}+K_s} S_r \sum_{i=T_{s-1}+1}^r Z_i \mid T_{s-1}, K_s \right] \right] \\
&= \mathbb{E} \left[\sum_{r=T_{s-1}+1}^{T_{s-1}+K_s} \mathbb{E} \left[(\alpha_s (1 - Z_r) + \beta_s Z_r) \sum_{i=T_{s-1}+1}^r Z_i \mid T_{s-1}, K_s \right] \right] \\
&= \mathbb{E} \left[\sum_{r=T_{s-1}+1}^{T_{s-1}+K_s} \mathbb{E} \left[\alpha_s \sum_{i=T_{s-1}+1}^r Z_i - \frac{1}{2} \alpha_s Z_r \sum_{i=T_{s-1}+1}^r Z_i \mid T_{s-1}, K_s \right] \right]
\end{aligned}$$

where the last equality holds because $\beta_s = \alpha_s/2$.

Now, note

$$\mathbb{E} \left[\sum_{i=T_{s-1}+1}^r Z_i \mid T_{s-1}, K_s \right] = \sum_{i=T_{s-1}+1}^r \mathbb{P} [Z_i = 1] = (r - T_{s-1}) p_s^2$$

and

$$\begin{aligned}
\mathbb{E} \left[Z_r \sum_{i=T_{s-1}+1}^r Z_i \mid T_{s-1}, K_s \right] &= \mathbb{E} \left[Z_r^2 \mid T_{s-1}, K_s \right] + \sum_{i=T_{s-1}+1}^{r-1} \mathbb{E} [Z_r Z_i \mid T_{s-1}, K_s] \\
&= \mathbb{P} [Z_r = 1] + \sum_{i=T_{s-1}+1}^{r-1} \mathbb{P} [Z_r = 1] \mathbb{P} [Z_i = 1] \\
&= p_s^2 + (r - T_{s-1} - 1) p_s^4.
\end{aligned}$$

Putting the pieces together, we obtain

$$\begin{aligned}
\mathbb{E} [A_s^{(2)}] &= \mathbb{E} \left[\sum_{r=T_{s-1}+1}^{T_{s-1}+K_s} \alpha_s (r - T_{s-1}) p_s^2 - \frac{1}{2} \alpha_s [p_s^2 + (r - T_{s-1} - 1) p_s^4] \right] \\
&= \mathbb{E} \left[\sum_{r=T_{s-1}+1}^{T_{s-1}+K_s} \alpha_s p_s^2 (1 - \frac{1}{2} p_s^2) (r - T_{s-1}) - \frac{1}{2} \alpha_s p_s^2 (1 - p_s^2) \right] \\
&= \mathbb{E} \left[\alpha_s p_s^2 (1 - \frac{1}{2} p_s^2) \sum_{i=1}^{K_s} i - \frac{1}{2} \alpha_s p_s^2 (1 - p_s^2) K_s \right] \\
&= \mathbb{E} \left[\frac{1}{2} \alpha_s p_s^2 (1 - p_s^2) K_s^2 + \frac{1}{4} \alpha_s p_s^4 K_s (K_s + 1) \right] \\
&= \frac{1}{8} p_s^2 \left(1 - \frac{1}{2} p_s^2 \right) \alpha_s \mathbb{E}[X_s^2] + O(n)
\end{aligned}$$

where the last equality holds because $K_s = \lfloor X_s/2 \rfloor$.

Using the upper bounds for $\mathbb{E}[A_s^{(1)}]$ and $\mathbb{E}[A_s^{(2)}]$, we have

$$\mathbb{E}[A_s] = \frac{1}{4} \left(1 - \frac{1}{2} p_s \right) \left(1 - \frac{1}{2} p_s^2 \right) p_s \alpha_s \mathbb{E}[X_s^2] + O(n).$$

□

We next show an upper bound for the expected value of B_s .

Lemma 6.4.3. *For every super-epoch $1 \leq s \leq \log_2(n)$, we have*

$$\mathbb{E}[B_s] \leq n p_s 2^{s+1} \prod_{i=0}^{s-1} q_i + 4^{s+1} \mathbb{P}[X_s > 0].$$

Proof. In the last epoch of super-epoch s , the two matching sets are of cardinality at least 2^s and at most 2^{s+1} , hence, from (6.4.24), we have

$$B_s \leq R_{T_{s-1}+K_s} \mathbf{1}_{\{X_s > 0\}} 4^{s+1}. \quad (6.4.28)$$

From (6.4.21), it follows that $R_{T_{s-1}+K_s} \leq R_{T_{s-1}} \leq \lfloor X_s^1/2 \rfloor + 1$ where recall $X_s^1 \mid X_s \sim \text{Bin}(X_s, p_s)$, and hence

$$\mathbb{E}[R_{T_{s-1}+K_s} \mid X_s = x] \leq \frac{1}{2} p_s x + 1.$$

Now, note

$$\begin{aligned}
\mathbb{E}[R_{T_{s-1}+K_s} \mid X_s > 0] &= \frac{\mathbb{E}[R_{T_{s-1}+K_s} \mathbf{1}_{\{X_s > 0\}}]}{\mathbb{P}[X_s > 0]} \\
&= \frac{\sum_{x=1}^{\infty} \mathbb{E}[R_{T_{s-1}+K_s} \mid X_s = x] \mathbb{P}[X_s = x]}{\mathbb{P}[X_s > 0]} \\
&\leq \frac{1}{2} p_s \frac{\sum_{x=1}^{\infty} x \mathbb{P}[X_s = x]}{\mathbb{P}[X_s > 0]} + 1 \\
&= \frac{1}{2} p_s \frac{\mathbb{E}[X_s \mathbf{1}_{\{X_s > 0\}}]}{\mathbb{P}[X_s > 0]} + 1 \\
&= \frac{1}{2} p_s \frac{\mathbb{E}[X_s]}{\mathbb{P}[X_s > 0]} + 1.
\end{aligned}$$

Hence, we have

$$\mathbb{E}[R_{T_{s-1}+K_s} \mathbf{1}_{\{X_s > 0\}}] \leq \frac{1}{2} p_s \mathbb{E}[X_s] + \mathbb{P}[X_s > 0]$$

which combined with $\mathbb{E}[X_s] \leq n(1/2^s) \prod_{i=0}^{s-1} q_i$ that follows from (6.3.1), yields

$$\mathbb{E}[R_{T_{s-1}+K_s} \mathbf{1}_{\{X_s > 0\}}] \leq n p_s \frac{1}{2^{s+1}} \prod_{i=0}^{s-1} q_i + \mathbb{P}[X_s > 0].$$

Combining with (6.4.28), we have

$$\mathbb{E}[B_s] \leq n p_s 2^{s+1} \prod_{i=0}^{s-1} q_i + 4^{s+1} \mathbb{P}[X_s > 0].$$

□

Now, note by Markov's inequality and (6.4.19), for a constant $c > 2/(\sqrt{2} - 1)$,

$$\sum_{s=1}^{\log_2(n)} 4^{s+1} \mathbb{P}[X_s > 0] \leq \sum_{s=1}^{\log_2(n)} 4^{s+1} \mathbb{E}[X_s] \leq 8 \left(\frac{1}{2(c/(c+2))^2 - 1} + \frac{c}{p} \right) n$$

and, obviously, we have

$$\sum_{s=1}^{\log_2(n)} 4^{s+1} \mathbb{P}[X_s > 0] \leq \sum_{s=1}^{\log_2(n)} 4^{s+1} \leq \frac{4}{3} n^2.$$

Putting the pieces together, we have

$$\mathbb{E}[R] \leq \sum_{s=1}^{\log_2(n)} \mathbb{E}[A_s + B_s] \leq a_{n,p}n^2 + b_{n,p}n^2 + c_{n,p}n^2 + O(n \log(n))$$

where

$$a_{n,p} = \frac{1}{4} \sum_{s=1}^{\log_2(n)} \left(1 - \frac{1}{2}p_s\right) \left(1 - \frac{1}{2}p_s^2\right) p_s \left(\prod_{i=0}^{s-1} q_i^2 + \frac{1}{n} \sum_{i=0}^{s-1} 2^{i+1} \binom{i}{j=0} q_j \right) \left(\prod_{j=i+1}^{s-1} q_j^2 \right),$$

$$b_{n,p} = \frac{1}{n} \sum_{s=1}^{\log_2(n)} 2^{s+1} p_s \prod_{i=0}^{s-1} q_i,$$

and

$$c_{n,p} = \min \left\{ \frac{c_1}{np}, \frac{4}{3} \right\}$$

where c_1 is a constant such that $c_1 > 16/(\sqrt{2} - 1)$.

6.4.6 Properties of $a_{n,p}$, $b_{n,p}$ and $c_{n,p}$

We consider properties of $a_{n,p}$, $b_{n,p}$ and $c_{n,p}$ that appear in the regret upper bound in Theorem 6.3.5. For $c_{n,p}$ we note that $c_{n,p} = o(1)$ whenever $p = \omega(1/n)$. In other words, $c_{n,p}$ diminishes with n whenever p is a constant or it goes to 0 as $\omega(1/n)$. In the following, we focus on showing some properties for $a_{n,p}$ and $b_{n,p}$.

Let $a'_{n,p}$ and $a''_{n,p}$ be the two terms in $a_{n,p}$, i.e.

$$a'_{n,p} := \frac{1}{4} \sum_{s=1}^{\log_2(n)} \left(1 - \frac{1}{2}p_s\right) \left(1 - \frac{1}{2}p_s^2\right) p_s \prod_{i=0}^{s-1} q_i^2 \quad (6.4.29)$$

and

$$a''_{n,p} := \frac{1}{4} \sum_{s=1}^{\log_2(n)} \left(1 - \frac{1}{2}p_s\right) \left(1 - \frac{1}{2}p_s^2\right) p_s \frac{1}{n} \sum_{i=0}^{s-1} 2^{i+1} \binom{i}{j=0} q_j \left(\prod_{j=i+1}^{s-1} q_j^2 \right). \quad (6.4.30)$$

We will use the following basic facts. Note that

$$p_i = \frac{p2^i}{1-p+p2^i} \geq 1 - \frac{1-p}{p} 2^{-i}.$$

The right-hand side in the last inequality is positive if and only if $i \geq \log_2((1 -$

$p)/p)$). The latter condition holds for all $i \geq 0$ provided that $1/2 \leq p \leq 1$.

It follows that for every $i \geq \log_2((1-p)/p)$, we have

$$q_i = 1 - p_i^2 \leq 2 \frac{1-p}{p} 2^{-i}.$$

We next consider properties of $a_{n,p}$ and $b_{n,p}$, for three different cases depending on the value of parameter p .

Case $1/2 < p < 1$ From (6.4.29) we have

$$a'_{n,p} \leq \frac{1}{4} \sum_{s=1}^{\log_2(n)} \prod_{i=0}^{s-1} q_i^2$$

and, note that

$$\prod_{i=0}^{s-1} q_i^2 \leq \prod_{i=0}^{s-1} \left(2 \frac{1-p}{p}\right)^2 4^{-i} = \left(8 \left(\frac{1-p}{p}\right)^2\right)^s 2^{-s^2}.$$

Hence, it follows

$$a'_{n,p} \leq \frac{1}{4} \sum_{s=1}^{\log_2(n)} \left(8 \left(\frac{1-p}{p}\right)^2\right)^s 2^{-s^2}. \quad (6.4.31)$$

We use the following facts to bound the sum in (6.4.31). By straightforward calculus, we have

$$\int_a^b e^{cx} e^{-dx^2} dx = \frac{\sqrt{\pi}}{\sqrt{d}} e^{\frac{c^2}{4d}} \left[\Phi\left(\sqrt{2d}\left(b - \frac{c}{2d}\right)\right) - \Phi\left(\sqrt{2d}\left(a - \frac{c}{2d}\right)\right) \right] \quad (6.4.32)$$

where $\Phi(x)$ is the cumulative distribution function of a normal random variable.

It follows

$$\int_a^b e^{cx} e^{-dx^2} dx \leq \frac{\sqrt{\pi}}{\sqrt{d}} e^{\frac{c^2}{4d}} \Phi^c\left(\sqrt{2d}\left(a - \frac{c}{2d}\right)\right)$$

where $\Phi^c(x) = 1 - \Phi(x)$.

The following is a well known bound:

$$\Phi^c(x) \leq \frac{1}{\sqrt{2\pi}} \frac{1}{x} e^{-\frac{x^2}{2}} \text{ for } x > 0.$$

If $a - c/(2d) > 0$, we have

$$\int_a^b e^{cx} e^{-dx^2} dx \leq \frac{1}{2d} \frac{1}{a - \frac{c}{2d}} e^{-(da^2 - ca)}. \quad (6.4.33)$$

Note that

$$\begin{aligned} \sum_{s=1}^{\log_2(n)} \left(8 \left(\frac{1-p}{p} \right)^2 \right)^s 2^{-s^2} &\leq 4 \left(\frac{1-p}{p} \right)^2 + 4 \left(\frac{1-p}{p} \right)^4 + \left(\frac{1-p}{p} \right)^6 \\ &\quad + \sum_{s=4}^{\log_2(n)} \left(\left(\frac{1-p}{p} \right)^2 \right)^s 2^{-\frac{1}{4}s^2}. \end{aligned}$$

Now, by applying (6.4.33) with $c = \log((1-p)/p)$ and $d = (1/4) \log(2)$, we have

$$\begin{aligned} \sum_{s=4}^{\log_2(n)} \left(\left(\frac{1-p}{p} \right)^2 \right)^s 2^{-\frac{1}{4}s^2} &\leq \int_3^{\log_2(n)-1} \left(\left(\frac{1-p}{p} \right)^2 \right)^x 2^{-\frac{1}{2}x^2} dx \\ &\leq \frac{2}{\log(2)} \frac{1}{3 - 4 \log_2 \left(\frac{1-p}{p} \right)} \\ &\quad \exp \left(- \left(\frac{9}{4} \log(2) - 6 \log(2) \log_2 \left(\frac{1-p}{p} \right) \right) \right) \\ &= \frac{1}{2^{5/4} \log(2)} \frac{1}{3 + 4 \log_2 \left(\frac{p}{1-p} \right)} \left(\frac{1-p}{p} \right)^6 \\ &\leq \frac{1}{3 \cdot 2^{5/4} \log(2)} \left(\frac{1-p}{p} \right)^6. \end{aligned}$$

We have shown that

$$\sum_{s=1}^{\log_2(n)} \left(8 \left(\frac{1-p}{p} \right)^2 \right)^s 2^{-s^2} \leq 4 \left(\frac{1-p}{p} \right)^2 + 4 \left(\frac{1-p}{p} \right)^4 + \left(1 + \frac{1}{3 \cdot 2^{5/4} \log(2)} \right) \left(\frac{1-p}{p} \right)^6.$$

Thus, we have

$$a'_{n,p} \leq \left(\frac{1-p}{p} \right)^2 + \left(\frac{1-p}{p} \right)^4 + \frac{1}{4} \left(1 + \frac{1}{3 \cdot 2^{5/4} \log(2)} \right) \left(\frac{1-p}{p} \right)^6.$$

From this, observe that $a'_{n,p} = O((1-p)^2)$.

Next, we consider $a''_{n,p}$ and $b_{n,p}$. Using (6.4.30), we have

$$\begin{aligned}
a''_{n,p} &\leq \frac{1}{4} \sum_{s=1}^{\log_2(n)} \frac{1}{n} \sum_{i=0}^{s-1} 2^{i+1} \left(\prod_{j=0}^i q_j \right) \left(\prod_{j=i+1}^{s-1} q_j^2 \right) \\
&\leq \frac{1}{4} \sum_{s=1}^{\log_2(n)} \frac{1}{n} \sum_{i=0}^{s-1} 2^{i+1} q_{i+1} \left(\prod_{j=0}^{s-1} q_j \right) \\
&\leq \frac{1}{4} \sum_{s=1}^{\log_2(n)} \frac{1}{n} \sum_{i=0}^{s-1} 2^{i+1} \left(2^{\frac{1-p}{p}} 2^{-(i+1)} \right) \left(\prod_{j=0}^{s-1} q_j \right) \\
&= \frac{1}{4} \sum_{s=1}^{\log_2(n)} \frac{s}{n} \left(2^{\frac{1-p}{p}} \right) \left(\prod_{j=0}^{s-1} q_j \right) \\
&\leq \frac{1}{2} \frac{\log_2(n)}{n} \frac{1-p}{p} \sum_{s=1}^{\log_2(n)} \prod_{j=0}^{s-1} q_j.
\end{aligned}$$

From definition of $b_{n,p}$, we have

$$b_{n,p} \leq 2 \frac{1}{n} \sum_{s=1}^{\log_2(n)} 2^s \prod_{i=0}^{s-1} q_i.$$

Now, note that

$$\begin{aligned}
\prod_{i=0}^{s-1} q_i &\leq \prod_{i=0}^{s-1} \left(2^{\frac{1-p}{p}} \right) 2^{-i} \\
&= \left(2^{\frac{1-p}{p}} \right)^s 2^{-\frac{s(s-1)}{2}} \\
&= \left((\sqrt{2})^3 \frac{1-p}{p} \right)^s (\sqrt{2})^{-s^2}.
\end{aligned}$$

Summing up, we have

$$\begin{aligned}
\sum_{s=1}^{\log_2(n)} \prod_{i=0}^{s-1} q_i &\leq 2 \left(\frac{1-p}{p} \right) + 2 \left(\frac{1-p}{p} \right)^2 + \sum_{s=3}^{\log_2(n)} \left(\frac{1}{(\sqrt{2})^{s-3}} \frac{1-p}{p} \right)^s \\
&\leq 2 \left(\frac{1-p}{p} \right) + 2 \left(\frac{1-p}{p} \right)^2 + \min \left\{ \frac{\frac{1-p}{p}}{1 - \frac{1-p}{p}}, \frac{1}{1 - 1/\sqrt{2}} \right\}.
\end{aligned}$$

$$\begin{aligned}
\sum_{s=1}^{\log_2(n)} 2^s \prod_{i=0}^{s-1} q_i &\leq \sum_{s=1}^{\log_2(n)} \left(2^{\frac{5-s}{2}} \frac{1-p}{p} \right)^s \\
&\leq \sum_{s=1}^5 \left(2^{\frac{5-s}{2}} \frac{1-p}{p} \right)^s + \sum_{s=5}^{\log_2(n)} \left(\frac{1-p}{p} \right)^s \\
&\leq \sum_{s=1}^5 \left(2^{\frac{5-s}{2}} \frac{1-p}{p} \right)^s + \frac{(1-p)^5}{p^4(2p-1)} \\
&= O(1-p).
\end{aligned}$$

It follows that

$$a''_{n,p} = O\left((1-p)^2 \frac{\log(n)}{n} \right)$$

and

$$b_{n,p} = O\left((1-p) \frac{1}{n} \right).$$

Summing $a'_{n,p}$ and $a''_{n,p}$, we have

$$a_{n,p} \leq \left(\left(\frac{1-p}{p} \right)^2 + \left(\frac{1-p}{p} \right)^4 + \frac{1}{4} \left(1 + \frac{1}{3 \cdot 2^{5/4} \log(2)} \right) \left(\frac{1-p}{p} \right)^6 \right) \left(1 + O\left(\frac{\log(n)}{n} \right) \right).$$

Case $p = O(1/n)$ From (6.4.29) and (6.4.30), note that

$$a'_{n,p} \leq \frac{1}{4} \sum_{s=1}^{\log_2(n)} p_s \text{ and } a''_{n,p} \leq \frac{1}{2} \sum_{s=1}^{\log_2(n)} p_s$$

and, from the definition of $b_{n,p}$, we have

$$b_{n,p} \leq 2 \sum_{s=1}^{\log_2(n)} p_s.$$

We will use the following basic fact

$$\int_a^b \frac{p2^s}{1-p+p2^s} ds = \log_2 \left(\frac{1-p+p2^b}{1-p+p2^a} \right). \quad (6.4.34)$$

Using this basic fact, we have

$$\begin{aligned}
\sum_{s=1}^{\log_2(n)} p_s &= \sum_{s=1}^{\log_2(n)} \frac{p2^s}{1-p+p2^s} \\
&\leq \int_1^{\log_2(n)+1} \frac{p2^s}{1-p+p2^s} ds \\
&= \log_2 \left(\frac{1-p+2np}{1+p} \right) \\
&\leq \log_2(1+2np) \\
&\leq \log_2(1+2c).
\end{aligned}$$

Hence, it follows that

$$a'_{n,p} = O(1), a''_{n,p} = O(1), \text{ and } b_{n,p} = O(1).$$

Case $\omega(1/n) \leq p < 1/2$ We first upper bound $a'_{n,p}$. Let c_2 be a fixed constant in $(0, 1)$ and $s^* = (1/2) \log_2(1/c_2) + \log_2((1-p)/p)$. From (6.4.29), we have

$$a'_{n,p} \leq \frac{1}{4} \sum_{s=1}^{\log_2(n)} p_s \prod_{i=0}^{s-1} q_i^2 \leq \frac{1}{4} \sum_{s=1}^{s^*} p_s + \frac{1}{4} \sum_{s=s^*+1}^{\log_2(n)} \prod_{i=s^*}^{s-1} q_i^2.$$

Note that

$$\begin{aligned}
\prod_{i=s^*}^{s-1} q_i^2 &\leq \prod_{i=s^*}^{s-1} \left(2 \frac{1-p}{p} \right)^2 4^{-i} \\
&= \left(4 \left(\frac{1-p}{p} \right)^2 \right)^{s-s^*} 2^{-((s-1)s - (s^*-1)s^*)} \\
&= \left(8 \left(\frac{1-p}{p} \right)^2 \right)^{s-s^*} 2^{-(s^2-s^*)} \\
&= \left(\frac{1}{2^{s+s^*-3}} \left(\frac{1-p}{p} \right)^2 \right)^{s-s^*}.
\end{aligned}$$

For every $s \geq s^* + 1$, we have

$$\frac{1}{2^{s+s^*-3}} \left(\frac{1-p}{p} \right)^2 \leq \frac{1}{2^{2(s^*-1)}} \left(\frac{1-p}{p} \right)^2 = c_2.$$

Hence, it follows

$$\sum_{s=s^*+1}^{\log_2(n)} \prod_{i=s^*}^{s-1} q_i^2 \leq \sum_{s=s^*+1}^{\log_2(n)} c_2^s \leq \frac{1}{1-c_2}.$$

Now, note also

$$\begin{aligned} \sum_{s=1}^{s^*} p_s &= \sum_{s=1}^{s^*} \frac{p2^s}{1-p+p2^s} \\ &\leq \int_1^{s^*+1} \frac{p2^s}{1-p+p2^s} ds \\ &= \log_2 \left(\frac{1-p+p2^{s^*+1}}{1+p} \right) \\ &\leq \log_2(1+p2^{s^*+1}) \\ &\leq \log_2(1+2/\sqrt{c_2}) \end{aligned}$$

where we used (6.4.34).

We have thus shown that

$$a'_{n,p} \leq \frac{1}{4} \left(\log_2 \left(1 + \frac{1}{\sqrt{c_2}} \right) + \frac{1}{1-c_2} \right).$$

We next consider $a''_{n,p}$ defined in (6.4.30). Fix a constant c in $(0,1)$ and let $s^* = \log_2(1/c_2) + \log_2((1-p)/p) + 1$. We break the sum in the definition of $a''_{n,p}$ in two parts, A and B , that can be bounded as follows:

$$\begin{aligned}
A &:= \sum_{s=1}^{s^*} \left(1 - \frac{1}{2}p_s\right) \left(1 - \frac{1}{2}p_s^2\right) p_s \frac{1}{n} \sum_{i=0}^{s-1} 2^{i+1} \binom{i}{j=0} q_j \binom{s-1}{j=i+1} q_j^2 \\
&\leq \sum_{s=1}^{s^*} p_s \frac{1}{n} \sum_{i=0}^{s-1} 2^{i+1} \\
&\leq \sum_{s=1}^{s^*} p_s \frac{1}{n} 2(2^s - 1) \\
&\leq 2 \frac{2^{s^*} - 1}{n} \sum_{s=1}^{s^*} p_s \\
&\leq 2 \frac{2^{s^*} - 1}{n} \log_2 \left(\frac{1 - p + p2^{s^*+1}}{1 + p} \right) \\
&\leq 2 \frac{2^{s^*} - 1}{n} \log_2 (1 + p2^{s^*+1})
\end{aligned}$$

$$\begin{aligned}
B &:= \sum_{s=s^*+1}^{\log_2(n)} \left(1 - \frac{1}{2}p_s\right) \left(1 - \frac{1}{2}p_s^2\right) p_s \frac{1}{n} \sum_{i=0}^{s-1} 2^{i+1} \binom{i}{j=0} q_j \binom{s-1}{j=i+1} q_j^2 \\
&\leq \sum_{s=s^*+1}^{\log_2(n)} \frac{1}{n} \sum_{i=0}^{s-1} 2^{i+1} \binom{s-1}{j=s^*} q_j \\
&\leq 2 \sum_{s=s^*+1}^{\log_2(n)} \binom{s-1}{j=s^*} q_j.
\end{aligned}$$

Note that

$$\begin{aligned}
\prod_{i=s^*}^{s-1} q_i &\leq \prod_{i=s^*}^{s-1} \left(2 \frac{1-p}{p}\right) 2^{-i} \\
&= \left(2 \left(\frac{1-p}{p}\right)\right)^{s-s^*} 2^{-\frac{1}{2}((s-1)s - (s^*-1)s^*)} \\
&= \left((\sqrt{2})^3 \left(\frac{1-p}{p}\right)\right)^{s-s^*} (\sqrt{2})^{-(s^2 - s^*)} \\
&= \left(\frac{1}{(\sqrt{2})^{s+s^*-3}} \left(\frac{1-p}{p}\right)\right)^{s-s^*}.
\end{aligned}$$

For every $s \geq s^* + 1$, we have

$$\frac{1}{(\sqrt{2})^{s+s^*-3}} \left(\frac{1-p}{p} \right) \leq \frac{1}{(\sqrt{2})^{2(s^*-1)}} \left(\frac{1-p}{p} \right) = c_2.$$

It follows that

$$B \leq 2 \frac{1}{1-c_2}$$

and

$$\begin{aligned} A &\leq 2 \frac{2^{\frac{1-p}{c_2}} - 1}{n} \log_2 \left(1 + 4(1-p) \frac{1}{c_2} \right) \\ &\leq \frac{4}{c_2} \frac{1}{pn} \log_2 \left(1 + \frac{4}{c_2} \right). \end{aligned}$$

We have thus shown that

$$a''_{n,p} \leq \frac{1}{c_2} \frac{1}{pn} \log_2 \left(1 + \frac{4}{c_2} \right) + \frac{1}{2} \frac{1}{1-c_2} = \frac{1}{2} \frac{1}{1-c_2} + o(1).$$

We next consider $b_{n,p}$. This is similar to bounding $a''_{n,p}$ by noting that

$$\begin{aligned} b_{n,p} &\leq \frac{1}{n} \sum_{s=1}^{s^*} 2^{s+1} p_s + \frac{1}{n} \sum_{s=s^*+1}^{\log_2(n)} 2^{s+1} \prod_{i=s^*}^{s-1} q_i \\ &\leq \frac{2(2^{s^*} - 1)}{n} \sum_{s=1}^{s^*} p_s + 2 \sum_{s=s^*+1}^{\log_2(n)} \prod_{i=s^*}^{s-1} q_i. \end{aligned}$$

By the bounds established for $a''_{n,p}$, we have

$$b_{n,p} \leq \frac{1}{c_2} \frac{1}{pn} \log_2 \left(1 + \frac{4}{c_2} \right) + \frac{1}{2} \frac{1}{1-c_2} = \frac{1}{2} \frac{1}{1-c_2} + o(1).$$

Chapter 7

Conclusions

We conclude with a summary of the work in this thesis and discussions on future research.

Cooperative non-stochastic multi-armed bandits We presented new results for the collaborative multi-agent non-stochastic multi-armed bandit with communication delays. We showed a lower bound on the regret of each individual agent and proposed two algorithms (CFTRL and DFTRL) together with their regret upper bounds. CFTRL provides an optimal regret of each individual agent with respect to the scaling with the number of arms. DFTRL has an optimal average regret with respect to the scaling with the edge-delay. Our numerical results validate our theoretical bounds and demonstrate that significant performance gains can be achieved by our two algorithms compared to state-of-the-art algorithms.

There are several open research questions for future research. The first question is to consider the existence of a decentralized algorithm which can provide $O((1/\sqrt{|\mathcal{N}(v)|})\sqrt{KT})$ individual regret for each agent v . It is unclear whether a center-based communication protocol is necessary to achieve this regret. The second question is to consider whether an algorithm exists with an optimal scaling with the number of arms and the edge-delay parameter. The third question is to understand the effect of edge-delay heterogeneity on individual regrets of agents.

Doubly adversarial federated bandits We studied doubly adversarial federated bandits, a new adversarial (non-stochastic) setting for federated bandits, which complement prior studies on stochastic federated bandits. Firstly, we derived regret lower bounds for any federated bandit algorithm when the agents have access to full-information or bandit feedback. These regret lower bounds relate the hardness of the problem to the algebraic connectivity of the network through which the agents communicate. Then we proposed the FEDEXP3 algorithm which is a federated version of the Exp3 algorithm. We showed that there is only a small polynomial gap between the regret upper bound of FEDEXP3 and the lower bound. Numerical experiments performed by using both

synthetic and real-word datasets demonstrated that FEDEXP3 can outperform the state-of-the-art stochastic federated bandit algorithm by a significant margin in non-stationary environments.

We point out some interesting avenues for future research on doubly adversarial federated bandits. The first is to close the gap between the regret upper bound of FEDEXP3 algorithm and the lower bounds shown in this chapter. The second is to extend the doubly adversarial assumption to federated linear bandit problems, where the doubly adversarial assumption could replace the stochastic assumption on the noise in the linear model.

Communication-regret trade-off in federated online learning In this study, we examine the online convex optimization in a federated learning setting, where multiple agents collaborate to minimize a convex objective function composed of convex functions distributed over a T time horizon and N agents. To this end, we propose a novel algorithm called FEDOCO, which offers a communication-efficient solution with a cumulative regret upper bound of $o(T)$ and communication complexity of $o(\sqrt{T})$ with high probability.

As future research avenues, we aim to determine the lower bound of communication complexity for algorithms that ensure sub-linear regret upper bounds in federated online convex optimization problems. Additionally, we plan to extend our stochastic communication protocol to more challenging non-stochastic multi-agent multi-armed bandit problems.

Greedy Bayes incremental matching strategies We discuss the sequential learning problem of a maximum value matching on a set fixed of nodes with different value functions. We have identified greedy Bayes incremental matching strategies for Boolean value functions AND and OR, and analysed their regret. Our analysis reveals how regret depends on the population size and its profile with respect to node types. For AND function, the dynamics of posterior beliefs can be represented by a stochastic process of matching sets which are merged or dissolved under greedy matching strategy.

Future work may consider the problem under different relaxations of assumptions, for different domains of latent node variables, noisy observations of match values, and other definitions of incremental matchings that allow for a larger number of re-matches per round.

References

- Mridul Agarwal, Vaneet Aggarwal, and Kamyar Azizzadenesheli. 2022. Multi-Agent Multi-Armed Bandits with Limited Communication. *Journal of Machine Learning Research* 23, 212 (2022), 1–24.
- Mridul Agarwal, Vaneet Aggarwal, Christopher J. Quinn, and Abhishek K. Umrawal. 2021. Stochastic Top-K Subset Bandits with Linear Space and Non-Linear Feedback. In *Proceedings of the 32nd International Conference on Algorithmic Learning Theory (Proceedings of Machine Learning Research, Vol. 132)*, Vitaly Feldman, Katrina Ligett, and Sivan Sabato (Eds.). PMLR, 306–339.
- Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. 2017. QSGD: Communication-efficient SGD via gradient quantization and encoding. *Advances in neural information processing systems* 30 (2017).
- Sanjeev Arora, Satish Rao, and Umesh Vazirani. 2009. Expander flows, geometric embeddings and graph partitioning. *J. ACM* 56, 2 (2009), 1–37.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. 2002a. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47 (2002), 235–256.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. 2002b. The nonstochastic multiarmed bandit problem. *SIAM journal on computing* 32, 1 (2002), 48–77.
- Baruch Awerbuch and Robert Kleinberg. 2008. Competitive collaborative learning. *J. Comput. System Sci.* 74, 8 (2008), 1271–1288.
- Mohammad Babaeizadeh, Iuri Frosio, Stephen Tyree, Jason Clemons, and Jan Kautz. 2016. Reinforcement Learning through Asynchronous Advantage Actor-Critic on a GPU. In *International Conference on Learning Representations*.
- Yogev Bar-On and Yishay Mansour. 2019. Individual regret in cooperative non-stochastic multi-armed bandits. *Advances in Neural Information Processing Systems* 32 (2019).
- Ranieri Baraglia, Patrizio Dazzi, Matteo Mordacchini, and Laura Ricci. 2013. A peer-to-peer recommender system for self-emerging user communities based on gossip overlays. *J. Comput. System Sci.* 79, 2 (2013), 291–308.

- Avrim Blum, John Hopcroft, and Ravindran Kannan. 2020. *Foundations of data science*. Cambridge University Press.
- Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloe Kiddon, Jakub Konečný, Stefano Mazzocchi, Brendan McMahan, et al. 2019. Towards federated learning at scale: System design. *Proceedings of machine learning and systems* 1 (2019), 374–388.
- Stephen Boyd, Persi Diaconis, and Lin Xiao. 2004. Fastest mixing Markov chain on a graph. *SIAM review* 46, 4 (2004), 667–689.
- G. Carlier and I. Ekeland. 2010. Matching for teams. *Economic Theory* 42, 2 (2010), 397–418.
- Nicolò Cesa-Bianchi, Tommaso Cesari, and Claire Monteleoni. 2020. Cooperative online learning: Keeping your neighbors updated. In *Algorithmic Learning Theory*. PMLR, 234–250.
- Nicolò Cesa-Bianchi, Yoav Freund, David Haussler, David P Helmbold, Robert E Schapire, and Manfred K Warmuth. 1997. How to use expert advice. *J. ACM* 44, 3 (1997), 427–485.
- Nicolò Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. 2019. Delay and Cooperation in Nonstochastic Bandits. *Journal of Machine Learning Research* 20, 17 (2019), 1–38.
- Nicolò Cesa-Bianchi, Claudio Gentile, Yishay Mansour, and Alberto Minora. 2016. Delay and cooperation in nonstochastic bandits. In *Conference on Learning Theory*. PMLR, 605–622.
- Nicolò Cesa-Bianchi and Gábor Lugosi. 2006. *Prediction, learning, and games*. Cambridge university press.
- Nicola Cesa-Bianchi and Gabor Lugosi. 2012. Combinatorial bandits. *J. Comput. System Sci.* 78, 5 (2012), 1404–1422.
- Wei Chen, Wei Hu, Fu Li, Jian Li, Yu Liu, and Pinyan Lu. 2016. Combinatorial Multi-Armed Bandit with General Reward Functions. In *Proceedings of the 30th International Conference on Neural Information Processing Systems (Barcelona, Spain) (NIPS'16)*. Curran Associates Inc., Red Hook, NY, USA, 1659–1667.

- Tianshu Chu, Jie Wang, Lara Codecà, and Zhaojian Li. 2019. Multi-agent deep reinforcement learning for large-scale traffic signal control. *IEEE Transactions on Intelligent Transportation Systems* 21, 3 (2019), 1086–1095.
- Richard Combes, M. Sadegh Talebi, Alexandre Proutiere, and Marc Lelarge. 2015. Combinatorial Bandits Revisited. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2* (Montreal, Canada). MIT Press, Cambridge, MA, USA, 2116–2124.
- Abhimanyu Dubey and AlexSandy’ Pentland. 2020. Differentially-private federated linear bandits. *Advances in Neural Information Processing Systems* 33 (2020), 6003–6014.
- John C Duchi, Alekh Agarwal, and Martin J Wainwright. 2011. Dual averaging for distributed optimization: Convergence analysis and network scaling. *IEEE Transactions on Automatic control* 57, 3 (2011), 592–606.
- Matthew Eichhorn, Siddhartha Banerjee, and David Kempe. 2022. Online Team Formation under Different Synergies. <https://doi.org/10.48550/ARXIV.2210.05795>
- Genevieve E Flaspohler, Francesco Orabona, Judah Cohen, Soukayna Mouatadid, Miruna Oprescu, Paulo Orenstein, and Lester Mackey. 2021. Online learning with optimism and delay. In *International Conference on Machine Learning*. PMLR, 3363–3373.
- Karol Flores-Szwagrzak and Rafael Treibich. 2020. Teamwork and Individual Productivity. *Management Science* 66, 6 (2020), 2523–2544.
- Thore Graepel, Tom Minka, and Ralf Herbrich. 2006. TrueSkill(tm): A Bayesian Skill Rating System. In *Proc. of NIPS 2006*, Vol. 19. 569–576.
- Aric Hagberg, Pieter Swart, and Daniel S Chult. 2008. *Exploring network structure, dynamics, and function using NetworkX*. Technical Report. Los Alamos National Lab.(LANL), Los Alamos, NM (United States).
- Jenny Hamer, Mehryar Mohri, and Ananda Theertha Suresh. 2020. Fedboost: A communication-efficient algorithm for federated learning. In *International Conference on Machine Learning*. PMLR, 3973–3983.

- Charles R Harris, K Jarrod Millman, Stéfan J Van Der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. 2020. Array programming with NumPy. *Nature* 585, 7825 (2020), 357–362.
- Elad Hazan et al. 2016. Introduction to online convex optimization. *Foundations and Trends® in Optimization* 2, 3-4 (2016), 157–325.
- Jean-Baptiste Hiriart-Urruty and Claude Lemarechal. 2010. *Convex Analysis and Minimization Algorithms II: Advanced Theory and Bundle Methods: 306* (softcover reprint of hardcover 1st ed. 1993 ed.). Springer, Berlin Heidelberg.
- Saghar Hosseini, Airlie Chapman, and Mehran Mesbahi. 2013. Online distributed optimization via dual averaging. In *52nd IEEE Conference on Decision and Control*. IEEE, 1484–1489.
- Ruiquan Huang, Weiqiang Wu, Jing Yang, and Cong Shen. 2021. Federated linear contextual bandits. *Advances in Neural Information Processing Systems* 34 (2021), 27057–27068.
- Ramesh Johari, Vijay Kamble, Anilesh K. Krishnaswamy, and Hannah Li. 2018. Exploration vs. Exploitation in Team Formation. In *Web and Internet Economics - 14th International Conference, WINE 2018, Oxford, UK, December 15-17, 2018, Proceedings (Lecture Notes in Computer Science, Vol. 11316)*, George Christodoulou and Tobias Harks (Eds.). Springer, 451.
- Pooria Joulani, Andras Gyorgy, and Csaba Szepesvári. 2013. Online learning under delayed feedback. In *International Conference on Machine Learning*. PMLR, 1453–1461.
- Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. 2021. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning* 14, 1–2 (2021), 1–210.
- Vijay Kamble, Anilesh K. Krishnaswamy, and Hannah Li. 2018. Exploration vs. Exploitation in Team Formation. In *WINE*.
- Soumya Kar, H Vincent Poor, and Shuguang Cui. 2011. Bandit problems in networks: Asymptotically efficient distributed allocation rules. In *2011 50th IEEE*

- Conference on Decision and Control and European Control Conference*. IEEE, 1771–1778.
- Jon Kleinberg and Maithra Raghu. 2018. Team Performance with Test Scores. *ACM Trans. Econ. Comput.* 6, 3-4, Article 17 (Oct. 2018), 26 pages.
- Ravi Kumar Kolla, Krishna Jagannathan, and Aditya Gopalan. 2018. Collaborative learning of stochastic bandits over a social network. *IEEE/ACM Transactions on Networking* 26, 4 (2018), 1782–1795.
- Branislav Kveton, Zheng Wen, Azin Ashkan, and Csaba Szepesvari. 2015. Tight Regret Bounds for Stochastic Combinatorial Semi-Bandits. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, Guy Lebanon and S. V. N. Vishwanathan (Eds.), Vol. 38. 535–543.
- Peter Landgren, Vaibhav Srivastava, and Naomi Ehrich Leonard. 2016. Distributed cooperative decision-making in multiarmed bandits: Frequentist and bayesian algorithms. In *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 167–172.
- Tor Lattimore and Csaba Szepesvári. 2020. *Bandit Algorithms*. Cambridge University Press. <https://doi.org/10.1017/9781108571401>
- Dabeen Lee, Milan Vojnovic, and Se-Young Yun. 2022. Test Score Algorithms for Budgeted Stochastic Utility Maximization. *INFORMS Journal on Optimization* 0, 0 (2022), null.
- Lihong Li, Wei Chu, John Langford, and Robert E Schapire. 2010. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*. 661–670.
- Somchaya Liemhetcharat and Manuela Veloso. 2012. Modeling and Learning Synergy for Team Formation with Heterogeneous Agents. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 1 (Valencia, Spain) (AAMAS '12)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 365–374.
- Weichao Mao, Haoran Qiu, Chen Wang, Hubertus Franke, Zbigniew Kalbarczyk, Ravi Iyer, and Tamer Basar. 2022. A Mean-Field Game Approach to Cloud Resource Management with Function Approximation. In *Advances in Neural Information Processing Systems*.

- David Martínez-Rubio, Varun Kanade, and Patrick Rebeschini. 2019. Decentralized cooperative stochastic bandits. *Advances in Neural Information Processing Systems* 32 (2019).
- Aranyak Mehta, Uri Nadav, Alexandros Psomas, and Aviad Rubinfeld. 2020. Hitting the High Notes: Subset Selection for Maximizing Expected Order Statistics. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Michael Mitzenmacher and Eli Upfal. 2017. *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*. Cambridge university press.
- Yurii Nesterov. 2009. Primal-dual subgradient methods for convex problems. *Mathematical programming* 120, 1 (2009), 221–259.
- George A. Neuman, Stephen H. Wagner, and Neil D. Christiansen. 1999. The Relationship between Work-Team Personality Composition and the Job Performance of Teams. *Group & Organization Management* 24, 1 (1999), 28–45.
- Mathew Penrose. 2003. *Random geometric graphs*. Vol. 5. OUP Oxford.
- Arun Rajkumar, Koyel Mukherjee, and Theja Tulabandhula. 2017. Learning to Partition Using Score Based Compatibilities. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS '17)*. 574–582.
- Clémence Réda, Sattar Vakili, and Emilie Kaufmann. 2022. Near-Optimal Collaborative Learning in Bandits. In *Advances in Neural Information Processing Systems*.
- Amirhossein Reisizadeh, Aryan Mokhtari, Hamed Hassani, Ali Jadbabaie, and Ramtin Pedarsani. 2020. Fedpaq: A communication-efficient federated learning method with periodic averaging and quantization. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 2021–2031.
- Idan Rejwan and Yishay Mansour. 2020. Top- k Combinatorial Bandits with Full-Bandit Feedback. In *Proceedings of the 31st International Conference on Algorithmic Learning Theory (Proceedings of Machine Learning Research, Vol. 117)*, Aryeh Kon-torovich and Gergely Neu (Eds.). PMLR, 752–776.

- Daniel Rothchild, Ashwinee Panda, Enayat Ullah, Nikita Ivkin, Ion Stoica, Vladimir Braverman, Joseph Gonzalez, and Raman Arora. 2020. Fetchsgd: Communication-efficient federated learning with sketching. In *International Conference on Machine Learning*. PMLR, 8253–8265.
- Kevin Scaman, Francis Bach, Sébastien Bubeck, Yin Lee, and Laurent Massoulié. 2019. Optimal convergence rates for convex distributed optimization in networks. *Journal of Machine Learning Research* 20 (2019), 1–31.
- Frank Seide, Hao Fu, Jasha Droppo, Gang Li, and Dong Yu. 2014. 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech dnns. In *Fifteenth annual conference of the international speech communication association*.
- Shreyas Sekar, Milan Vojnovic, and Se-Young Yun. 2021. A Test Score-Based Approach to Stochastic Submodular Optimization. *Management Science* 67, 2 (2021), 1075–1092.
- Flore Sentenac, Jialin Yi, Clement Calauzenes, Vianney Perchet, and Milan Vojnovic. 2021. Pure Exploration and Regret Minimization in Matching Bandits. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 9434–9442.
- Ohad Shamir. 2014. Fundamental limits of online and distributed algorithms for statistical learning and estimation. *Advances in Neural Information Processing Systems* 27 (2014).
- Chengshuai Shi, Cong Shen, and Jing Yang. 2021. Federated multi-armed bandits with personalization. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 2917–2925.
- Max Simchowitz, Kevin Jamieson, and Benjamin Recht. 2016. Best-of-K-bandits. In *29th Annual Conference on Learning Theory*, Vitaly Feldman, Alexander Rakhlin, and Ohad Shamir (Eds.), Vol. 49. 1440–1489.
- Virginia Smith, Simone Forte, Ma Chenxin, Martin Takáč, Michael I Jordan, and Martin Jaggi. 2018. CoCoA: A general framework for communication-efficient distributed optimization. *Journal of Machine Learning Research* 18 (2018), 230.

- Nikko Strom. 2015. Scalable Distributed DNN Training Using Commodity GPU Cloud Computing. In *Sixteenth annual conference of the international speech communication association*.
- Balazs Szorenyi, Róbert Busa-Fekete, István Hegedus, Róbert Ormándi, Márk Jelasity, and Balázs Kégl. 2013. Gossip-based distributed stochastic bandit algorithms. In *International Conference on Machine Learning*. PMLR, 19–27.
- Tobias Sommer Thune, Nicolò Cesa-Bianchi, and Yevgeny Seldin. 2019. Non-stochastic multiarmed bandits with unrestricted delays. *Advances in Neural Information Processing Systems* 32 (2019).
- Constantino Tsallis. 1988. Possible generalization of Boltzmann-Gibbs statistics. *Journal of statistical physics* 52, 1 (1988), 479–487.
- Yogatheesan Varatharajah and Brent Berry. 2022. A Contextual-bandit-based approach for informed decision-making in clinical trials. *Life* 12, 8 (2022), 1277.
- Yuanyu Wan, Wei-Wei Tu, and Lijun Zhang. 2020. Projection-free Distributed Online Convex Optimization with $O(\sqrt{T})$ Communication Complexity. In *International Conference on Machine Learning*. PMLR, 9818–9828.
- Jianqiao Wangni, Jialei Wang, Ji Liu, and Tong Zhang. 2018. Gradient sparsification for communication-efficient distributed optimization. *Advances in Neural Information Processing Systems* 31 (2018).
- Nicholas C. Wormald. 1995. Differential Equations for Random Processes and Random Graphs. *Ann. Appl. Probab.* 5, 4 (11 1995), 1217–1235.
- Ruofan Wu, Junmin Zhong, Brent Abraham Wallace, Xiang Gao, He Huang, and Jennie Si. 2022. Human-Robotic Prosthesis as Collaborating Agents for Symmetrical Walking. In *Advances in Neural Information Processing Systems*.
- Lin Xiao. 2009. Dual averaging method for regularized stochastic learning and online optimization. *Advances in Neural Information Processing Systems* 22 (2009).
- Jialin Yi and Milan Vojnovic. 2022. On Regret-optimal Cooperative Nonstochastic Multi-armed Bandits. *arXiv preprint arXiv:2211.17154* (2022).
- Jialin Yi and Milan Vojnović. 2023. Doubly Adversarial Federated Bandits. *arXiv preprint arXiv:2301.09223* (2023).

- Jingwei Yi, Fangzhao Wu, Chuhan Wu, Ruixuan Liu, Guangzhong Sun, and Xing Xie. 2021. Efficient-FedRec: Efficient Federated Learning Framework for Privacy-Preserving News Recommendation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 2814–2824.
- Wenpeng Zhang, Peilin Zhao, Wenwu Zhu, Steven CH Hoi, and Tong Zhang. 2017. Projection-free distributed online learning in networks. In *International Conference on Machine Learning*. PMLR, 4054–4062.
- Zhaowei Zhu, Jingxuan Zhu, Ji Liu, and Yang Liu. 2021. Federated bandit: A gossiping approach. In *Abstract Proceedings of the 2021 ACM SIGMETRICS/International Conference on Measurement and Modeling of Computer Systems*. 3–4.
- Julian Zimmert and Yevgeny Seldin. 2020. An optimal algorithm for adversarial bandits with arbitrary delays. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 3285–3294.
- Julian Zimmert and Yevgeny Seldin. 2021. Tsallis-INF: An Optimal Algorithm for Stochastic and Adversarial Bandits. *J. Mach. Learn. Res.* 22, 28 (2021), 1–49.