
Multi-layered Network Exploration via Random Walks: From Offline Optimization to Online Learning

Xutong Liu¹ Jinhang Zuo² Xiaowei Chen³ Wei Chen⁴ John C.S. Lui¹

Abstract

Multi-layered network exploration (MuLaNE) problem is an important problem abstracted from many applications. In MuLaNE, there are multiple network layers where each node has an importance weight and each layer is explored by a random walk. The MuLaNE task is to allocate total random walk budget B into each network layer so that the total weights of the unique nodes visited by random walks are maximized. We systematically study this problem from offline optimization to online learning. For the offline optimization setting where the network structure and node weights are known, we provide greedy based constant-ratio approximation algorithms for overlapping networks, and greedy or dynamic-programming based optimal solutions for non-overlapping networks. For the online learning setting, neither the network structure nor the node weights are known initially. We adapt the combinatorial multi-armed bandit framework and design algorithms to learn random walk related parameters and node weights while optimizing the budget allocation in multiple rounds, and prove that they achieve logarithmic regret bounds. Finally, we conduct experiments on a real-world social network dataset to validate our theoretical results.

1. Introduction

Network exploration is a fundamental paradigm of searching/exploring in order to discover information and resources available at nodes in a network, and random walk is often

used as an effective tool for network exploration (Lv et al., 2002; Gleich, 2015; Wilder et al., 2018). In this paper, we study the multi-layered network exploration via random walks problem, which can model many real-world applications, including resource searching in peer-to-peer (P2P) networks and web surfing in online social networks (OSNs).

In P2P networks, a user wants to find resources that are scattered on nodes (i.e. peers) in different P2P networks via some platform-specified strategies. A commonly used search strategy is based on multiple random walks (Lv et al., 2002). Since different resources have different importance, the search quality of different random walkers varies. Moreover, the resource-search process typically has a life span, i.e., a total time-limit or hop-limit for random walks. So the user’s goal is to decide how to allocate limited budgets to different random walkers to find as many important resources as possible. Another application is the web surfing, where users want to find information in different OSNs, by looking at posts via others’ home-pages (Lerman & Jones, 2006). Once the user arrives at one of his friends’ home-pages in a particular OSN, he could find some information and continues to browse the home-pages of his friends’ friends, which can be regarded as a random walk process. Since the user only has a finite duration in web surfing, the similar question is how to allocate his time in exploring different OSNs so to get the maximum amount of useful information.

We abstract the above application scenarios as the **Multi-Layered Network Exploration** (MuLaNE) problem. In MuLaNE, we model the overall network to be explored (e.g., combining different OSNs) as a multi-layered network \mathcal{G} that consists of m layers L_1, \dots, L_m . Each layer L_i (e.g., a single OSN) is represented by a weighted directed graph $\mathcal{G}_i(\mathcal{V}_i, \mathcal{E}_i, w_i)$, where \mathcal{V}_i is the set of nodes to be explored (e.g., users’ home-pages with importance weight σ_u for $u \in \mathcal{V}_i$), $\mathcal{E}_i \subseteq \mathcal{V}_i \times \mathcal{V}_i$ is the set of directed edges (e.g., social links), and w_i is the edge weight function on edges \mathcal{E}_i . Each layer L_i is associated with an explorer (or random walker) W_i and a fixed starting distribution α_i on nodes \mathcal{V}_i . Explorer W_i starts on a node in \mathcal{G}_i following the distribution α_i , and then walks on network \mathcal{G}_i following outgoing edges with probability proportional to edge weights. We assume that each random walk step will cost one unit of the bud-

¹Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong SAR, China ²Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, USA ³Bytedance, Mountain View, CA, USA ⁴Microsoft Research, Beijing, China. Correspondence to: Xutong Liu <liuxt@cse.cuhk.edu.hk>, Wei Chen <weic@microsoft.com>, John C.S. Lui <cslui@cse.cuhk.edu.hk>.

get¹. Given a total budget constraint B and individual layer budget constraint c_i , the MuLaNE task is to find the optimal budget allocation $\mathbf{k}=(k_1, \dots, k_m)$ with $\sum_{i=1}^m k_i \leq B$, $0 \leq k_i \leq c_i$ such that the expected total weights of *unique* nodes visited are maximized.

In real applications, the network structure \mathcal{G} , the node weights σ and the starting distributions α_i 's may not be known in advance, so we consider both the *offline optimization* cases when \mathcal{G} , σ and α_i 's are known and the *online learning* case when \mathcal{G} , σ and α_i 's are unknown. Moreover, different layers may have overlapping vertices (e.g., homepage of the same user may appear in different OSNs), and the starting distributions α_i 's may or may not be stationary distributions. In this paper, we provide a systematic study of all these case combinations.

For the offline optimization, we first consider the general overlapping setting and provide constant approximation algorithms based on the lattice submodularity property for non-stationary α_i 's and the diminishing-return (DR) submodularity property for stationary α_i 's. For the special non-overlapping setting, we design a dynamic programming algorithm that finds the exact optimal solution for MuLaNE.

In the online learning, we conduct multiple rounds of exploration, each of which has the same budget constraints B and c_i 's. After each round of exploration, the total weights of unique nodes visited in this round are the reward for this round, and the trajectory of every explorer in every layer and the importance weights of visited nodes are observed as the feedback, which could be used to learn information about the network for the benefit of future explorations.

We adapt the combinatorial multi-armed bandit (CMAB) framework and the CUCB algorithm (Chen et al., 2016) to our setting, and design online learning algorithms to minimize the regret, which is the difference between the cumulative reward achieved by the exact or approximate offline algorithm and that achieved by our learning algorithm, over T rounds. We show that directly learning the graph structure \mathcal{G} is inefficient. Instead, we define intermediate random variables in MuLaNE as the base arms in CMAB and learn these intermediate parameters, which can sufficiently determine the rewards to guide our budget allocation. Moreover, the node weight is not revealed until this node is first visited, which further complicates the design of online exploration algorithms. For the overlapping case, we adapt the CUCB algorithm to address the unrevealed node weights and the extra constraint of monotonicity for the intermediate parameters. We further improve the analysis by leveraging on special properties in our setting and show logarithmic regret bounds in this case. For the non-overlapping case,

¹Our model can be extended to move multiple steps with one unit of budget.

we define more efficient intermediate parameters and use the exact offline algorithm, and thus achieve a better regret bound. Finally, we conduct experiments on a real-world multi-layered social network dataset to validate the effectiveness of both our offline and online algorithms.

Our contributions can be summarized as follows: (1) We are the first to model the multi-layered network exploration via random walks problem (MuLaNE) as an abstraction for many real-world applications (2) We provide a systematic study of the MuLaNE problem via theoretical analysis and empirical validation by considering offline and online settings for both overlapping and non-overlapping multi-layered networks. Due to the space limit, proofs are included in the supplementary material.

1.1. Related Work

Network exploration via random walks has been studied in various application contexts such as community detection (Pons & Latapy, 2005), centrality measuring (Gleich, 2015), large-scale network sampling (Li et al., 2019), and influence maximization (Wilder et al., 2018). Multiple random walks has also been used, e.g. in (Lv et al., 2002) for query resolution in peer-to-peer networks. However, none of these studies address the budgeted MuLaNE problem. MuLaNE is also related to the influence maximization (IM) problem (Kempe et al., 2003; Chen et al., 2009) and can be viewed as a special variant of IM. Different from the standard Independent Cascade (IC) model (Wang et al., 2012) in IM, which randomly broadcasts to all its neighbors, the propagation process of MuLaNE is a random walk that selects one neighbor. Moreover, each unit of budget in MuLaNE is used to propagate one random-walk step, not to select one seed node as in IC.

The offline budget allocation problem has been studied by Alon et al. (2012); Soma et al. (2014), the latter of which propose the lattice submodularity and a constant approximation algorithm based on this property. Our offline overlapping MuLaNE setting is based on a similar approach, but we provide a better approximation ratio compared to the original analysis in (Alon et al., 2012). Moreover, we further show that the stationary setting enjoys DR-submodularity, leading to an efficient algorithm with a better approximation ratio.

The multi-armed bandit (MAB) problem is first studied by Robbins (1952) and then extended by many studies (cf. (Bubeck & Cesa-Bianchi, 2012)). Our online MuLaNE setting fits into the general Combinatorial MAB (CMAB) framework of (Gai et al., 2012; Chen et al., 2016). CUCB is proposed as a general algorithm for CMAB (Chen et al., 2016). Our study includes several adaptations, such as handling unknown node weights, defining intermediate random variables as base arms and so on.

Chen et al. (2018) study the community exploration problem,

which is essentially a special case of MuLaNE with non-overlapping complete graphs. As a result, their technique is different and much simpler than ours.

The rest of the paper is organized as follows. Section 2 states the settings of MuLaNE; Section 3 states the equivalent bipartite coverage model to derive the explicit form for our reward function; Section 4 states offline algorithms for four offline settings with provable approximation guarantee and running time analysis; Section 5 states two online learning algorithms with regret analysis; Empirical results are shown in Section 6; and Section 7 concludes the paper.

2. Problem Settings

Basic Notations. In this paper, we use \mathbb{R} and \mathbb{Z} to denote the sets of real numbers and integers, respectively, and the associated subscript ≥ 0 and > 0 denote their non-negative and positive subsets, respectively. Suppose we want to explore a multi-layered network $\mathcal{G}(\mathcal{V}, \mathcal{E}, \boldsymbol{\sigma})$, consisting of m layers L_1, \dots, L_m and N unique nodes \mathcal{V} with fixed node weights $\boldsymbol{\sigma} \in [0, 1]^{\mathcal{V}}$. Let $[m]$ denote the set $\{1, 2, \dots, m\}$. Each layer $L_i, i \in [m]$, represents a subgraph of \mathcal{G} we could explore and is modeled as a weighted digraph $\mathcal{G}_i(\mathcal{V}_i, \mathcal{E}_i; w_i)$, where $\mathcal{V}_i \subseteq \mathcal{V}$ is the set of vertices in layer L_i , $\mathcal{E}_i \subseteq \mathcal{V}_i \times \mathcal{V}_i$ is the set of edges in \mathcal{G}_i , and $w_i : \mathcal{E}_i \rightarrow \mathbb{R}_{>0}$ is the edge weight function associating each edge of \mathcal{G}_i . For any two layers L_i and L_j , we say they are *non-overlapping* if $\mathcal{V}_i \cap \mathcal{V}_j = \emptyset$. \mathcal{G} is called a *non-overlapping* multi-layered network if any two layers are non-overlapping; otherwise \mathcal{G} is an *overlapping* multi-layered network. Let $n_i = |\mathcal{V}_i|$ denote the size of layer L_i , and the adjacency matrix of \mathcal{G}_i is defined as $\mathbf{A}_i \in \mathbb{R}_{\geq 0}^{n_i \times n_i}$, where $\mathbf{A}_i[u, v] = w_i((u, v))$ if $(u, v) \in \mathcal{E}_i$ and 0 otherwise.

Exploration Rule. Each layer L_i is associated with an explorer W_i , a budget $k_i \in \mathbb{Z}_{\geq 0}$ and an initial starting distribution $\boldsymbol{\alpha}_i = (\alpha_{i,u})_{u \in \mathcal{V}_i} \in \mathbb{R}_{\geq 0}^{n_i}$ with $\sum_{u \in \mathcal{V}_i} \alpha_{i,u} = 1$. The exploration process of W_i is identical to applying *weighted random walks* within layer L_i . Specifically, the explorer W_i starts the exploration from a random node $v_1 \in \mathcal{V}_i$ with probability α_{i,v_1} ; it consumes one unit of budget and continues the exploration by walking from v_1 to one of its out-neighbors, say v_2 , with probability $P_i[v_1, v_2]$, where $P_i \in \mathbb{R}_{\geq 0}^{n_i \times n_i}$ is the transition probability matrix, and $P_i[u, v] := \mathbf{A}_i[u, v] / (\sum_{w: (u,w) \in \mathcal{E}_i} \mathbf{A}_i[u, w])$. Consider visiting the initial node v_1 as the first step, the process is repeated until the explorer W_i walks k_i steps and visits k_i nodes (with possibly duplicated nodes). Note that one can easily generalize our results by moving $\lambda_i \in \mathbb{Z}_{>0}$ steps with one unit of budget for W_i , and for simplicity, we set $\lambda_i = 1$.

Reward Function. We define the exploration trajectory for explorer W_i after exploring k_i nodes as $\Phi(i, k_i) := (X_{i,1}, \dots, X_{i,k_i})$, where $X_{i,j} \in \mathcal{V}_i$ denotes the node visited

at j -th step, and $\Pr(X_{i,1}=u) = \alpha_{i,u}, u \in \mathcal{V}_i$. The reward for $\Phi(i, k_i)$ is defined as the total weights of *unique* nodes visited by W_i , i.e., $\sum_{v \in \cup_{j=1}^{k_i} \{X_{i,j}\}} \sigma_v$. Considering trajectories of all W_i 's, the total reward is the total weights of *unique* nodes visited by all random walkers, i.e., $\sum_{v \in \cup_{i=1}^m \cup_{j=1}^{k_i} \{X_{i,j}\}} \sigma_v$. For notational simplicity, let $\mathcal{G} := (\mathcal{G}_1, \dots, \mathcal{G}_m)$, $\boldsymbol{\alpha} := (\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_m)$ and $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_{|\mathcal{V}|})$ be the parameters of a problem instance and $\mathbf{k} := (k_1, \dots, k_m)$ be the allocation vector. For overlapping multi-layered network \mathcal{G} , the total expected reward is

$$r_{\mathcal{G}, \boldsymbol{\alpha}, \boldsymbol{\sigma}}(\mathbf{k}) := \mathbb{E}_{\Phi(1, k_1), \dots, \Phi(m, k_m)} \left[\sum_{v \in \cup_{i=1}^m \cup_{j=1}^{k_i} \{X_{i,j}\}} \sigma_v \right], \quad (1)$$

where its explicit formula will be discussed later in Sec. 3. For non-overlapping \mathcal{G} , the reward function can be simplified and written as the summation over separated layers,

$$r_{\mathcal{G}, \boldsymbol{\alpha}, \boldsymbol{\sigma}}(\mathbf{k}) := \sum_{i=1}^m \mathbb{E}_{\Phi(i, k_i)} \left[\sum_{v \in \cup_{j=1}^{k_i} \{X_{i,j}\}} \sigma_v \right]. \quad (2)$$

Problem Formulation. We are interested in the *budget allocation problem*: how to allocate the total budget $B \in \mathbb{Z}_{\geq 0}$ to the m explorers so as to maximize the total weights of unique nodes visited. Also, we assume there is a budget constraint $\mathbf{c} := (c_1, \dots, c_m)$ such that the allocated budget k_i should not exceed c_i , i.e., $k_i \leq c_i$, for $i \in [m]$.

Definition 1. Given graph structures $\mathcal{G} := (\mathcal{G}_1, \dots, \mathcal{G}_m)$, starting distributions $\boldsymbol{\alpha} := (\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_m)$, total budget B and budget constraints $\mathbf{c} := (c_1, \dots, c_m)$, the **Multi-Layered Network Exploration problem**, denoted as **MuLaNE**, is formulated as the following optimization problem,

$$\max_{\mathbf{k}} r_{\mathcal{G}, \boldsymbol{\alpha}, \boldsymbol{\sigma}}(\mathbf{k}) \text{ s.t. } \mathbf{k} \in \mathbb{Z}_{\geq 0}^m \leq \mathbf{c}, \sum_{i=1}^m k_i \leq B, \quad (3)$$

where $r_{\mathcal{G}, \boldsymbol{\alpha}, \boldsymbol{\sigma}}(\mathbf{k})$ is given by Eq. (1) or Eq. (2). We also need to consider the following settings:

Offline Setting. For the offline setting, all problem instance parameters $(\mathcal{G}, \boldsymbol{\alpha}, \boldsymbol{\sigma}, \mathbf{c}, B)$ are given, and we aim to find the optimal budget allocation \mathbf{k}^* determined by Eq. (3).

Online Setting. For the online setting, we consider T -round explorations. Before the exploration, we only know an upper bound of the total number of nodes in \mathcal{G} and the number of layers m ,² but we do not know about the network structure \mathcal{G} , the starting distributions $\boldsymbol{\alpha}$ or node weights $\boldsymbol{\sigma}$. In round $t \in [T]$, we choose the budget allocation $\mathbf{k}_t := (k_{t,1}, \dots, k_{t,m})$ only based on observations from previous rounds, where $k_{t,i} \in \mathbb{Z}_{\geq 0}$ is the budget allocated to the i -th layer and $\sum_{i=1}^m k_{t,i} \leq B, \mathbf{k}_t \leq \mathbf{c}$. Define \mathbf{k}_t as the *action* taken in round t . By taking the action \mathbf{k}_t we mean that we interact with the environment and the random explorer W_i , which is part of the environment, would

²More precisely we only need to know the number of independent explorers. If two explorers explore on the same layer, it is equivalent as two layers with identical graph structures.

explore $k_{t,i}$ steps and generate the exploration trajectory $\Phi(i, k_{t,i}) = (X_{i,1}, \dots, X_{i,k_{t,i}})$. After we take the action \mathbf{k}_t , the exploration trajectory $\Phi(i, k_{t,i})$ for each layer L_i as well as the fixed importance weight σ_u of $u \in \Phi(i, k_{t,i})$ is revealed as the *feedback*³, which we leverage on to learn parameters related to the graph structures \mathcal{G} , the starting distribution α and the node weights σ , so that we can select better actions in future rounds. The reward we gain in round t is the total weights of unique nodes visited by all random explorers in round t . Our goal is to design an efficient online learning algorithm A to give us guidance on taking actions and gain as much cumulative reward as possible in T rounds.

In general, the online algorithm has to deal with the exploration-exploitation tradeoff. The cumulative regret is a commonly used metric to evaluate the performance of an online learning algorithm A . Formally, the T -round $((\xi, \beta)$ -approximation) regret of A is:

$$\text{Reg}_{\mathcal{G}, \alpha, \sigma}(T) = \xi \beta T \cdot r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k}^*) - \mathbb{E} \left[\sum_{t=1}^T r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k}_t^A) \right], \quad (4)$$

where $r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k}^*)$ is the reward value for the optimal budget allocation \mathbf{k}^* , \mathbf{k}_t^A is budget allocation selected by the learning algorithm A in round t , the expectation is taken over the randomness of the algorithm and the exploration trajectories in all T rounds, and (ξ, β) is the approximation guarantee of the offline oracle as explained below. Similar to other online learning frameworks (Chen et al., 2016; Wang & Chen, 2017), we assume that the online learning algorithm has access to an offline (ξ, β) -approximation oracle, which for problem instance $(\mathcal{G}, \alpha, \sigma, c, B)$ outputs an action \mathbf{k} such that $\Pr(r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k}) \geq \xi \cdot r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k}^*)) \geq \beta$. We also remark that the actual oracle we use takes certain intermediate parameters as inputs instead of \mathcal{G} and α .

Submodularity and DR-Submodularity Over Integer lattices. To solve the MuLaNE problem, we leverage on the submodular and DR-submodular properties of the reward function. For any $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_{\geq 0}^m$, we denote $\mathbf{x} \vee \mathbf{y}, \mathbf{x} \wedge \mathbf{y} \in \mathbb{Z}_{\geq 0}^m$ as the coordinate-wise maximum and minimum of these two vectors, i.e., $(\mathbf{x} \vee \mathbf{y})_i = \max\{x_i, y_i\}, (\mathbf{x} \wedge \mathbf{y})_i = \min\{x_i, y_i\}$. We define a function $f : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R}$ over the integer lattice $\mathbb{Z}_{\geq 0}^m$ as a *submodular* function if the following inequality holds for any $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_{\geq 0}^m$:

$$f(\mathbf{x} \vee \mathbf{y}) + f(\mathbf{x} \wedge \mathbf{y}) \leq f(\mathbf{x}) + f(\mathbf{y}). \quad (5)$$

Let $\text{supp}^+(\mathbf{x} - \mathbf{y})$ denote the set $I = \{i \in [m] : x_i > y_i\}$, $\chi_i = (0, \dots, 1, \dots, 0)$ be the one-hot vector whose i -th element is 1 and 0 otherwise, and $\mathbf{x} \geq \mathbf{y}$ means $x_i \geq y_i$ for all $i \in [m]$. We define a function $f : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R}$ as a *DR-submodular* (diminishing return submodular) function

³We further consider random node weights with unknown mean vector σ , see the Appendix F.3 for details.

if the following inequality holds for any $\mathbf{x} \leq \mathbf{y}$ and $i \in [m]$,

$$f(\mathbf{y} + \chi_i) - f(\mathbf{y}) \leq f(\mathbf{x} + \chi_i) - f(\mathbf{x}). \quad (6)$$

We say a function f is *monotone* if for any $\mathbf{x} \leq \mathbf{y}$, $f(\mathbf{x}) \leq f(\mathbf{y})$. Note that for a function $f : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R}$, submodularity does not imply DR-submodularity over integer lattices. In fact, the former is weaker than the latter, that is, a DR-submodular function is always a submodular function, but not vice versa. However, for a typical submodular function $f : \{0, 1\}^m \rightarrow \mathbb{R}$ defined on a set, they are equivalent.

3. Equivalent Bipartite Coverage Model

In order to derive the explicit formulation of the reward function $r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k})$ in Eq. (1), we construct an undirected bipartite coverage graph $\mathcal{B}(\mathcal{W}, \mathcal{V}, \mathcal{E}')$, where $\mathcal{W} = \{W_1, \dots, W_m\}$ denotes m random explorers, $\mathcal{V} = \bigcup_{i \in [m]} \mathcal{V}_i$ denotes all possible distinct nodes to be explored in \mathcal{G} , and the edge set $\mathcal{E}' = \{(W_i, u) | u \in \mathcal{V}_i, i \in [m]\}$ indicating whether node u could be visited by W_i . For each edge $(W_i, u) \in \mathcal{E}'$, we associate it with $c_i + 1$ visiting probabilities denoted as $P_{i,u}(k_i)$ for $k_i \in \{0\} \cup [c_i]$. $P_{i,u}(k_i)$ represents the probability that the node u is visited by the random walker W_i given the budget k_i , i.e., $P_{i,u}(k_i) = \Pr(u \in \Phi(i, k_i))$. Since W_i can never visit the node outside the i -th layer (i.e., $u \notin \mathcal{G}_i$), we set $P_{i,u}(k_i) = 0$ if $(W_i, u) \notin \mathcal{E}'$. Then, given budget allocation \mathbf{k} , the probability of a node u visited by at least one random explorer is $\Pr(u, \mathbf{k}) = 1 - \prod_{i \in [m]} (1 - P_{i,u}(k_i))$ because each W_i has the independent probability $P_{i,u}(k_i)$ to visit u .

By summing over all possible nodes, the reward function is

$$r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k}) = \sum_{u \in \mathcal{V}} \sigma_u \left(1 - \prod_{i \in [m]} (1 - P_{i,u}(k_i)) \right) \quad (7)$$

According to Eq. (2), for non-overlapping multi-layered network, we can rewrite the reward function as:

$$r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k}) = \sum_{i \in [m]} \sum_{u \in \mathcal{V}_i} \sigma_u P_{i,u}(k_i). \quad (8)$$

Remark. The bipartite coverage model is needed to integrate the graph structure and the random walk exploration mechanisms into $P_{i,u}(k_i)$ to determine the reward function. It also handles the scenario where each layer is explored by multiple random walkers (see Appendix A).

3.1. Properties of the visiting probability $P_{i,u}(k_i)$

The quantities $P_{i,u}(k_i)$'s in Eq. (7) and (8) are crucial for both our offline and online algorithms, and thus we provide their analytical formulas and properties here. In order to analyze the property of the reward function, we apply the absorbing Markov Chain technique to calculate $P_{i,u}(k_i)$. For $u \notin \mathcal{G}_i$, $P_{i,u}(k_i) = 0$; For $u \in \mathcal{G}_i$, we create an absorbing Markov Chain $\mathbf{P}_i(\mathbf{u}) \in \mathbb{R}^{n_i \times n_i}$ by setting the target

node u as the absorbing node. We derive the corresponding transition matrix $P_i(\mathbf{u})$ as $P_i(\mathbf{u})[v, \cdot] = \chi_u^\top$ if $v = u$ and $P_i(\mathbf{u})[v, \cdot] = P_i[v, \cdot]$ otherwise, where $P_i[v, \cdot]$ denotes the row vector corresponding to the node v of P_i , and $\chi_u = (0, \dots, 0, 1, 0, \dots, 0)^\top$ denotes the one-hot vector with 1 at the u -th entry and 0 elsewhere.⁴ Intuitively, $P_i(\mathbf{u})$ corresponds to the transition probability matrix of \mathcal{G}_i after removing all out-edges of u and adding a self loop to itself in the original graph \mathcal{G}_i . The random walker W_i will be trapped in u if it ever visits u . We observe that $P_{i,u}(k_i)$ equals to the probability the random walker stays in the absorbing node u at step $k_i \in \mathbb{Z}_{>0}$ (trivially, $P_{i,u}(0) = 0$),

$$P_{i,u}(k_i) = \alpha_i^\top P_i(\mathbf{u})^{k_i-1} \chi_u. \quad (9)$$

Then define the marginal gain of $P_{i,u}(\cdot)$ at step $k_i \geq 1$ as,

$$g_{i,u}(k_i) = P_{i,u}(k_i) - P_{i,u}(k_i - 1). \quad (10)$$

The physical meaning of $g_{i,u}(k_i)$ is the probability that node u is visited exactly at the k_i -th step and not visited before k_i . Now, we can show $g_{i,u}(k_i)$ is non-negative,

Lemma 1. $g_{i,u}(k_i) \geq 0$ for any $i \in [m]$, $u \in \mathcal{V}_i$, $k_i \in \mathbb{Z}_{>0}$.

This means $P_{i,u}(k_i)$ is non-decreasing with respect to step k_i . In other words, the more budgets we allocate to W_i , the higher probability W_i can visit u in k_i steps.

Starting from arbitrary distributions: Although $P_{i,u}(k_i)$ is monotone (non-decreasing) w.r.t k_i , $g_{i,u}(k_i)$ may not be monotonic non-increasing under arbitrary starting distributions. Intuitively, there may exist one critical step k_i such that $P_{i,u}(k_i)$ suddenly increases by a large value (see Appendix B.1). This means that $P_{i,u}(k_i)$ lacks the diminishing return (or “discretely concave”) property, which many problems rely on to provide good optimization (Kapralov et al., 2013; Soma & Yoshida, 2015). In other words, we are dealing with a more challenging non-concave optimization problem for the discrete budget allocation.

Starting from the stationary distribution: If our walkers start with the stationary distributions, the following property holds: $g_{i,u}(k_i)$ is non-increasing w.r.t k_i .

Lemma 2. $g_{i,u}(k_i + 1) - g_{i,u}(k_i) \leq 0$ for any $i \in [m]$, $u \in \mathcal{V}_i$, $k_i \in \mathbb{Z}_{>0}$, if $\alpha_i = \pi_i$, where $\pi_i^\top P_i = \pi_i^\top$.

4. Offline Optimization for MuLaNE

In this section, we first consider the general case, where layers are overlapping and starting distributions are arbitrary. Next, we consider the starting distribution is the stationary distribution, and give solutions with better solution quality and time complexity. Then we analyze special cases where layers are non-overlapping and give optimal solutions for

⁴When related to matrix operations, we treat all vectors as column vectors by default.

arbitrary distributions. The summary for offline models and algorithmic results are presented in Table 1.

4.1. Overlapping MuLaNE

Starting from arbitrary distributions. Based on the equivalent bipartite coverage model, one can observe that our problem formulation is a generalization of the Probabilistic Maximum Coverage (PMC) (Chen et al., 2016) problem, which is NP-hard and has a $(1 - 1/e)$ approximation based on submodular set function maximization. However, our problem is more general in that we want to select multi-sets from \mathcal{W} with budget constraints, and the reward function in general does not have the DR-submodular property for an arbitrary starting distribution. Nevertheless, we have the following lemma to solve our problem.

Lemma 3. For any network \mathcal{G} , distribution α and weights $\sigma, r_{\mathcal{G},\alpha,\sigma}(\cdot) : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R}$ is monotone and submodular.

Leveraging on the monotone submodular property, we design a Budget Effective Greedy algorithm (Alg. 1). The core of Alg. 1 is the BEG procedure. Let $\delta(i, b, \mathbf{k})$ be the per-unit marginal gain $(r_{\mathcal{G},\alpha,\sigma}(\mathbf{k} + b\chi_i) - r_{\mathcal{G},\alpha,\sigma}(\mathbf{k}))/b$ for allocating b more budgets to layer i , which equals to

$$\sum_{u \in \mathcal{V}} \sigma_u \prod_{j \neq i} (1 - P_{j,u}(k_j)) (P_{i,u}(k_i + b) - P_{i,u}(k_i)) / b. \quad (11)$$

BEG procedure consists of two parts and maintains a queue \mathcal{Q} , where any pair $(i, b) \in \mathcal{Q}$ represents a tentative plan of allocating b more budgets to layer i . The first part is built around the while loop (line 6-10), where each iteration greedily selects the (i, b) pair in \mathcal{Q} such that the per-unit marginal gain $\delta(i, b, \mathbf{k})$ is maximized. The second part is a for loop (line 12-13), where in the i -th round we attempt to allocate all c_i budgets to layer i and replace the current best budget allocation if we have a larger reward.

Theorem 1. Algorithm 1 obtains a $(1 - e^{-\eta}) \approx 0.357$ -approximate solution, where η is the solution of equation $e^\eta = 2 - \eta$, to the overlapping MuLaNE problem.

Line 1 uses $O(m \|c\|_\infty n_{max}^3)$ time to pre-calculate visiting probabilities based on Eq. (9), where $n_{max} = \max_i |\mathcal{V}_i|$. In the BEG procedure, the while loop contains B iterations, the size of queue \mathcal{Q} is $O(m \|c\|_\infty)$, and line 7 uses $O(n_{max})$ to calculate $\delta(i, b, \mathbf{k})$ by proper pre-computation and update (see Appendix C.3), thus the time complexity of Algorithm 1 is $O(B \|c\|_\infty m n_{max} + m \|c\|_\infty n_{max}^3)$.

Remark 1. The idea of combining the greedy algorithm with enumerating solutions on one layer is also adopted in previous works (Khuller et al., 1999; Alon et al., 2012), but they only give a $\frac{1}{2}(1 - e^{-1}) \approx 0.316$ -approximation analysis. In this paper, we provide a novel analysis with a better $(1 - e^{-\eta}) \approx 0.357$ -approximation, see Appendix C.1 for details.

Table 1. Summary of the offline models and algorithms.

Overlapping?	Starting distribution	Algorithm	Apprx ratio	Time complexity
✓	Arbitrary	Budget Effective Greedy	$(1 - e^{-\eta})^5$	$O(B \ c\ _\infty mn_{max} + \ c\ _\infty mn_{max}^3)$
✓	Stationary	Myopic Greedy	$(1 - 1/e)$	$O(Bmn_{max} + \ c\ _\infty mn_{max}^3)$
×	Arbitrary	Dynamic Programming	1	$O(B \ c\ _\infty m + \ c\ _\infty mn_{max}^3)$
×	Stationary	Myopic Greedy	1	$O(B \log m + \ c\ _\infty mn_{max}^3)$

Algorithm 1 Budget Effective Greedy (BEG) Algorithm for the Overlapping MuLaNE

Input: Network \mathcal{G} , starting distributions α , node weights σ , budget B , constraints c .
Output: Budget allocation \mathbf{k} .

- 1: Compute visiting probabilities $(P_{i,u}(b))_{i \in [m], u \in \mathcal{V}, b \in [c_i]}$ according to Eq. (9).
- 2: $\mathbf{k} \leftarrow \text{BEG}((P_{i,u}(b))_{i \in [m], u \in \mathcal{V}, b \in [c_i]}, \sigma, B, c)$.
- 3: **Procedure** BEG($(P_{i,u}(b))_{i \in [m], u \in \mathcal{V}, b \in [c_i]}, \sigma, B, c$)
- 4: Let $\mathbf{k} := (k_1, \dots, k_m) \leftarrow \mathbf{0}, K \leftarrow B$.
- 5: Let $\mathcal{Q} \leftarrow \{(i, b_i) \mid i \in [m], 1 \leq b_i \leq c_i\}$.
- 6: **while** $K > 0$ and $\mathcal{Q} \neq \emptyset$ **do**
- 7: $(i^*, b^*) \leftarrow \arg \max_{(i,b) \in \mathcal{Q}} \delta(i, b, \mathbf{k})/b \triangleright \text{Eq. (11)}$
- 8: $k_{i^*} \leftarrow k_{i^*} + b^*, K \leftarrow K - b^*$.
- 9: Modify all pairs $(i, b) \in \mathcal{Q}$ to $(i, b - b^*)$.
- 10: Remove all paris $(i, b) \in \mathcal{Q}$ such that $b \leq 0$.
- 11: **end while**
- 12: **for** $i \in [m]$ **do**
- 13: **if** $r_{\mathcal{G}, \alpha, \sigma}(c_i \chi_i) > r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k})$, **then** $\mathbf{k} \leftarrow c_i \chi_i$.
- 14: **end for**
- 15: **return** $\mathbf{k} := (k_1, \dots, k_m)$.
- 16: **end Procedure**

Remark 2. Another algorithm (Alon et al., 2012) with a better approximation ratio is to use partial enumeration techniques (i.e., BEGE), which can achieve $(1 - 1/e)$ approximation ratio. This is the best possible solution in polynomial time unless $P=NP$. But the time complexity is prohibitively high in $O(B^4 m^4 \|c\|_\infty n_{max} + m \|c\|_\infty n_{max}^3)$. The algorithm and the analysis are provided in the Appendix D.2.

Starting from the stationary distribution. We also consider the special case where each random explorer W_i starts from the stationary distribution π_i with $\pi_i^\top P = \pi_i^\top$. In this case, we have the following stronger DR-submodularity.

Lemma 4. For any network \mathcal{G} , stationary distributions π and node weights σ , function $r_{\mathcal{G}, \pi, \sigma}(\cdot) : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R}$ is monotone and DR-submodular.

Since the reward function is DR-submodular, the BEG procedure can be replaced by the simple MG procedure in Alg. 2 with a better approximation ratio. The time complexity is also improved to $O(Bmn_{max} + m \|c\|_\infty n_{max}^3)$.

Theorem 2. Algorithm 2 obtains a $(1 - 1/e)$ -approximate

Algorithm 2 Myopic Greedy (MG) Algorithm for MuLaNE

- 1: Same input, output and line 1-2 as in Alg. 1, except replacing BEG with MG procedure below.
- 2: **Procedure** MG($(P_{i,u}(b))_{i \in [m], u \in \mathcal{V}, b \in [c_i]}, \sigma, B, c$)
- 3: Let $\mathbf{k} := (k_1, \dots, k_m) \leftarrow \mathbf{0}, K \leftarrow B$.
- 4: **while** $K > 0$ **do**
- 5: $i^* \leftarrow \arg \max_{i \in [m], k_i + 1 \leq c_i} \delta(i, 1, \mathbf{k}) \triangleright \text{Eq. (11)}$
- 6: $k_{i^*} \leftarrow k_{i^*} + 1, K \leftarrow K - 1$.
- 7: **end while**
- 8: **return** $\mathbf{k} = (k_1, \dots, k_m)$.
- 9: **end Procedure**

solution to the overlapping MulaNE with the stationary starting distributions.

4.2. Non-overlapping MuLaNE

For non-overlapping MuLaNE, we are able to achieve the exact optimal solution: for the stationary starting distribution, a slight modification of the greedy algorithm Alg. 2 gives the optimal solution, while for an arbitrary starting distribution, we design a dynamic programming algorithm to compute the optimal solution. Due to the space constraint, the details are included in Appendix E.

5. Online Learning for MuLaNE

In the online setting, we continue to study both overlapping and non-overlapping MuLaNE problem. However, the network structure \mathcal{G} , the distribution α and node weights σ are not known a priori. Instead, the only information revealed to the decision maker includes total budget B , the number of layers m , the number of the target nodes $|\mathcal{V}|$ (or an upper bound of it) and the budget constraint c .

5.1. Online Algorithm for Overlapping MuLaNE

For the unknown network structure and starting distributions, we bypass the transition matrices P_i and directly estimate the visiting probabilities $P_{i,u}(b) \in [0, 1]$. This avoids the analysis for how the estimated P_i affects the performance of online algorithms, which could be unbounded since we consider general graph structures and the reward function given by Eq. (7) is highly non-linear in P_i . Moreover, we can save the matrix calculation by directly using $P_{i,u}(b)$,

which is efficient even for large networks.

Specifically, we maintain a set of base arms $\mathcal{A} = \{(i, u, b) | i \in [m], u \in \mathcal{V}, b \in [c_i]\}$, where the total number $|\mathcal{A}| = \sum_{i \in [m]} c_i |\mathcal{V}|$. For each base arm $(i, u, b) \in \mathcal{A}$, we denote $\mu_{i,u,b}$ as the true value of each base arm, i.e., $\mu_{i,u,b} = P_{i,u}(b)$. For the unknown node weights, we maintain the optimistic weight $\bar{\sigma}_u = 1$ if $u \in \mathcal{V}$ has not been visited. After u is first visited and its true value σ_u is revealed, we replace $\bar{\sigma}_u$ with σ_u . With a little abuse of the notation, we use $\boldsymbol{\mu}$ to denote the unknown intermediate parameters and $r_{\boldsymbol{\mu}, \boldsymbol{\sigma}}(\mathbf{k})$ to denote the reward $r_{\mathcal{G}, \boldsymbol{\alpha}, \boldsymbol{\sigma}}(\mathbf{k})$.

We present our algorithm in Alg. 3, which is an adaptation of the CUCB algorithm for the general Combinatorial Multi-arm Bandit (CMAB) framework (Chen et al., 2016) to our setting. Notice that we use \mathcal{A} as defined above in the algorithm, and \mathcal{A} is defined using \mathcal{V} , which is the set of node ids and should not be known before the learning process starts. This is not an issue, because at the beginning we can create $|\mathcal{V}|$ (which is known) placeholders for the node ids, and once a new node is visited, we immediately replace one of the placeholders with the new node id. Thus \mathcal{V} appearing in the above definition of \mathcal{A} is just for notational convenience.

In Alg. 3, we maintain an unbiased estimation of the visiting probability $P_{i,u}(b)$, denoted as $\hat{\mu}_{i,u,b}$. Let $T_{i,u,b}$ record the number of times arm (i, u, b) is played so far and $\bar{\sigma}_v$ denote the optimistic importance weight. In each round $t \geq 1$, we compute the confidence radius $\rho_{i,u,b}$ in line 5, which controls the level of exploration. The confidence radius is larger when the arm (i, u, b) is not explored often (i.e. $T_{i,u,b}$ is small), and thus motivates more exploration. Due to the randomness of the exploration process, the upper confidence bound (UCB) value $\tilde{\mu}_{i,u,b}$ could be decreasing w.r.t b , but our offline oracle BEG can only accept non-decreasing UCB values (otherwise the $1 - e^{-\eta}$ approximation is not guaranteed). Therefore, in line 8, we increase the UCB value $\tilde{\mu}_{i,u,b}$ and set it to be $\max_{j \in [b]} \tilde{\mu}_{i,u,j}$. This is the adaption of the CUCB algorithm to fix our case, and thus we name our algorithm as CUCB-MAX. After we apply the $1 - e^{-\eta}$ approximate solution \mathbf{k} given by the BEG oracle (i.e., Alg. 1), we get m trajectories as feedbacks. In line 11, we can update the unknown weights of visited nodes. In line 13, for base arms (i, u, b) with $b \leq k_i$, we update corresponding statistics by the Bernoulli random variable $Y_{i,u,b} \in \{0, 1\}$ indicating whether node u is visited by W_i in first b steps.

Regret Analysis. We define the reward gap $\Delta_{\mathbf{k}} = \max(0, \xi r_{\boldsymbol{\mu}, \boldsymbol{\sigma}}(\mathbf{k}^*) - r_{\boldsymbol{\mu}, \boldsymbol{\sigma}}(\mathbf{k}))$ for all feasible action \mathbf{k} satisfying $\sum_{i=1}^m k_i = B$, $0 \leq k_i \leq c_i$, where \mathbf{k}^* is the optimal solution for parameters $\boldsymbol{\mu}$, $\boldsymbol{\sigma}$ and $\xi = 1 - e^{-\eta}$ is the approximation ratio of the $(\xi, 1)$ -approximate offline oracle. For each base arm (i, u, b) , we define $\Delta_{\min}^{i,u,b} = \min_{\Delta_{\mathbf{k}} > 0, k_i = b} \Delta_{\mathbf{k}}$ and $\Delta_{\max}^{i,u,b} = \max_{\Delta_{\mathbf{k}} > 0, k_i = b} \Delta_{\mathbf{k}}$. As a convention, if there

Algorithm 3 CUCB-MAX Algorithm for the MuLaNE

Input: Budget B , number of layers m , number of nodes $|\mathcal{V}|$, constraints \mathbf{c} , offline oracle BEG.

- 1: For each arm $(i, u, b) \in \mathcal{A}$, $T_{i,u,b} \leftarrow 0$, $\hat{\mu}_{i,u,b} \leftarrow 0$.
- 2: For each node $v \in \mathcal{V}$, $\bar{\sigma}_v \leftarrow 1$.
- 3: **for** $t = 1, 2, 3, \dots, T$ **do**
- 4: **for** $(i, u, b) \in \mathcal{A}$ **do**
- 5: $\rho_{i,u,b} \leftarrow \sqrt{3 \ln t / (2T_{i,u,b})}$.
- 6: $\tilde{\mu}_{i,u,b} \leftarrow \min\{\hat{\mu}_{i,u,b} + \rho_{i,u,b}, 1\}$.
- 7: **end for**
- 8: For $(i, u, b) \in \mathcal{A}$, $\bar{\mu}_{i,u,b} \leftarrow \max_{j \in [b]} \tilde{\mu}_{i,u,j}$.
- 9: $\mathbf{k} \leftarrow \text{BEG}((\bar{\mu}_{i,u,b})_{(i,u,b) \in \mathcal{A}}, (\bar{\sigma}_v)_{v \in \mathcal{V}}, B, \mathbf{c})$.
- 10: Apply budget allocation \mathbf{k} , which gives trajectories $\mathbf{X} := (X_{i,1}, \dots, X_{i,k_i})_{i \in [m]}$ as feedbacks.
- 11: For any visited node $v \in \bigcup_{i \in [m]} \{X_{i,1}, \dots, X_{i,k_i}\}$, receive its node weight σ_v and set $\bar{\sigma}_v \leftarrow \sigma_v$.
- 12: For any $(i, u, b) \in \tau := \{(i, u, b) \in \mathcal{A} | b \leq k_i\}$, $Y_{i,u,b} \leftarrow 1$ if $u \in \{X_{i,1}, \dots, X_{i,b}\}$ and 0 otherwise.
- 13: For $(i, u, b) \in \tau$, update $T_{i,u,b}$ and $\hat{\mu}_{i,u,b}$:
 $T_{i,u,b} \leftarrow T_{i,u,b} + 1$, $\hat{\mu}_{i,u,b} \leftarrow \hat{\mu}_{i,u,b} + (Y_{i,u,b} - \hat{\mu}_{i,u,b}) / T_{i,u,b}$.
- 14: **end for**

is no action \mathbf{k} with $k_i = b$ such that $\Delta_{\mathbf{k}} > 0$, we define $\Delta_{\min}^{i,u,b} = \infty$ and $\Delta_{\max}^{i,u,b} = 0$. Let $\Delta_{\min} = \min_{(i,u,b) \in \mathcal{A}} \Delta_{\min}^{i,u,b}$ and $\Delta_{\max} = \max_{(i,u,b) \in \mathcal{A}} \Delta_{\max}^{i,u,b}$. The following theorem summarizes the regret bound for Alg. 3.

Theorem 3. *Algorithm 3 has the following distribution-dependent $(1 - e^{-\eta}, 1)$ approximation regret,*

$$\text{Reg}_{\boldsymbol{\mu}, \boldsymbol{\sigma}}(T) \leq \sum_{(i,u,b) \in \mathcal{A}} \frac{108m|\mathcal{V}| \ln T}{\Delta_{\min}^{i,u,b}} + 2|\mathcal{A}| + \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max}.$$

Remark 1. Looking at the above distribution dependent bound, we have the $O(\log T)$ approximation regret, which is asymptotically tight. Coefficient $m|\mathcal{V}|$ in the leading term corresponds to the number of edges in the complete bipartite coverage graph. Notice that we cannot use the true edge number $\sum_{i \in [m]} |\mathcal{V}_i|$, because the learning algorithm does not know which nodes are contained in each layer, and has to explore all visiting possibilities given by the default complete bipartite graph. The set of base arms \mathcal{A} has some redundancy due to the correlation between these base arms, and thus it is unclear if the summation over all base arms in the regret bound is tight. For the non-overlapping case, we further reduce the number of base arms to achieve better regret bounds, but for the overlapping case, how to further reduce base arms to achieve a tighter regret bound is a challenging open question left for the future work.

Remark 2. The $(1 - e^{-\eta}, 1)$ approximate regret is determined by the offline oracle BEG that we plug in Line 9 and can be replaced by $(1 - 1/e, 1)$ regret using BEGE or even by the exact regret if the oracle can obtain the optimal budget allocation. The usage of BEG is a trade-off we

make between computational efficiency and learning efficiency and empirically, it performs well as we shall see in Section 6.

Remark 3. The full proof of the above theorem is included in the Appendix F, where we rely on the following properties of $r_{\mu,\sigma}(\mathbf{k})$ to bound the regret.

Property 1. (Monotonicity). *The reward $r_{\mu,\sigma}(\mathbf{k})$ is monotonically increasing, i.e., for any budget allocation \mathbf{k} , any two vectors $\mu = (\mu_{i,u,b})_{(i,u,b) \in \mathcal{A}}$, $\mu' = (\mu'_{i,u,b})_{(i,u,b) \in \mathcal{A}}$ and any node weights σ, σ' , we have $r_{\mu,\sigma}(\mathbf{k}) \leq r_{\mu',\sigma'}(\mathbf{k})$, if $\mu_{i,u,b} \leq \mu'_{i,u,b}$ and $\sigma_v \leq \sigma'_v, \forall (i, u, b) \in \mathcal{A}, v \in \mathcal{V}$.*

Property 2. (1-Norm Bounded Smoothness). *The reward function $r_{\mu,\sigma}(\mathbf{k})$ satisfies the 1-norm bounded smoothness condition, i.e., for any budget allocation \mathbf{k} , any two vectors $\mu = (\mu_{i,u,b})_{(i,u,b) \in \mathcal{A}}$, $\mu' = (\mu'_{i,u,b})_{(i,u,b) \in \mathcal{A}}$ and any node weights σ, σ' , we have $|r_{\mu,\sigma}(\mathbf{k}) - r_{\mu',\sigma'}(\mathbf{k})| \leq \sum_{i \in [m], u \in \mathcal{V}, b = k_i} (\sigma_u |\mu_{i,u,b} - \mu'_{i,u,b}| + |\sigma_u - \sigma'_u| \mu'_{i,u,b})$.*

We emphasize that our algorithm and analysis differ from the original CUCB algorithm (Chen et al., 2016) as follows. First, we have the additional regret caused by the over-estimated weights for unvisited nodes, i.e., $|\sigma_u - \sigma'_u| \mu'_{i,u,b}$ term in property 2. We carefully bound this term based on the observation that $\mu'_{i,u,b}$ is small and decreasing quickly before u is first visited. Next, we have to take the max (line 8) to guarantee the UCB value is monotone w.r.t b since our BEG oracle can only output $(1 - e^{-\eta}, 1)$ -approximation with monotone inputs. Due to the above operation, (i, u, b) 's UCB value depends on the feedback from all arms (i, u, j) for $j \leq b$ (set τ in line 12). So we should update all these arms (line 13) to guarantee that the estimates to all these arms are accurate enough. Finally, directly following the standard CMAB result would have a larger regret, because arms in τ are defined as triggered arms, but only arms in $\tau' = \{(i, u, b) \in \mathcal{A} | k_i = b\}$ affect the rewards. So we conceptually view arms in τ' as triggered arms and use a tighter 1-Norm Bounded Smoothness condition as given above to derive a tighter regret bound. This improves the coefficient of the leading $\ln T$ term in the distribution dependent regret by a factor of $|\tau|/|\tau'| = O(B/m)$, and the $1/\Delta_{\min}^{i,u,b}$ term is smaller since the original definition would have $\Delta_{\min}^{i,u,b} = \min_{\Delta_{\mathbf{k}} > 0, b \leq k_i} \Delta_{\mathbf{k}}$.

5.2. Online Algorithm for Non-overlapping Case

For the non-overlapping case, we set *layer-wise marginal gains* as our base arms. Concretely, we maintain a set of base arms $\mathcal{A} = \{(i, b) | i \in [m], b \in [c_i]\}$. For each base arm $(i, b) \in \mathcal{A}$, let $\mu_{i,b} = \sum_{u \in \mathcal{V}} \sigma_u (P_{i,u}(b) - P_{i,u}(b-1))$ be the true marginal gain of assigning budget b in layer i . We apply the standard CUCB algorithm to this setting and call the resulting algorithm CUCB-MG (Algorithm 9 in Appendix G). Note that in the non-overlapping setting we can

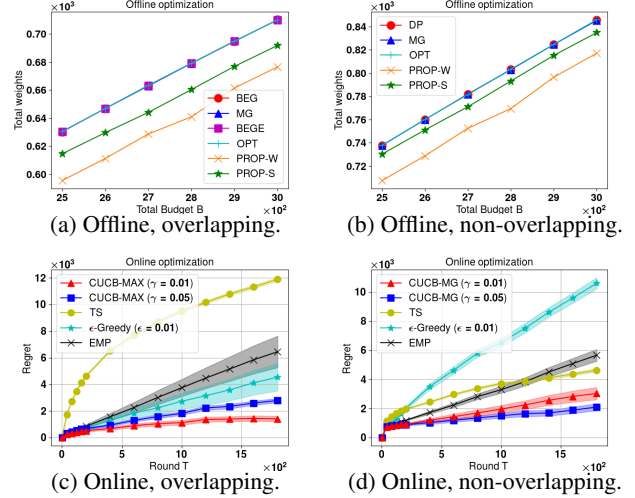


Figure 1. Above: total weights of unique nodes visited for offline algorithms. Below: regret for online algorithms when $B = 3000$.

Table 2. Statistics for FF-TW-YT network

Layer	FriendFeed	Twitter	YouTube
# of vertices	5,540	5,702	663
# of edges	31,921	42,327	614

solve the offline problem exactly, so we can use the exact offline oracle to solve the online problem and achieve an exact regret bound. This is the major advantage over the overlapping setting where we can only achieve an approximate bound. Define $\Delta_{\mathbf{k}} = r_{\mu,\sigma}(\mathbf{k}^*) - r_{\mu,\sigma}(\mathbf{k})$ for all feasible action \mathbf{k} , and $\Delta_{\min}^{i,b} = \min_{\Delta_{\mathbf{k}} > 0, k_i > b} \Delta_{\mathbf{k}}$, CUCB-MG has a $O(\sum_{(i,b) \in \mathcal{A}} 48B \ln T / \Delta_{\min}^{i,b})$ regret bound.

6. Experiments

Dataset and settings. We conduct experiments on a real-world multi-layered network FF-TW-YT, which contains $m = 3$ layers representing users' social connections in FriendFeed (FF), Twitter (TW) and YouTube (YT) (Dickison et al., 2016). In total, FF-TW-YT has 6,407 distinct vertices representing users and 74,836 directed edges representing connections (“who follows whom”) among users. The statistics for each layer is summarized in Table 2. We transform the FF-TW-YT network $\mathcal{G}(\mathcal{V}, \mathcal{E})$ into a symmetric directed network (by adding a new edge (v, u) if $(u, v) \in \mathcal{E}$ but $(v, u) \notin \mathcal{E}$) because a user can be visited via her followers or followees. Each edge weight is set to be 1 and the node weights are set to be $\sigma_u \in \{0, 0.5, 1\}$ uniformly at random. Each random walker always starts from the smallest node-id in each layer and we set constraints c_i equal to the total budget B . Note that in order to test the non-overlapping case, we use the same network but relabel node-ids so that

Table 3. Running time (seconds) for offline and online algorithms.

	B=2.6k	B=2.8k	B=3.0k		B=2.6k	B=2.8k	B=3.0k	Overlapping?	✓	×
BEG	0.274	0.316	0.363	DP	0.038	0.044	0.050	BEG	1.22	NA
BEGE	34.37	45.51	59.20	MG	0.008	0.009	0.010	BEGE	60.03	NA
OPT	91.16	98.42	105.63	OPT	70.10	75.78	81.11	DP	NA	0.86
								OPT	107.33	82.01

(a) Running time of offline algorithms for the overlapping case with different budgets B.

(b) Running time of offline algorithms for the non-overlapping case with different budgets B.

(c) Per-round running time for CUCB-MAX (or CUCB-MG) with different oracles when B=3.0k.

they do not overlap between different layers. To handle the randomness, we repeat 10,000 times and present the averaged total weights of unique nodes visited for offline optimization. We calculate the regret by comparing with the *optimal* solution, which is stronger than comparing with $(1 - e^{-\eta}, 1)$ -approximate solution as defined in Eq. (4). We average over 200 independent experiments to provide the mean regret with 95% confidence interval. To evaluate the computational efficiency, we also present the running time for both offline and online algorithms in Table 3.

Algorithms in comparison. For the offline setting, we present the results for Alg. 1 (denoted as BEG), Alg. 2 (denoted as MG), BEG with partial enumeration (denoted as BEGE) and Alg. 7 (denoted as DP). We provide two baselines PROP-S and PROP-W, which allocates the budget proportional to the layer size and proportional to the total weights if we allocate $B/3$ budgets to that layer, respectively. The optimal solution (denoted as OPT) is also provided by enumerating all possible budget allocations. For online settings, we consider CUCB-MAX (Alg. 3) and CUCB-MG (Alg. 9) algorithms. We shrink the confidence interval by γ , i.e., $\rho_{i,u,b} \leftarrow \gamma \rho_{i,u,b}$, to speed up the learning, though our theoretical regret bound requires $\gamma = 1$. For baselines, we consider the EMP algorithm which always allocates according to the empirical mean, and the ϵ -Greedy algorithm which allocates budgets according to empirical mean with probability $1 - \epsilon$ and allocates all B budgets to the i -th layer with probability ϵ/m . We also compare with the Thompson sampling (TS) method (Wang & Chen, 2018), which uses Beta distribution $\text{Beta}(\alpha, \beta)$ (where $\alpha = \beta = 1$ initially) as prior distribution for each base arm.

Experimental results. We show the results for MuLaNE problems in Figure 1. For the offline overlapping case, both BEG and MG outperform two baselines PROP-W and PROP-S in receiving total weights. Although not guaranteed by the theory, BEG are empirically close to BEGE and the optimal solution (OPT). As for the computational efficiency, in Table 3a, the BEG is at least two orders of magnitude (e.g., 163 times when B=3.0k) faster than BEGE and OPT. Combining that the reward of BEG is empirically close to BEGE and the optimal solution, this shows that BEG is empirically better than BEGE and OPT. For the offline non-

overlapping case, the results are similar, but the difference is that we have the theoretical guarantee for the optimality of DP. For the online setting, all CUCB-MAX/CUCB-MG curves outperform the baselines. This demonstrates empirically that CUCB-MAX algorithm can effectively learn the unknown parameters while optimizing the objective. For the computational efficiency of online learning algorithms, since the running time for algorithms with the same oracle is similar, we present the running time for CUCB-MAX with different oracles in Table 3c. CUCB-MAX with BEG is 50 times faster than BEGE, which is consistent with our theoretical analysis. The results for different budgets B are consistent with $B = 3000$, which are included in the Appendix H.2. Results and analysis for stationary starting distributions are also in the Appendix H.1.

7. Conclusions and Future Work

This paper formulates the multi-layered network exploration via random walks (MuLaNE) as a budget allocation problem, requiring that the total weights of distinct nodes visited on the multi-layered network is maximized. For the offline setting, we propose four algorithms for MuLaNE according to the specification of multi-layered network (overlapping or non-overlapping) and starting distributions (arbitrary or stationary), each of which has a provable guarantee on approximation factors and running time. We further study the online setting where network structure and the node weights are not known a priori. We propose the CUCB-MAX algorithm for overlapping MuLaNE and the CUCB-MG algorithm for the non-overlapping case, both of which are bounded by a $O(\log T)$ (approximate) regret. Finally, we conduct experiments on a social network dataset to show the empirical performance of our algorithms.

There are many compelling directions for the future study. For example, it would be interesting to extend our problem where the decision maker can jointly optimize the starting distribution and the budget allocation. One could also study the adaptive MuLaNE by using the feedback from the exploration results of the previous steps to determine the exploration strategy for future steps.

8. Acknowledgement

The work of John C.S. Lui was supported in part by the GRF 14200420.

References

- Alon, N., Gamzu, I., and Tennenholtz, M. Optimizing budget allocation among channels and influencers. In *Proceedings of the 21st international conference on World Wide Web*, pp. 381–388. ACM, 2012.
- Bubeck, S. and Cesa-Bianchi, N. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122, 2012.
- Chen, W., Wang, Y., and Yang, S. Efficient influence maximization in social networks. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 199–208, 2009.
- Chen, W., Wang, Y., Yuan, Y., and Wang, Q. Combinatorial multi-armed bandit and its extension to probabilistically triggered arms. *The Journal of Machine Learning Research*, 17(1):1746–1778, 2016.
- Chen, X., Huang, W., Chen, W., and Lui, J. C. Community exploration: From offline optimization to online learning. In *Advances in Neural Information Processing Systems*, pp. 5474–5483, 2018.
- Dickison, M. E., Magnani, M., and Rossi, L. *Multilayer social networks*. Cambridge University Press, 2016.
- Dubhashi, D. P. and Panconesi, A. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009.
- Gai, Y., Krishnamachari, B., and Jain, R. Combinatorial network optimization with unknown variables: Multi-armed bandits with linear rewards and individual observations. *IEEE/ACM Transactions on Networking (TON)*, 20(5):1466–1478, 2012.
- Gleich, D. F. Pagerank beyond the web. *SIAM Review*, 57(3):321–363, 2015.
- Kapralov, M., Post, I., and Vondrák, J. Online submodular welfare maximization: Greedy is optimal. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pp. 1216–1225. SIAM, 2013.
- Kempe, D., Kleinberg, J., and Tardos, É. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 137–146, 2003.
- Khuller, S., Moss, A., and Naor, J. S. The budgeted maximum coverage problem. *Information processing letters*, 70(1):39–45, 1999.
- Kijima, M. *Markov processes for stochastic modeling*, volume 6. CRC Press, 1997.
- Krause, A., Leskovec, J., Guestrin, C., VanBriesen, J., and Faloutsos, C. Efficient sensor placement optimization for securing large water distribution networks. *Journal of Water Resources Planning and Management*, 134(6):516–526, 2008.
- Lerman, K. and Jones, L. Social browsing on flickr. *arXiv preprint cs/0612047*, 2006.
- Li, Y., Wu, Z., Lin, S., Xie, H., Lv, M., Xu, Y., and Lui, J. C. Walking with perception: Efficient random walk sampling via common neighbor awareness. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, pp. 962–973. IEEE, 2019.
- Lv, Q., Cao, P., Cohen, E., Li, K., and Shenker, S. Search and replication in unstructured peer-to-peer networks. In *Proceedings of the 16th international conference on Supercomputing*, pp. 84–95, 2002.
- Pons, P. and Latapy, M. Computing communities in large networks using random walks. In *International symposium on computer and information sciences*, pp. 284–293. Springer, 2005.
- Raginsky, M., Sason, I., et al. Concentration of measure inequalities in information theory, communications, and coding. *Foundations and Trends® in Communications and Information Theory*, 10(1-2):1–246, 2013.
- Robbins, H. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535, 1952.
- Serfozo, R. *Basics of applied stochastic processes*. Springer Science & Business Media, 2009.
- Soma, T. and Yoshida, Y. A generalization of submodular cover via the diminishing return property on the integer lattice. In *Advances in Neural Information Processing Systems*, pp. 847–855, 2015.
- Soma, T., Kakimura, N., Inaba, K., and Kawarabayashi, K.-i. Optimal budget allocation: Theoretical guarantee and efficient algorithm. In *International Conference on Machine Learning*, pp. 351–359, 2014.
- Wang, C., Chen, W., and Wang, Y. Scalable influence maximization for independent cascade model in large-scale social networks. *Data Mining and Knowledge Discovery*, 25(3):545–576, 2012.

Wang, Q. and Chen, W. Improving regret bounds for combinatorial semi-bandits with probabilistically triggered arms and its applications. In *Advances in Neural Information Processing Systems*, pp. 1161–1171, 2017.

Wang, S. and Chen, W. Thompson sampling for combinatorial semi-bandits. In *International Conference on Machine Learning*, pp. 5114–5122, 2018.

Wilder, B., Immorlica, N., Rice, E., and Tambe, M. Maximizing influence in an unknown social network. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

Supplementary Material

The supplementary material is organized as follows.

We first discuss how we handle the multiple random walks in Section A. We provide proofs and examples for properties of the visiting probability in Section B. Proofs of offline optimization for overlapping MuLaNE are provided in Section C. We provide the detailed budget effective greedy algorithm with partial enumeration (BEGE) and its analysis in Section D. Proofs of offline optimization for non-overlapping MuLaNE are provided in Section E. We state the detailed analysis of online learning for overlapping MuLaNE in Section F. We state the detailed analysis of online learning for non-overlapping MuLaNE in Section G. Supplemental experiments are provided in Section H.

A. Handling the Multiple Random Walkers




Figure 2. Two types of multi-layered networks.

We can handle the scenario where each layer is explored by multiple random walkers using the bipartite coverage model. Without loss of generality, suppose we want to add a new random walker W'_1 to the layer L_1 . A new node W'_1 can be easily added to \mathcal{L} representing the new random explorer, and new edges $\{(W'_1, u) | u \in \mathcal{V}_1\}$ are added to \mathcal{E}' . Thus, we can use the same algorithms and analysis to solve the optimization problem on the newly constructed bipartite coverage graph.

B. Proofs and Examples for Properties of the Visiting Probability $P_{i,u}(k_i)$

B.1. Starting From Arbitrary Distributions

Lemma 1. $g_{i,u}(k_i) \geq 0$ for any $i \in [m]$, $u \in \mathcal{V}_i$, $k_i \in \mathbb{Z}_{>0}$.

Proof. For analysis, we use the following equivalent formulation for $k_i \geq 2$, (trivially $P_{i,u}(0) = 0$ and $P_{i,u}(1) = \alpha_{i,u}$),

$$P_{i,u}(k_i) = ((\alpha_i)_{-u}^\top (\Gamma_i)_{uk_i} + \alpha_{i,u}), \quad (12)$$

where $(\Gamma_i)_{uk_i} = (\mathbf{I} + \mathbf{T}_i(\mathbf{u}) + \dots + \mathbf{T}_i(\mathbf{u})^{k_i-2})(\mathbf{p}_i)_u$, $(\mathbf{p}_i)_u = (\mathbf{P}_i[\cdot, u])_{-u}$, $\mathbf{T}_i(\mathbf{u}) = (\mathbf{P}_i)_{-u, -u}$, (we use Eq.

(2.67) in (Kijima, 1997) to derive Γ_{uk_i}). Note that $\mathbf{p}_{-u} \in \mathbb{R}^{n-1}$ is the vector obtained by deleting the u -th element from $\mathbf{p} \in \mathbb{R}^n$, and $\mathbf{P}_{-u, -v} \in \mathbb{R}^{(n-1) \times (n-1)}$ is the matrix obtained by deleting the u -th row and the v -th column from $\mathbf{P} \in \mathbb{R}^{n \times n}$.

We further derive the marginal gain of $P_{i,u}(\cdot)$ at step $k_i \geq 2$ (trivially $g_{i,u}(1) = \alpha_{i,u}$), as,

$$g_{i,u}(k_i) = (\alpha_i)_{-u}^\top \mathbf{T}_i(\mathbf{u})^{k_i-2} (\mathbf{p}_i)_u. \quad (13)$$

Now, we can show $g_{i,u}(k_i)$ is non-negative because any element of $(\alpha_i)_{-u}^\top$, $\mathbf{T}_i(\mathbf{u})$ and $(\mathbf{p}_i)_u$ are non-negative, which means $P_{i,u}(k_i)$ is non-decreasing with respect to step k_i . ■

An example showing $g_{i,u}(k_i)$ is not monotone: Consider a path P_3 with three nodes as the i -th layer $\mathcal{G}_i(\mathcal{V}_i, \mathcal{E}_i)$, where $\mathcal{V}_i = \{u, v, w\}$ and $\mathcal{E}_i = \{(u, v), (v, u), (v, w), (w, v)\}$. If the W_i always starts from the left-most node u and chooses the right-most node w as our target node. Then, $g_{i,w}(1) = g_{i,w}(2) = 0$ but $g_{i,w}(3) > 0$ since at least three steps are needed to visit the node w , which shows that $g_{i,w}(k_i)$ is not always non-increasing.

B.2. Starting From the Stationary Distribution

Lemma 2. $g_{i,u}(k_i + 1) - g_{i,u}(k_i) \leq 0$ for any $i \in [m]$, $u \in \mathcal{V}_i$, $k_i \in \mathbb{Z}_{>0}$, if $\alpha_i = \pi_i$, where $\pi_i^\top \mathbf{P}_i = \pi_i^\top$.

Proof. Consider any layer L_i with transition probability matrix \mathbf{P}_i , if we start from the stationary distribution π_i , the probability that node $u \in \mathcal{V}_i$ is ever visited in the first k_i steps is $P_{i,u}(k_i) = (\pi_i)_{-u}^\top (\Gamma_i)_{uk_i} + \pi_{i,u}$, where $(\Gamma_i)_{uk_i} = (\mathbf{I} + \mathbf{T}_i(\mathbf{u}) + \dots + \mathbf{T}_i(\mathbf{u})^{k_i-2})(\mathbf{p}_i)_u$, $\mathbf{T}_i(\mathbf{u}) = (\mathbf{P}_i)_{-u, -u}$, $(\mathbf{p}_i)_u = (\mathbf{P}_i[\cdot, u])_{-u}$. Then the marginal gain for node u is $g_{i,u}(k_i) = P_{i,u}(k_i) - P_{i,u}(k_i - 1) = (\pi_i)_{-u}^\top \mathbf{T}_i(\mathbf{u})^{k_i-2} (\mathbf{p}_i)_u$ for $k_i \geq 2$ and $g_{i,u}(k_i) = \pi_{i,u}$ when $k_i = 1$.

Define the margin of the marginal gain as $\Delta_{i,u}(k_i) = g_{i,u}(k_i + 1) - g_{i,u}(k_i)$. When $k = 1$, $\Delta_{i,u}(1) = g_{i,u}(2) - g_{i,u}(1) = (\pi_i)_{-u}^\top \cdot (\mathbf{p}_i)_u - \pi_{i,u} = \pi_{i,u} - \pi_{i,u} \mathbf{P}[u, u] - \pi_{i,u} = -\pi_{i,u} \mathbf{P}[u, u] \leq 0$. When $k_i \geq 2$, $\Delta_{i,u}(k_i) = g_{i,u}(k_i + 1) - g_{i,u}(k_i) = (\pi_i)_{-u}^\top (\mathbf{T}_i(\mathbf{u})^{k_i-1} - \mathbf{T}_i(\mathbf{u})^{k_i-2})(\mathbf{p}_i)_u$. Because $\pi_i^\top \mathbf{P}_i = \pi_i^\top$, we have $(\pi_i)_{-u}^\top - (\pi_i)_{-u}^\top \mathbf{P}_i(u) = \pi_{i,u} (\mathbf{q}_i)_{-u}^\top$, where $(\mathbf{q}_i)_{-u}^\top = (\mathbf{P}_i[u, \cdot])_{-u}$. Thus, $\Delta_{i,u}(k_i) = -(\pi_{i,u}) (\mathbf{q}_i)_{-u}^\top \mathbf{T}_i(\mathbf{u})^{k_i-2} (\mathbf{p}_i)_u \leq 0$ because any element in π_i , $(\mathbf{q}_i)_{-u}^\top$, $(\mathbf{p}_i)_u$ and $\mathbf{P}_i(u)$ is non-negative. ■

More interestingly, the stationary distribution π_i is the *only* starting distribution for \mathbf{P}_i such that any $u \in \mathcal{V}_i$, $k_i \in \mathbb{Z}_{>0}$, $g_{i,u}(k_i + 1) - g_{i,u}(k_i) \leq 0$ when \mathbf{P}_i is ergodic and there are no self loops in \mathcal{G}_i .

Proof. With a little abuse of the notation, we use P to denote the transition probability matrix $P_i \in \mathbb{R}^{n \times n}$ and let $P_{i,j}$ be the element in the i -th row and the j -th column. Since P is ergodic, according to Theorem 54 in (Serfozo, 2009), there exists a unique and positive stationary distribution $\pi = (\pi_1, \dots, \pi_n)$, i.e., $\pi P = \pi$ and $\pi > \mathbf{0}$. Any starting distribution α can be represented by $\pi + \epsilon$, where $\epsilon = (\epsilon_1, \dots, \epsilon_n)$ is a perturbation vector, and $-\pi_j \leq \epsilon_j \leq 1 - \pi_j, j \in [n]$. We have the following equation for the margin of marginal gains for node u in the first two steps,

$$\begin{aligned} \Delta_u &= g_{i,u}(1) - g_{i,u}(2) = (\pi_u + \epsilon_u) - \sum_{j \neq u} (\pi_j + \epsilon_j) P_{j,u} \\ &= \pi_u P_{uu} + \epsilon_u - \sum_{j \neq u} \epsilon_j P_{j,u}. \end{aligned} \quad (14)$$

Since \mathcal{G} has no self loops, i.e., $P_{u,u} = 0$, we have $\Delta_u = \epsilon_u - \sum_{j \neq u} \epsilon_j P_{j,u}$, and we can verify that $\sum_{u \in [n]} \Delta_u = \sum_{u \in [n]} (1 - \sum_{j \neq u} P_{u,j}) \epsilon_u = 0$. Therefore, we have to guarantee $\Delta_u = 0$ for all u , otherwise there will exist a node u such that $\Delta_u < 0$. To ensure $\Delta_u = 0$, we need to ensure $\epsilon P = \epsilon$ by rephrasing the equations $\Delta_u = 0$ for all u . Again, since there exists a unique and positive stationary distribution for P and $\sum_{u \in [n]} \epsilon_u = 0$, we can derive $\epsilon = \beta \pi$, where β has to be 0. Therefore, combined with Lemma 2, the stationary distribution is the only starting distribution such that for any P_i , any $u \in \mathcal{V}_i, k_i \in \mathbb{Z}_{>0}$, $g_{i,u}(k_i + 1) - g_{i,u}(k_i) \leq 0$. ■

C. Proofs of Offline Optimization for Overlapping MuLaNE

C.1. Starting from the arbitrary distribution

Lemma 3. For any network \mathcal{G} , distribution α and weights $\sigma, r_{\mathcal{G}, \alpha, \sigma}(\cdot) : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R}$ is monotone and submodular.

Proof. By definition, we need to show $r_{\mathcal{G}, \alpha, \sigma}(\mathbf{x} \wedge \mathbf{y}) + r_{\mathcal{G}, \alpha, \sigma}(\mathbf{x} \vee \mathbf{y}) \leq r_{\mathcal{G}, \alpha, \sigma}(\mathbf{x}) + r_{\mathcal{G}, \alpha, \sigma}(\mathbf{y})$ for any $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_{\geq 0}^m$, and $r_{\mathcal{G}, \alpha, \sigma}(\mathbf{x}) \leq r_{\mathcal{G}, \alpha, \sigma}(\mathbf{y})$ if $\mathbf{x} \leq \mathbf{y}$.

(Monotonicity.) By Eq. (7), $r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k}) = \sum_{v \in \mathcal{V}} \sigma_v (1 - \prod_{i \in [m]} (1 - P_{i,v}(k_i)))$. Since $P_{j,v}(x_j) \leq P_{j,v}(x_j + 1)$, for any $j \in [m], v \in \mathcal{V}, x_j \in \mathbb{Z}_{\geq 0}$, we have

$$\begin{aligned} &r_{\mathcal{G}, \alpha, \sigma}(\mathbf{x} + \chi_j) - r_{\mathcal{G}, \alpha, \sigma}(\mathbf{x}) \\ &= \sum_{v \in \mathcal{V}} \sigma_v \left(\left(\prod_{i \neq j} (1 - P_{i,v}(x_i)) \right) (P_{j,v}(x_j + 1) - P_{j,v}(x_j)) \right) \\ &\geq 0, \end{aligned}$$

for any $\mathbf{x} \in \mathbb{Z}_{\geq 0}^m, j \in [m]$. Then we can use the above inequality repeatedly to show that $r_{\mathcal{G}, \alpha, \sigma}(\mathbf{x}) \leq r_{\mathcal{G}, \alpha, \sigma}(\mathbf{y})$ when $\mathbf{x} \leq \mathbf{y}$.

(Submodularity.) For submodular property, it is sufficient to prove $(1 - \prod_{i \in [m]} (1 - P_{i,v}(k_i)))$ is submodular for any $v \in \mathcal{V}$, because a positive weighted sum of submodular function is still submodular. We will rely on the following lemma to prove the submodularity of $r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k})$.

Lemma 5. Function $f : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R}$ is submodular if and only if

$$\begin{aligned} f(\mathbf{x} + \chi_i) - f(\mathbf{x}) &\geq f(\mathbf{x} + \chi_j + \chi_i) - f(\mathbf{x} + \chi_j), \\ \text{for any } \mathbf{x} \in \mathbb{Z}_{\geq 0}^m \text{ and } i \neq j. \end{aligned}$$

Proof of Lemma 5. (If part.) We first prove, if inequality (14) holds, the following inequality holds,

$$f(\mathbf{x} + \chi_i) - f(\mathbf{x}) \geq f(\mathbf{y} + \chi_i) - f(\mathbf{y}) \quad (15)$$

for any $\mathbf{x} \leq \mathbf{y}$ and $i \in [m]$ such that $x_i = y_i$.

Let $I_0 = \{i \in [m] : x_i = y_i\}, I_1 = \{i \in [m] : x_i < y_i\}$. For any $\mathbf{x} \leq \mathbf{y}$, we denote the elements in I_1 by i_1, \dots, i_s and write $\mathbf{y} = \mathbf{x} + \sum_{j=1}^s \alpha_j \chi_{i_j}$, where $\alpha_j = y_j - x_j$. For any $i \in I_0$, we have

$$\begin{aligned} &f(\mathbf{x} + \chi_i) - f(\mathbf{x}) \\ &\geq f(\mathbf{x} + \chi_{i_1} + \chi_i) - f(\mathbf{x} + \chi_{i_1}) \\ &\geq f(\mathbf{x} + 2\chi_{i_1} + \chi_i) - f(\mathbf{x} + 2\chi_{i_1}) \\ &\geq \dots \\ &\geq f(\mathbf{x} + \alpha_{i_1} \chi_{i_1} + \chi_i) - f(\mathbf{x} + \alpha_{i_1} \chi_{i_1}) \\ &\geq f(\mathbf{x} + \alpha_{i_1} \chi_{i_1} + \chi_{i_2} + \chi_i) - f(\mathbf{x} + \alpha_{i_1} \chi_{i_1} + \chi_{i_2}) \\ &\geq f(\mathbf{x} + \alpha_{i_1} \chi_{i_1} + \alpha_{i_2} \chi_{i_2} + \chi_i) - f(\mathbf{x} + \alpha_{i_1} \chi_{i_1} + \alpha_{i_2} \chi_{i_2}) \\ &\geq \dots \\ &\geq f(\mathbf{x} + \sum_{j=1}^s \alpha_j \chi_{i_j} + \chi_i) - f(\mathbf{x} + \sum_{j=1}^s \alpha_j \chi_{i_j}) \\ &= f(\mathbf{y} + \chi_i) - f(\mathbf{y}). \end{aligned} \quad (16)$$

Then for any $i \in I_0$ and $a \in \mathbb{Z}_{\geq 0}$, we have

$$\begin{aligned} &f(\mathbf{x} + a\chi_i) - f(\mathbf{x}) \\ &= \sum_{j=1}^a (f(\mathbf{x} + j\chi_i) - f(\mathbf{x} + (j-1)\chi_i)) \\ &\geq \sum_{j=1}^a (f(\mathbf{y} + j\chi_i) - f(\mathbf{y} + (j-1)\chi_i)) \\ &\geq f(\mathbf{y} + a\chi_i) - f(\mathbf{y}). \end{aligned} \quad (17)$$

The first inequality holds because of Inequality (16), the fact $\mathbf{x} + (j-1)\chi_i \leq \mathbf{y} + (j-1)\chi_i$ and $x_i + j - 1 = y_i + j - 1$.

Then for any $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_{\geq 0}^m$, let $I_2 = \{i \in [m] : x_i > y_i\} =$

$\{i_1, \dots, i_s\}$. We have

$$\begin{aligned}
 & f(\mathbf{x}) - f(\mathbf{x} \wedge \mathbf{y}) \\
 &= \sum_{l=1}^s \left(f\left(\mathbf{x} \wedge \mathbf{y} + \sum_{j=1}^l (x_{i_j} - y_{i_j}) \chi_{i_j}\right) \right. \\
 &\quad \left. - f\left(\mathbf{x} \wedge \mathbf{y} + \sum_{j=1}^{l-1} (x_{i_j} - y_{i_j}) \chi_{i_j}\right) \right) \\
 &\geq \sum_{l=1}^s \left(f\left(\mathbf{y} + \sum_{j=1}^l (x_{i_j} - y_{i_j}) \chi_{i_j}\right) \right. \\
 &\quad \left. - f\left(\mathbf{y} + \sum_{j=1}^{l-1} (x_{i_j} - y_{i_j}) \chi_{i_j}\right) \right) \\
 &= f(\mathbf{x} \vee \mathbf{y}) - f(\mathbf{y}).
 \end{aligned}$$

The inequality is derived from Inequality (17) because $\mathbf{y} + \sum_{j=1}^{l-1} (x_{i_j} - y_{i_j}) \chi_{i_j} \geq \mathbf{x} \wedge \mathbf{y} + \sum_{j=1}^{l-1} (x_{i_j} - y_{i_j}) \chi_{i_j}$ for any $0 \leq l \leq s$ and $(\mathbf{x} \wedge \mathbf{y})_{i_l} = \mathbf{y}_{i_l}$, which concludes the if part.

(Only if part.) Assume f is submodular, let $\mathbf{a} = \mathbf{x} + \chi_i$, $\mathbf{b} = \mathbf{x} + \chi_j$, $i \neq j$, we have $\mathbf{a} \vee \mathbf{b} = \mathbf{x} + \chi_i + \chi_j$, $\mathbf{a} \wedge \mathbf{b} = \mathbf{x}$. $f(\mathbf{x} + \chi_i) - f(\mathbf{x}) = f(\mathbf{a}) - f(\mathbf{a} \wedge \mathbf{b}) \geq f(\mathbf{a} \vee \mathbf{b}) - f(\mathbf{b}) = f(\mathbf{x} + \chi_j + \chi_i) - f(\mathbf{x} + \chi_j)$.

■

Then, by Lemma 5 and the explicit formula of the reward function given by Eq. (7), we can prove $g(\mathbf{x}, v) - g(\mathbf{x} + \chi_l, v) \geq g(\mathbf{x} + \chi_j, v) - g(\mathbf{x} + \chi_l + \chi_j, v)$ for any $\mathbf{x} \in \mathbb{Z}_{\geq 0}^m$, $v \in \mathcal{V}$ and $l \neq j \in [m]$, where $g(\mathbf{x}, v) = \prod_{i \in [m]} (1 - P_{i,v}(x_i))$. This holds due to the fact that the left hand side equals to $\left(\prod_{i \in [m] \setminus \{j,l\}} (1 - P_{i,v}(x_i))\right) (1 - P_{j,v}(k_j)) (P_{l,v}(k_l + 1) - P_{l,v}(k_l))$ and the right hand side equals to $\left(\prod_{i \in [m] \setminus \{j,l\}} (1 - P_{i,v}(x_i))\right) (1 - P_{j,v}(k_j + 1)) (P_{l,v}(k_l + 1) - P_{l,v}(k_l))$, and the left hand side is larger or equal to the right hand side because $(1 - P_{j,v}(k_j)) \geq (1 - P_{j,v}(k_j + 1))$. By summation over all nodes $v \in \mathcal{V}$ with node weights $\sigma_v \in [0, 1]$, we can prove the reward function is submodular. ■

Theorem 1. *Algorithm 1 obtains a $(1 - e^{-\eta}) \approx 0.357$ -approximate solution, where η is the solution of equation $e^\eta = 2 - \eta$, to the overlapping MuLaNE problem.*

Proof. For theoretical analysis, we first give a modified version of Alg. 1 in Alg. 4. Both algorithms provide the same solution \mathbf{k} given the same problem instance $(\mathcal{G}, \alpha, \sigma, B, c)$. To see this fact, Alg. 4 considers invalid tentative allocations (i^*, b^*) (adding it will exceed the total budget constraint B) and remove them in line 13 of Alg. 4, while Alg. 1 only

Algorithm 4 Equivalent Budget Effective Greedy Algorithm (BEG) for the Overlapping MuLaNE.

Input: Network \mathcal{G} , starting distributions α , node weights σ , budget B , constraints c .

Output: Budget allocation \mathbf{k} .

- 1: Compute visiting probabilities $(P_{i,u}(b))_{i \in [m], u \in \mathcal{V}, b \in [c_i]}$ according to Eq. (9).
- 2: $\mathbf{k} \leftarrow \text{BEG}((P_{i,u}(b))_{i \in [m], u \in \mathcal{V}, b \in [c_i]}, \sigma, B, c)$.
- 3: **Procedure** $\text{BEG}((P_{i,u}(b))_{i \in [m], u \in \mathcal{V}, b \in [c_i]}, \sigma, B, c)$
- 4: Let $\mathbf{k} := (k_1, \dots, k_m) \leftarrow \mathbf{0}$, $K \leftarrow B$.
- 5: Let $\mathcal{Q} \leftarrow \{(i, b_i) \mid i \in [m], 1 \leq b_i \leq c_i\}$.
- 6: **while** $K > 0$ and $\mathcal{Q} \neq \emptyset$ **do**
- 7: $(i^*, b^*) \leftarrow \arg \max_{(i,b) \in \mathcal{Q}} \delta(i, b, \mathbf{k})/b \triangleright \text{Eq. (11)}$
- 8: **if** $b^* \leq K$ **then**
- 9: $k_{i^*} \leftarrow k_{i^*} + b^*$, $K \leftarrow K - b^*$.
- 10: Modify pairs $(i^*, b) \in \mathcal{Q}$ to $(i^*, b - b^*)$.
- 11: Remove pairs $(i^*, b) \in \mathcal{Q}$ such that $b \leq 0$.
- 12: **else**
- 13: Remove (i^*, b^*) from \mathcal{Q} .
- 14: **end if**
- 15: **end while**
- 16: **for** $i \in [m]$ **do**
- 17: **if** $r_{\mathcal{G}, \alpha, \sigma}(c_i \chi_i) > r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k})$, **then** $\mathbf{k} \leftarrow c_i \chi_i$.
- 18: **end for** **return** $\mathbf{k} := (k_1, \dots, k_m)$.
- 19: **end Procedure**

considers valid allocations by directly removing invalid allocations in advance in line 10 of Alg. 1.

With a little abuse of the notation, we use $r(\mathbf{k})$ to represent the reward $r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k})$ of a given problem instance. Let $\mathbf{k}^* \in \mathbb{Z}_{\geq 0}^m$ denote the optimal budget allocation and $\mathbf{k}^j \in \mathbb{Z}_{\geq 0}^m$ denote the budget allocation before entering the j -th iteration of the while loop (line 6-13) in Alg. 4. After entering the j -th iteration, the algorithm tries to extend the current budget allocation \mathbf{k}^j by choosing the pair (i^*, b^*) in line 7, which we denote as (i^j, b^j) . Let s be the first iteration we can not extend the current solution, i.e., $\mathbf{k}^s = \mathbf{k}^{s+1}$ and $\mathbf{k}^j < \mathbf{k}^{j+1}$ for $j = 1, \dots, s-1$. If we can always extend the current solution, we set s to be $B+1$. For analysis, we temporarily add (i^s, b^s) (in the algorithm, this pair is removed by line 13 in the s -th iteration) to form a "virtual" budget allocation $\mathbf{k}^{s+1} = \mathbf{k}^s + b^s \chi_{i^s}$. Let $\mathbf{k}_g \in \mathbb{Z}_{\geq 0}^m$ denote the solution returned by Alg. 4.

We first introduce lemmas describing two important properties given by the submodularity over the integer lattice.

Lemma 6. (Soma et al., 2014). *Let $f : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R}$ be a*

submodular function. For any $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_{\geq 0}^m$, we have,

$$\begin{aligned} f(\mathbf{x} \vee \mathbf{y}) &\leq f(\mathbf{x}) + \sum_{i \in \text{supp}^+(\mathbf{y} - \mathbf{x})} (f(\mathbf{x} + (y_i - x_i)\chi_i) - f(\mathbf{x})). \end{aligned} \quad (18)$$

Lemma 7. (Soma et al., 2014). Let $f : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R}$ be a monotone submodular function. For any $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_{\geq 0}^m$ with $\mathbf{x} \leq \mathbf{y}$ and $i \in [m]$ we have,

$$f(\mathbf{x} \vee k\chi_i) - f(\mathbf{x}) \geq f(\mathbf{y} \vee k\chi_i) - f(\mathbf{y}). \quad (19)$$

Then, we have the following lemma.

Lemma 8. For $j = 1, \dots, s$,

$$r(\mathbf{k}^{j+1}) \geq (1 - \frac{b^j}{B})r(\mathbf{k}^j) + \frac{b^j}{B}r(\mathbf{k}^*). \quad (20)$$

Proof of Lemma 8. This is because

$$\begin{aligned} r(\mathbf{k}^*) &\leq r(\mathbf{k}^* \vee \mathbf{k}^j) \\ &\leq r(\mathbf{k}^j) + \sum_{i \in \text{supp}^+(\mathbf{k}^* - \mathbf{k}^j)} (r(\mathbf{k}^j + (k_i^* - k_i^j)\chi_i) - r(\mathbf{k}^j)) \\ &= r(\mathbf{k}^j) + \sum_{i \in \text{supp}^+(\mathbf{k}^* - \mathbf{k}^j)} (r(\mathbf{k}^j + \alpha_i\chi_i) - r(\mathbf{k}^j)) \\ &\quad (\text{Let } \alpha_i = k_i^* - k_i^j) \\ &\leq r(\mathbf{k}^j) + \sum_{i \in \text{supp}^+(\mathbf{k}^* - \mathbf{k}^j)} \left(\alpha_i \frac{r(\mathbf{k}^{j+1}) - r(\mathbf{k}^j)}{b^j} \right) \\ &\leq r(\mathbf{k}^j) + B \left(\frac{r(\mathbf{k}^{j+1}) - r(\mathbf{k}^j)}{b^j} \right). \end{aligned} \quad (21)$$

The second inequality comes from Lemma 6, the third inequality holds because of the greedy procedure in line 7 and the last inequality holds because $\sum_{i \in \text{supp}^+(\mathbf{k}^* - \mathbf{k}^j)} \alpha_i \leq B$. By rearranging terms, Inequality (20) holds. ■

Next, We can prove the following lemma.

Lemma 9. For $l = 1, \dots, s$,

$$\begin{aligned} r(\mathbf{k}^{l+1}) &\geq r(\mathbf{k}^1) \prod_{j=1}^l \left(1 - \frac{b^j}{B} \right) \\ &\quad + r(\mathbf{k}^*) \left(1 - \prod_{j=1}^l \left(1 - \frac{b^j}{B} \right) \right). \end{aligned} \quad (22)$$

Proof of Lemma 9. We can prove this lemma by induction on l . When $l = 1$, the lemma holds due to Lemma 8. Assume that the lemma holds for $l-1$, we have the following inequality holds,

$$\begin{aligned} r(\mathbf{k}^{l+1}) &\geq (1 - \frac{b^l}{B})r(\mathbf{k}^l) + \frac{b^l}{B}r(\mathbf{k}^*) \\ &\geq (1 - \frac{b^l}{B})r(\mathbf{k}^1) \prod_{j=1}^{l-1} \left(1 - \frac{b^j}{B} \right) \\ &\quad + (1 - \frac{b^l}{B})r(\mathbf{k}^*) \left(1 - \prod_{j=1}^{l-1} \left(1 - \frac{b^j}{B} \right) \right) + \frac{b^l}{B}r(\mathbf{k}^*) \\ &= r(\mathbf{k}^1) \prod_{j=1}^l \left(1 - \frac{b^j}{B} \right) + r(\mathbf{k}^*) \left(1 - \prod_{j=1}^l \left(1 - \frac{b^j}{B} \right) \right), \end{aligned} \quad (23)$$

where the first inequality is due to Lemma 8 by setting $j = l$, and the second inequality is by the assumption for $l-1$. By induction, Lemma 9 holds. ■

We then consider the following cases.

Case 1. Suppose the total budget used for \mathbf{k}^s is larger or equal to ηB , i.e., $\sum_{j=1}^{s-1} b^j \geq \eta B$, where $\eta \in [0, 1]$.

We have the following inequality.

$$r(\mathbf{k}^s) \geq (1 - e^{-\eta})r(\mathbf{k}^*). \quad (24)$$

This is due to Lem. 9 by setting $l = s-1$, combined with the fact that $\mathbf{k}^1 = \mathbf{0}$ and $\prod_{j=1}^{s-1} \left(1 - \frac{b^j}{B} \right) \leq e^{-\eta}$. The later fact holds because,

$$\begin{aligned} \log \left(\prod_{j=1}^{s-1} \left(1 - \frac{b^j}{B} \right) \right) &= (s-1) \sum_{j=1}^{s-1} \frac{1}{s-1} \log \left(1 - \frac{b^j}{B} \right) \\ &\leq (s-1) \log \left(1 - \frac{1}{s-1} \sum_{j=1}^{s-1} \frac{b^j}{B} \right) \\ &\leq (s-1) \log \left(1 - \frac{\eta}{s-1} \right), \end{aligned}$$

where the first inequality holds because of the Jensen's Inequality (Raginsky et al., 2013) and the second inequality holds because $\sum_{j=1}^{s-1} \frac{b^j}{B} \geq \eta$. Then we can easily check $\prod_{j=1}^{s-1} \left(1 - \frac{b^j}{B} \right) \leq (1 - \frac{\eta}{s-1})^{s-1} \leq e^{-\eta}$.

Case 2. Suppose the total budget $\sum_{j=1}^{s-1} b^j \leq \eta B$. Then, we have $b^s > (1 - \eta)B$. We can prove the following inequality holds,

$$r(\mathbf{k}_g) \geq \left(1 - \frac{1}{2 - \eta} \right) r(\mathbf{k}^*). \quad (25)$$

This is due to Inequality (21), we have

$$\begin{aligned} r(\mathbf{k}^*) &\leq r(\mathbf{k}^s) + B \left(\frac{r(\mathbf{k}^{s+1}) - r(\mathbf{k}^s)}{b^s} \right) \\ &\leq r(\mathbf{k}^s) + \left(\frac{r(\mathbf{k}^{s+1}) - r(\mathbf{k}^s)}{1 - \eta} \right) \\ &\leq \left(1 + \frac{1}{1 - \eta} \right) r(\mathbf{k}_g), \end{aligned}$$

where the second inequality is due to $b^s > (1 - \eta)B$ and the last equality is due to the fact $r(\mathbf{k}^{s+1}) - r(\mathbf{k}^s) \leq r(\mathbf{k}_g)$, $r(\mathbf{k}^s) \leq r(\mathbf{k}_g)$. To see the above fact, we assume without loss of generality the pair (i^s, b^s) improves \mathbf{k}^s towards the optimal budget allocation, i.e., $k_{i^s}^{s+1} = b^s + k_{i^s}^s \leq k_{i^s}^*$. Otherwise, if $b^s + k_{i^s}^s > k_{i^s}^*$, we can safely delete (i^s, b^s) in the queue \mathcal{Q} and does not affect the greedy solution, the optimal solution and the analysis. Let $r_{max} = \max_{i \in [m]} r(c_i \chi_i)$, we have $r(\mathbf{k}^{s+1}) - r(\mathbf{k}^s) \leq r(\mathbf{k}^s \vee k_{i^s}^* \chi_{i^s}) - r(\mathbf{k}^s) \leq r(k_{i^s}^* \chi_{i^s}) - r(\mathbf{0}) \leq r_{max} \leq r(\mathbf{k}_g)$. Also, we can obtain that $r(\mathbf{k}^s) \leq r(\mathbf{k}_g)$ since $\mathbf{k}^s \leq \mathbf{k}_g$. By rearranging the terms, Inequality (25) holds.

Combining Inequality (24) and (25), we have

$$\begin{aligned} r(\mathbf{k}_g) &\geq \min_{\eta \in [0,1]} \max \left(1 - e^{-\eta}, 1 - \frac{1}{2 - \eta} \right) r(\mathbf{k}^*) \\ &\geq (1 - e^{-\eta}) r(\mathbf{k}^*), \end{aligned}$$

where η is the solution for equation $e^\eta = 2 - \eta$. ■

C.2. Starting From the Stationary Distribution

Lemma 4. *For any network \mathcal{G} , stationary distributions π and node weights σ , function $r_{\mathcal{G}, \pi, \sigma, (\cdot)} : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R}$ is monotone and DR-submodular.*

Proof. By definition, we need to show $r_{\mathcal{G}, \pi, \sigma}(\mathbf{y} + \chi_j) - r_{\mathcal{G}, \pi, \sigma}(\mathbf{y}) \leq r_{\mathcal{G}, \pi, \sigma}(\mathbf{x} + \chi_j) - r_{\mathcal{G}, \pi, \sigma}(\mathbf{x})$, and $r_{\mathcal{G}, \pi, \sigma}(\mathbf{x}) \leq r_{\mathcal{G}, \pi, \sigma}(\mathbf{y})$ for any $\mathbf{x} \leq \mathbf{y}$, $j \in [m]$.

We first introduce the following lemma to help us to prove the DR-submodularity.

Lemma 10. *Function f is DR-submodular if and only if*

$$\begin{aligned} f(\mathbf{x} + \chi_i) - f(\mathbf{x}) &\geq f(\mathbf{x} + \chi_j + \chi_i) - f(\mathbf{x} + \chi_j), \\ \text{for any } \mathbf{x} \in \mathbb{Z}_{\geq 0}^m. \end{aligned} \quad (26)$$

Proof of Lemma 10. (If part.) We can easily check this direction holds by using the similar argument for Inequality (16), where the only difference is we consider $i \in [m]$ instead of $i \in I_0$.

(Only if part.) We can set $\mathbf{y} := \mathbf{x}' + \chi_j$, $\mathbf{x} := \mathbf{x}' + \chi_i$, for any $i, j \in [m]$, $\mathbf{x}' \in \mathbb{Z}_{\geq 0}^m$, and use Inequality (15) to

show the only if part holds. ■

Since $r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k})$ is submodular for arbitrary starting distributions α , Inequality (14) holds. Then consider any layer $j \in [m]$ and budget allocation $\mathbf{x} \in \mathbb{Z}_{\geq 0}^m$, it is sufficient to show $r_{\mathcal{G}, \pi, \sigma}(\mathbf{x} + 2\chi_j) - r_{\mathcal{G}, \pi, \sigma}(\mathbf{x} + \chi_j) \leq r_{\mathcal{G}, \pi, \sigma}(\mathbf{x} + \chi_j) - r_{\mathcal{G}, \pi, \sigma}(\mathbf{x})$. Since the left hand side minus the right hand side equals to $\sum_{v \in \mathcal{V}} \sigma_v \left[\left(\prod_{i \in [m], i \neq j} (1 - P_{i,v}(x_i)) \right) \cdot (P_{j,v}(x_j) - 2P_{j,u}(x_j + 1) + P_{j,v}(x_j + 2)) \right]$, we only need to show $\sum_{v \in \mathcal{V}} \sigma_v \left[\left(\prod_{i \in [m], i \neq j} (1 - P_{i,v}(x_i)) \right) \cdot (g_{j,v}(x_j + 2) - g_{j,v}(x_j + 1)) \right] \leq 0$. The above inequality holds because of Lemma 2. ■

Theorem 2. *Algorithm 2 obtains a $(1 - 1/e)$ -approximate solution to the overlapping MulaNE with the stationary starting distributions.*

Proof. We can observe that line 7 always select the pair (i^*, b^*) with $b^* = 1$ because $r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k} + b\chi_i) - r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k} + (b - 1)\chi_i) \leq r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k} + \chi_i) - r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k})$ for arbitrary \mathbf{k} , i and b . Thus, we have $s = B + 1$ and by the similar argument for Inequality (24), we have $r(\mathbf{k}_g) = r(\mathbf{k}^{B+1}) \geq (1 - 1/e)r(\mathbf{k}^*)$, which completes the proof. ■

C.3. Efficiently Evaluating the Reward Function

One key issue to derive the budget allocation is to efficiently evaluate the reward function $r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k})$ and its marginal gains $\delta(i, b, \mathbf{k})$. Since we need to repetitively use the visiting probabilities $P_{i,u}(k_i)$, we pre-calculate $P_{i,u}(k_i)$ in $O(m \|c\|_\infty n_{max}^3)$ time in our algorithms based on Eq. (9), where $n_{max} = \max_i |\mathcal{V}_i|$. For the overlapping case, given the current budget allocation \mathbf{k} , we maintain a value $p_u = \prod_{i \in [m]} (1 - P_{i,u}(k_i))$ for each node $u \in \mathcal{V}$. The marginal gain $\delta(i, b, \mathbf{k}) = \sum_{u \in \mathcal{V}} \sigma_u \left[p_u \left(1 - \frac{1 - P_{i,u}(k_i + b)}{1 - P_{i,u}(k_i)} \right) \right] / b$ can be evaluated in $O(n_{max})$ time. Then, we update all $p_u = p_u \left(\frac{1 - P_{i,u}(k_i + b)}{1 - P_{i,u}(k_i)} \right)$ in $O(n_{max})$ after we allocate b more budgets to layer i . Therefore, we can use $O(n_{max})$ in total to evaluate $\delta(i, b, \mathbf{k})$. In practice, lazy evaluation (Krause et al., 2008) and parallel computing can be used to further accelerate our algorithm.

D. Budget Effective Greedy Algorithm With Partial Enumeration and Its Analysis

D.1. Algorithm

The algorithm is shown in Alg. 5.

D.2. Analysis

Theorem 4. *The Algorithm 5 obtains a $(1 - 1/e)$ -approximate solution to the overlapping MuLaNE with ar-*

Algorithm 5 Budget Effective Greedy Algorithm with Partial Enumeration (BEGE).

Input: Graph \mathcal{G} , starting distributions α , node weights σ , budget B , constraints c .

Output: Budget allocation \mathbf{k} .

- 1: Compute visiting probabilities $(P_{i,u}(b))_{i \in [m], u \in \mathcal{V}, b \in [c_i]}$ according to Eq. (7).
- 2: $\mathbf{k} \leftarrow \text{BEGE}((P_{i,u}(b))_{i \in [m], u \in \mathcal{V}, b \in [c_i]}, \sigma, B, c)$.

- 3: **Procedure** BEGE($(P_{i,u}(b))_{i \in [m], u \in \mathcal{V}, b \in [c_i]}, \sigma, B, c$)
- 4: Let $\mathbf{k}_{max} \leftarrow \mathbf{0}$.
- 5: $\mathcal{S} \leftarrow \{\mathbf{k} = (k_1, \dots, k_m) \mid 0 \leq k_i \leq c_i, \sum_{i \in [m]} k_i \leq B, \sum_{i \in [m]} \mathbb{I}\{k_i > 0\} \leq 3\}$. $\triangleright \mathcal{S}$ contains all partial solutions which allocate partial budgets to at most three layers
- 6: **for** $\mathbf{k} \in \mathcal{S}$ **do**
- 7: $K \leftarrow B - \sum_{i \in [m]} k_i$.
- 8: Let $\mathcal{Q} \leftarrow \{(i, b_i) \mid i \in [m], 1 \leq b_i \leq c_i - k_i\}$.
- 9: **while** $K > 0$ and $\mathcal{Q} \neq \emptyset$ **do**
- 10: $(i^*, b^*) \leftarrow \arg \max_{(i,b) \in \mathcal{Q}} \delta(i, b, \mathbf{k})/b$.
 $\triangleright \text{Eq. (11)}$
- 11: **if** $b^* \leq K$ **then**
- 12: $k_{i^*} \leftarrow k_{i^*} + b^*, K \leftarrow K - b^*$.
- 13: Modify all pairs $(i^*, b) \in \mathcal{Q}$ to $(i^*, b - b^*)$.
- 14: Remove all pairs $(i^*, b) \in \mathcal{Q}$ such that $b \leq 0$.
- 15: **else**
- 16: Remove (i^*, b^*) from \mathcal{Q} .
- 17: **end if**
- 18: **end while**
- 19: **if** $r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k}) > r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k}_{max})$, **then** $\mathbf{k}_{max} \leftarrow \mathbf{k}$.
- 20: **end forreturn** $\mathbf{k}_{max} := (k_1, \dots, k_m)$.
- 21: **end Procedure**

bitrary starting distributions.

Proof. Suppose that we start from a partial solution $\mathbf{k}^1 \in \mathbb{Z}_{\geq 0}^m \neq \mathbf{0}$. Let us first reorder the optimal solution \mathbf{k}^* according to a non-increasing marginal gain ordering. Namely, the marginal gain of the pair (s_1^*, b_1^*) with respect to the empty pair is the highest among all other pairs, the marginal gain of the pair (s_2^*, b_2^*) with respect to the solution consisting of pair (s_1^*, b_1^*) is the highest among all remaining pairs, and so on. To be concrete, $\mathbf{k}^* = \{(s_1^*, b_1^*), \dots, (s_m^*, b_m^*)\}$, where $s_i^* \in [m] \setminus \{s_1^*, \dots, s_{i-1}^*\}$ is selected to maximize the following equation:

$$\left(r \left(\sum_{j=1}^{i-1} b_j^* \chi_{s_j^*} + b^* \chi_{s_i^*} \right) - r \left(\sum_{j=1}^{i-1} b_j^* \chi_{s_j^*} \right) \right). \quad (27)$$

Then, we try to bound the "virtual" marginal gain $\Delta^s = r(\mathbf{k}^{s+1}) - r(\mathbf{k}^s)$, and recall s is the first iteration we can not extend the current solution. Without loss of generality, we assume the virtual pair (i^s, b^s) improves \mathbf{k}^s towards the

optimal budget allocation, i.e., $b^s + k_{i^s}^s \leq k_{i^s}^*$. Otherwise, if $b^s + k_{i^s}^s > k_{i^s}^*$, we can safely delete (i^s, b^s) in the queue \mathcal{Q} and does not affect the greedy solution, the optimal solution and the analysis.

If we start from the initial solution \mathbf{k}^1 such that \mathbf{k}^1 matches $(s_1^*, b_1^*), (s_2^*, b_2^*), (s_3^*, b_3^*)$, i.e., $\mathbf{k}^1_{s_i^*} = b_i^*$ for $i = 1, 2, 3$, we have $\Delta^s = r(\mathbf{k}^{s+1}) - r(\mathbf{k}^s) \leq r(\mathbf{k}^s \vee k_{i^s}^* \chi_{i^s}) - r(\mathbf{k}^s) \leq r(k_{i^s}^* \chi_{i^s}) - r(\mathbf{0}) \leq r(b_1^* \chi_{s_1^*})$. Moreover, $\Delta^s = r(\mathbf{k}^{s+1}) - r(\mathbf{k}^s) \leq r(\mathbf{k}^s \vee k_{i^s}^* \chi_{i^s}) - r(\mathbf{k}^s) = r(\mathbf{k}^s \vee b_1^* \chi_{s_1^*} \vee k_{i^s}^* \chi_{i^s}) - r(\mathbf{k}^s \vee b_1^* \chi_{s_1^*}) \leq r(b_1^* \chi_{s_1^*} \vee k_{i^s}^* \chi_{i^s}) - r(b_1^* \chi_{s_1^*}) \leq r(b_1^* \chi_{s_1^*} + b_2^* \chi_{s_2^*}) - r(b_1^* \chi_{s_1^*})$. Similarly, we have $\Delta^s \leq r(b_1^* \chi_{s_1^*} + b_2^* \chi_{s_2^*} + b_3^* \chi_{s_3^*}) - r(b_1^* \chi_{s_1^*} + b_2^* \chi_{s_2^*})$. By adding above inequalities, we have $\Delta^s \leq r(\mathbf{k}^1)/3$.

Now, we can use Lemma 9 by setting $l = s$ and the fact $\sum_{j=1}^s b^j \geq B$ to show $r(\mathbf{k}^{s+1}) \geq r(\mathbf{k}^1) + (1 - 1/e)(r(\mathbf{k}^s) - r(\mathbf{k}^1))$. Combining $\Delta^s \leq r(\mathbf{k}^1)/3$, we have $r(\mathbf{k}^s) \geq (1 - 1/3)r(\mathbf{k}^1) + (1 - 1/e)(r(\mathbf{k}^s) - r(\mathbf{k}^1)) \geq (1 - 1/e)r(\mathbf{k}^s)$, which completes the proof. \blacksquare

The time complexity of Alg. 5 is $O(B^4 m^4 \|c\|_\infty n_{max} + \|c\|_\infty m n_{max}^3)$. This is because the number of all partial solutions is $|\mathcal{S}| = O(B^3 m^3)$, and for any partial enumeration $\mathbf{k} \in \mathcal{S}$, the time complexity is the same order as the Alg. 1, i.e., $O(B \|c\|_\infty m n_{max})$, where $O(n_{max})$ is the time to evaluate $\delta(i, b, \mathbf{k})$ as discussed before.

E. Algorithms and Analysis of Offline Optimization for Non-overlapping MuLaNE

E.1. Starting From the Stationary Distribution

According to Eq. (8) and (10), the reward function can be rewritten as

$$r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k}) = \sum_{i \in [m]} \sum_{b \in [k_i]} g_i(b), \quad (28)$$

where $g_i(b)$ represents the *layer-level marginal gain* when we allocate one more budget (from $b - 1$ to b) to layer i , which is given by

$$g_i(b) = \sum_{v \in \mathcal{V}_i} \sigma_v (P_{i,v}(b) - P_{i,v}(b-1)). \quad (29)$$

From the definition of the reward function, we have two observations. First, budgets allocated to a layer will not affect the reward of other layers because layers are non-overlapping. Second, the layer-level marginal gain $g_i(b)$ for the i -th layer is non-increasing with respect to budget b . Based on these two observation, we propose the myopic greedy algorithm (Alg. 6), which is a slight modification of Alg. 2. Alg. 6 takes the multi-layered network $\mathcal{G} = (\mathcal{G}_1, \dots, \mathcal{G}_m)$, starting distributions $(\alpha_1, \dots, \alpha_m)$, node weights $\sigma = (\sigma_1, \dots, \sigma_{|\mathcal{V}|})$ and total budget B as inputs. It

consists of B rounds to compute the optimal budget allocation \mathbf{k}^* . In each round, line 7 in Alg. 6 selects the layer i^* with the largest layer-level marginal gain and allocate one budget to layer i^* until total B budgets are used up.

Algorithm 6 Myopic Greedy Algorithm for the Non-overlapping MuLaNE

Input: Network \mathcal{G} , starting distributions α , budget B , constraints \mathbf{c} .

Output: Budget allocation \mathbf{k} .

- 1: Compute visiting probabilities $(P_{i,u}(b))_{i \in [m], u \in \mathcal{V}, b \in [c_i]}$ according to Eq. (8).
 - 2: Compute layer-level marginal gain $g_i(b)_{i \in [m], b \in [c_i]}$ according to Eq. (29).
 - 3: $\mathbf{k} \leftarrow \text{MG}((g_i(b))_{i \in [m], b \in [c_i]}, B, \mathbf{c})$
 - 4: **Procedure** $\text{MG}((g_i(b))_{i \in [m], b \in [c_i]}, B, \mathbf{c})$
 - 5: Let $\mathbf{k} := \{k_1, \dots, k_m\} \leftarrow \mathbf{0}$, $K \leftarrow B$.
 - 6: **while** $K > 0$ **do**
 - 7: $i^* \leftarrow \arg \max_{i \in [m], k_i+1 \leq c_i} g_i(k_i+1) \cdot O(\log m)$
using the priority queue
 - 8: $k_{i^*} \leftarrow k_{i^*} + 1$, $K \leftarrow K - 1$.
 - 9: **end while** **return** $\mathbf{k} = (k_1, \dots, k_m)$.
 - 10: **end Procedure**
-

Theorem 5. *The Algorithm 6 obtains the optimal budget allocation to the non-overlapping MuLaNE with the stationary starting distributions.*

Proof. Define a two dimensional array $\mathbf{M} \in \mathbb{R}^{c_1+c_2+\dots+c_m}$, where (i, j) -th entry is the j -th step layer-level marginal gain for layer i , i.e., $\mathbf{M}[i, j] = \sum_{v \in \mathcal{V}_i} g_{i,v}(j)$, for $i \in [m], j \in [c_i]$. Given the budget allocation $\mathbf{k} = \{k_1, \dots, k_m\}$, the expected reward $r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k})$ can be written as the sum of elements in \mathbf{M} , i.e., $r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k}) = \sum_{i=1}^m \sum_{j=1}^{k_i} \mathbf{M}[i, j]$. Because $g_{i,v}(k_i)$ is non-increasing with respect to k_i , the element in each row or the layer-level marginal gain is non-increasing. Hence, the greedy method at step s choose the s -th largest element in \mathbf{M} . At step $s = B$, the greedy policy selects all B largest elements and the corresponding budget allocation maximizes the expected reward $r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k})$. ■

Alg. 6 uses $O(\|\mathbf{c}\|_{\infty} mn_{max}^3)$ to compute visiting probabilities $P_{i,u}(k_i)$ and $O(\|\mathbf{c}\|_{\infty} mn_{max})$ to compute layer-wise marginal gains $g_i(b)$. Then $\delta(i, 1, \mathbf{k})$ can be evaluated in $O(\log m)$ using the priority queue, which is repeated for B iterations. Therefore, the time complexity of Alg. 6 is $O(B \log m + \|\mathbf{c}\|_{\infty} mn_{max}^3)$.

E.2. Starting From Arbitrary Distributions

When random explorers start from any arbitrary distribution, Alg. 6 can not obtain the optimal solution. This is because the layer-level marginal gain $g_i(b)$ is *not* non-increasing

Algorithm 7 Dynamic Programming (DP) Algorithm for the Non-overlapping MuLaNE

Input: Network \mathcal{G} , starting distributions α , node weights σ , budget B , constraints \mathbf{c} .

Output: Budget allocation \mathbf{k} .

- 1: Compute visiting probabilities $(P_{i,u}(b))_{i \in [m], u \in \mathcal{V}, b \in [c_i]}$ according to Eq. (8).
 - 2: Compute layer-level marginal gain $(g_i(b))_{i \in [m], b \in [c_i]}$ according to Eq. (29).
 - 3: $\mathbf{k} \leftarrow \text{DP}((g_i(b))_{i \in [m], b \in [c_i]}, B, \mathbf{c})$
 - 4: **Procedure** $\text{DP}((g_i(b))_{i \in [m], b \in [c_i]}, B, \mathbf{c})$
 - 5: **for** $i \in 0 \cup [m], b \in 0 \cup [B]$. **do**
 - 6: Set $A[i, b] \leftarrow \mathbf{0}$, $V[i, b] \leftarrow 0$.
 - 7: **end for**
 - 8: **for** $i \leftarrow 0$ **to** $m - 1$ **do**
 - 9: **for** $b \in [B]$ **do**
 - 10: $j^* \leftarrow \arg \max_{j \in 0 \cup [c_{i+1}]} ((V[i, b - j] + r_{\mathcal{G}, \alpha, \sigma}(j \chi_{i+1})))$.
 - 11: $A[i + 1, b] \leftarrow A[i, b - j^*] + j^* \chi_{i+1}$.
 - 12: $V[i + 1, b] \leftarrow V[i, b - j^*] + r_{\mathcal{G}, \alpha, \sigma}(j^* \chi_{i+1})$.
 - 13: **end for**
 - 14: **end for** **return** $A[m, B]$.
 - 15: **end Procedure**
-

with respect to b . So we have to adopt a more general technique, dynamic programming (DP), to solve the budget allocation problem. The key idea is to keep a DP budget allocation table $A \in \mathbb{R}^{(m+1) \times (B+1)}$, where the (i, b) -th entry $A[i, b]$ saves the optimal budget allocation by allocating b budgets to the first i layers. Another DP table $V \in \mathbb{R}^{(m+1) \times (B+1)}$ saves the value of the optimal reward, where the (i, b) -th entry $V[i, b]$ corresponds to the optimal reward when setting $A[i, b]$ as the budget allocation. In the i' -th outer loop and b' -th inner loop in Alg 7, we have already obtained the optimal budget allocation for $A[i, b]$ with $i = 0, \dots, i'$, $b \in [B]$, which help us to find the optimal amount of budget j^* to $(i' + 1)$ -th layer (in line 10), and thus we can obtain the $A[i' + 1, b']$ and $V[i' + 1, b']$ accordingly.

Theorem 6. *Algorithm 7 obtains the optimal budget allocation to the non-overlapping MuLaNE with an arbitrary starting distribution.*

Alg. 7 uses $O(\|\mathbf{c}\|_{\infty} mn_{max}^3)$ to compute visiting probabilities $P_{i,u}(k_i)$ and $O(\|\mathbf{c}\|_{\infty} mn_{max})$ to compute layer-wise marginal gains $g_i(b)$. In the DP procedure, we can also precalculate all rewards $r_{\mathcal{G}, \alpha, \sigma}(j \chi_{i+1})$ in $O(\|\mathbf{c}\|_{\infty} m)$. Then DP procedure takes $O(\|\mathbf{c}\|_{\infty})$ to find j^* , which is repeated for Bm iterations. Therefore, the Alg. 7 has the time complexity of $O(B \|\mathbf{c}\|_{\infty} m + \|\mathbf{c}\|_{\infty} mn_{max}^3)$.

F. Analysis of Online Learning for Overlapping MuLaNE

F.1. Proof of 1-Norm Bounded Smoothness.

Property 1. (Monotonicity). *The reward $r_{\mu,\sigma}(\mathbf{k})$ is monotonically increasing, i.e., for any budget allocation \mathbf{k} , any two vectors $\mu = (\mu_{i,u,b})_{(i,u,b) \in \mathcal{A}}$, $\mu' = (\mu'_{i,u,b})_{(i,u,b) \in \mathcal{A}}$ and any node weights σ, σ' , we have $r_{\mu,\sigma}(\mathbf{k}) \leq r_{\mu',\sigma'}(\mathbf{k})$, if $\mu_{i,u,b} \leq \mu'_{i,u,b}$ and $\sigma_v \leq \sigma'_v, \forall (i, u, b) \in \mathcal{A}, v \in \mathcal{V}$.*

Proof. According to Eq. (7), and since $\mu_{i,u,b} \leq \mu'_{i,u,b}$, for $(i, u, b) \in \mathcal{A}$, we have $(1 - \prod_{i \in [m]} (1 - \mu_{i,u,k_i})) \leq (1 - \prod_{i \in [m]} (1 - \mu'_{i,u,k_i}))$. Therefore for any $\sigma_u \leq \sigma'_u$, we have $\sigma_u (1 - \prod_{i \in [m]} (1 - \mu_{i,u,k_i})) \leq \sigma'_u (1 - \prod_{i \in [m]} (1 - \mu'_{i,u,k_i}))$

By summing up both sides over $u \in \mathcal{V}$, we have $r_{\mathbf{k}}(\mu, \sigma) \leq r_{\mathbf{k}}(\mu', \sigma')$. ■

Property 2. (1-Norm Bounded Smoothness). *The reward function $r_{\mu,\sigma}(\mathbf{k})$ satisfies the 1-norm bounded smoothness condition, i.e., for any budget allocation \mathbf{k} , any two vectors $\mu = (\mu_{i,u,b})_{(i,u,b) \in \mathcal{A}}$, $\mu' = (\mu'_{i,u,b})_{(i,u,b) \in \mathcal{A}}$ and any node weights σ, σ' , we have $|r_{\mu,\sigma}(\mathbf{k}) - r_{\mu',\sigma'}(\mathbf{k})| \leq \sum_{i \in [m], u \in \mathcal{V}, b = k_i} (\sigma_u |\mu_{i,u,b} - \mu'_{i,u,b}| + |\sigma_u - \sigma'_u| \mu'_{i,u,b})$.*

Proof. The left-hand side:

$$\begin{aligned} & |r_{\mu,\sigma}(\mathbf{k}) - r_{\mu',\sigma'}(\mathbf{k})| \\ &= \left| \sum_{u \in \mathcal{V}} \left[\sigma_u \left(1 - \prod_{i \in [m]} (1 - \mu_{i,u,k_i}) \right) - \sigma'_u \left(1 - \prod_{i \in [m]} (1 - \mu'_{i,u,k_i}) \right) \right] \right| \\ &\leq \sum_{u \in \mathcal{V}} \left| \sigma_u (1 - \prod_{i \in [m]} q_{i,u}) - \sigma'_u (1 - \prod_{i \in [m]} q'_{i,u}) \right| \\ &= \sum_{u \in \mathcal{V}} \left| \sigma_u (1 - \prod_{i \in [m]} q_{i,u}) - \sigma_u (1 - \prod_{i \in [m]} q'_{i,u}) \right| \\ &\quad + \sum_{u \in \mathcal{V}} \left| \sigma_u (1 - \prod_{i \in [m]} q'_{i,u}) - \sigma'_u (1 - \prod_{i \in [m]} q'_{i,u}) \right| \\ &\leq \sum_{u \in \mathcal{V}} \sigma_u \left| \prod_{i \in [m]} q_{i,u} - \prod_{i \in [m]} q'_{i,u} \right| \\ &\quad + \sum_{u \in \mathcal{V}} |\sigma_u - \sigma'_u| (1 - \prod_{i \in [m]} q'_{i,u}), \end{aligned}$$

where $q_{i,u} = (1 - \mu_{i,u,k_i})$, $q'_{i,u} = (1 - \mu'_{i,u,k_i})$.

For the first term, it can be derived that,

$$\begin{aligned} & \left| \sigma_u \left(\prod_{i \in [m]} q_{i,u} - \prod_{i \in [m]} q'_{i,u} \right) \right| \\ &= \sigma_u \left| \sum_{i=1}^m \left(\left(\prod_{j=1}^{i-1} q'_{j,u} \right) (q_{i,u} - q'_{i,u}) \left(\prod_{j=i+1}^m q_{j,u} \right) \right) \right| \\ &\leq \sigma_u \sum_{i=1}^m \left| \left(\prod_{j=1}^{i-1} q'_{j,u} \right) (q_{i,u} - q'_{i,u}) \left(\prod_{j=i+1}^m q_{j,u} \right) \right| \\ &\leq \sigma_u \sum_{i=1}^m |q_{i,u} - q'_{i,u}| \\ &= \sigma_u \sum_{i \in [m]} |\mu_{i,u,k_i} - \mu'_{i,u,k_i}|, \end{aligned}$$

where last inequality holds because $q_{j,u}, q'_{j,u} \leq 1$ and $\sigma_u \in [0, 1]$, for any $j \in [m], u \in \mathcal{V}$.

For the second term, we have

$$\begin{aligned} & |\sigma_u - \sigma'_u| (1 - \prod_{i \in [m]} q'_{i,u}) \\ &= |\sigma_u - \sigma'_u| (1 - \prod_{i \in [m]} (1 - \mu'_{i,u,k_i})) \\ &\leq |\sigma_u - \sigma'_u| \sum_{i \in [m]} \mu'_{i,u,k_i}, \end{aligned}$$

where the inequality holds because of the Weierstrass product inequality. Plugging the above two terms into the previous inequality, 1-Norm Bounded Smoothness is satisfied. ■

F.2. Regret Analysis for Overlapping MuLaNE

In this section, we give the regret analysis for overlapping MuLaNE.

F.2.1. FACTS

We utilize the following tail bound in our analysis.

Fact 1 (Hoeffding's Inequality (Dubhashi & Panconesi, 2009)). *Let Y_1, \dots, Y_n be independent and identically distributed (i.e., i.i.d) random variables with common support $[0, 1]$ and mean μ . Let $Z = Y_1 + \dots + Y_n$. Then for all δ ,*

$$\Pr\{|Z - n\mu| \geq \delta\} \leq 2e^{-2\delta^2/n}.$$

F.2.2. PROOF DETAILS

Let \mathcal{A} denote the set containing all base arms, i.e., $\mathcal{A} = \{(i, u, b) | i \in [m], u \in \mathcal{V}, b \in [c_i]\}$. We add subscript t to denote the value of a variable at the end of round $t \in [T]$, e.g. $\bar{\sigma}_t, \hat{\mu}_{i,u,b,t}$, where T is the total number of rounds. For example, $T_{i,u,b,t}$ denotes the total number of times that arm

(i, u, b) is played at the end of round t . Let us first introduce a definition of an unlikely event, where $\hat{\mu}_{i,u,b,t-1}$ is not accurate as expected.

Definition 2. We say that the sampling is nice at the beginning of round t , if for every arm $(i, u, b) \in \mathcal{A}$, $|\hat{\mu}_{i,u,b,t-1} - \mu_{i,u,b}| < \rho_{i,u,b,t}$, where $\rho_{i,u,b,t} = \sqrt{\frac{3 \ln t}{2T_{i,u,b,t-1}}}$. Let \mathcal{N}_t be such event.

Lemma 11. For each round $t \geq 1$, $\Pr\{\neg \mathcal{N}_t\} \leq 2|\mathcal{A}|t^{-2}$

Proof. For each round $t \geq 1$, we have

$$\begin{aligned} & \Pr\{\neg \mathcal{N}_t\} \\ &= \Pr\{\exists (i, u, b) \in \mathcal{A}, |\hat{\mu}_{i,u,b,t-1} - \mu_{i,u,b}| \geq \sqrt{\frac{3 \ln t}{2T_{i,u,b,t-1}}}\} \\ &\leq \sum_{(i,u,b) \in \mathcal{A}} \Pr\{|\hat{\mu}_{i,u,b,t-1} - \mu_{i,u,b}| \geq \sqrt{\frac{3 \ln t}{2T_{i,u,b,t-1}}}\} \\ &= \sum_{(i,u,b) \in \mathcal{A}} \sum_{k=1}^{t-1} \Pr\{T_{i,u,b,t-1} = k \\ &\quad \wedge |\hat{\mu}_{i,u,b,t-1} - \mu_{i,u,b}| \geq \sqrt{\frac{3 \ln t}{2T_{i,u,b,t-1}}}\} \\ &\leq \sum_{(i,u,b) \in \mathcal{A}} \sum_{k=1}^{t-1} \frac{2}{k^3} \leq 2|\mathcal{A}|t^{-2}. \end{aligned}$$

Given $T_{i,u,b,t-1} = k$, $\hat{\mu}_{i,u,b,t-1}$ is the average of k i.i.d. random variables $Y_{i,u,b}^{[1]}, \dots, Y_{i,u,b}^{[k]}$, where $Y_{i,u,b}^{[j]}$ is the Bernoulli random variable (computed in Alg. 3 line 12) when the arm (i, u, b) is played for the j -th time during the execution. That is, $\hat{\mu}_{i,u,b,t-1} = \sum_{j=1}^k Y_{i,u,b}^{[j]}/k$. Note that the independence of $Y_{i,u,b}^{[1]}, \dots, Y_{i,u,b}^{[k]}$ comes from the fact we select a new starting position following the starting distribution α_i at the beginning of each round t . And the last inequality uses the Hoeffding Inequality (Fact 1). ■

Then we can use monotonicity and 1-norm bounded smoothness properties (Property 1 and Property 2) to bound the reward gap $\Delta_{\mathbf{k}_t} = \xi r_{\mu, \sigma}(\mathbf{k}^*) - r_{\mu, \sigma}(\mathbf{k}_t)$ between the optimal action \mathbf{k}^* and the action $\mathbf{k}_t = (k_{t,1}, \dots, k_{t,m})$ selected by our algorithm A . To achieve this, we introduce a positive real number $M_{i,u,b}$ for each arm (i, u, b) and define $M_{\mathbf{k}_t} = \max_{(i,u,b) \in S_t} M_{i,u,b}$, where $S_t = \{(i, u, b) \in \mathcal{A} | b = k_{t,i}\}$. Define $|\mathcal{E}'| = |S_t| = \sum_{i \in [m]} |\mathcal{V}| = m|\mathcal{V}|$, and let

$$\kappa_T(M, s) = \begin{cases} 2, & \text{if } s = 0, \\ 3\sqrt{\frac{6 \ln T}{s}}, & \text{if } 1 \leq s \leq \ell_T(M), \\ 0, & \text{if } s \geq \ell_T(M) + 1, \end{cases}$$

where

$$\ell_T(M) = \lfloor \frac{54|\mathcal{E}'|^2 \ln T}{M^2} \rfloor.$$

Lemma 12. If $\{\Delta_{\mathbf{k}_t} \geq M_{\mathbf{k}_t}\}$ and \mathcal{N}_t holds, we have

$$\Delta_{\mathbf{k}_t} \leq \sum_{(i,u,b) \in S_t} \kappa_T(M_{i,u,b}, T_{i,u,b,t-1}). \quad (30)$$

Proof. First, we can observe, for $(i, u, b) \in \mathcal{A}$, $\mu_{i,u,b}$ is upper bounded by $\bar{\mu}_{i,u,b,t}$, i.e., $\bar{\mu}_{i,u,b,t} \geq \mu_{i,u,b}$, when \mathcal{N}_t holds. This is due to

$$\begin{aligned} \bar{\mu}_{i,u,b,t} &= \max_{j \in [b]} \tilde{\mu}_{i,u,j,t} \\ &\geq \tilde{\mu}_{i,u,b,t} \\ &= \min\{\hat{\mu}_{i,u,b,t-1} + \rho_{i,u,b,t}, 1\} \\ &\geq \mu_{i,u,b}, \end{aligned}$$

where the last inequality comes from the fact that $|\hat{\mu}_{i,u,b,t-1} - \mu_{i,u,b}| < \rho_{i,u,b,t}$ when \mathcal{N}_t holds.

Then, notice that the right hand of Inequality (30) is non-negative, it is trivially satisfied when $\Delta_{\mathbf{k}_t} = 0$. We only need to consider $\Delta_{\mathbf{k}_t} > 0$. By \mathcal{N}_t and Property 1, We have the following inequalities,

$$r_{\bar{\mu}_t, \bar{\sigma}}(\mathbf{k}_t) \geq \xi r_{\bar{\mu}_t, \bar{\sigma}}(\mathbf{k}^*) \geq \xi r_{\mu, \sigma}(\mathbf{k}^*) = \Delta_{\mathbf{k}_t} + r_{\mu, \sigma}(\mathbf{k}_t).$$

The first inequality is due to the oracle outputs the (ξ, β) -approximate solution given parameters $\bar{\mu}_t$, where $\xi = 1 - e^{-\eta}$, $\beta = 1$ by definition. The second inequality is due to Condition 1 (monotonicity).

Then, by Condition 2, we have

$$\begin{aligned} \Delta_{\mathbf{k}_t} &\leq r_{\bar{\mu}_t, \bar{\sigma}}(\mathbf{k}_t) - r_{\mu, \sigma}(\mathbf{k}_t) \\ &\leq \sum_{(i,u,b) \in S_t} (\sigma_u |\bar{\mu}_{i,u,b,t} - \mu_{i,u,b}| + |\bar{\sigma}_u - \sigma_u| \bar{\mu}_{i,u,b,t}) \\ &\leq \sum_{(i,u,b) \in S_t} (|\bar{\mu}_{i,u,b,t} - \mu_{i,u,b}| + \rho_{i,u,b,t}), \end{aligned}$$

where the last inequality is due to the following observation: $\bar{\sigma}_u \neq \sigma_u$ if and only if u has not been visited so far, so $\hat{\mu}_{i,u,b,t-1} = 0$ for any $i \in [m], b \in [B]$ and $\bar{\mu}_{i,u,b,t} = \max_{j \in [b]} \{\hat{\mu}_{i,u,j,t-1} + \rho_{i,u,j,t}\} = \rho_{i,u,b,t}$.

Next, we can bound $\Delta_{\mathbf{k}_t}$ by bounding $\bar{\mu}_{i,u,b,t} - \mu_{i,u,b} + \rho_{i,u,b,t}$ by applying a transformation. As $\{\Delta_{\mathbf{k}_t} \geq M_{\mathbf{k}_t}\}$ holds by assumption, so $\sum_{(i,u,b) \in S_t} |\bar{\mu}_{i,u,b,t} - \mu_{i,u,b}| +$

$\rho_{i,u,b,t} \geq \Delta_{\mathbf{k}_t} \geq M_{\mathbf{k}_t}$. We have

$$\begin{aligned}
 \Delta_{\mathbf{k}_t} &\leq \sum_{(i,u,b) \in \mathcal{S}_t} (|\bar{\mu}_{i,u,b,t} - \mu_{i,u,b}| + \rho_{i,u,b,t}) \\
 &\leq -M_{\mathbf{k}_t} + 2 \sum_{(i,u,b) \in \mathcal{S}_t} |\bar{\mu}_{i,u,b,t} - \mu_{i,u,b}| + \rho_{i,u,b,t} \\
 &= 2 \sum_{(i,u,b) \in \mathcal{S}_t} (|\bar{\mu}_{i,u,b,t} - \mu_{i,u,b}| + \rho_{i,u,b,t} - \frac{M_{\mathbf{k}_t}}{2|\mathcal{E}'|}) \\
 &\leq 2 \sum_{(i,u,b) \in \mathcal{S}_t} (|\bar{\mu}_{i,u,b,t} - \mu_{i,u,b}| + \rho_{i,u,b,t} - \frac{M_{i,u,b}}{2|\mathcal{E}'|}). \tag{31}
 \end{aligned}$$

By \mathcal{N}_t , for all $(i, u, b) \in \mathcal{A}$, we have:

$$\begin{aligned}
 &|\bar{\mu}_{i,u,b,t} - \mu_{i,u,b}| + \rho_{i,u,b,t} - \frac{M_{i,u,b}}{2|\mathcal{E}'|} \\
 &= \tilde{\mu}_{i,u,l,t} - \mu_{i,u,b} + \rho_{i,u,b,t} - \frac{M_{i,u,b}}{2|\mathcal{E}'|} \\
 &\quad (l = \arg \max_{j \in [b]} \tilde{\mu}_{i,u,j,t}) \\
 &\leq \tilde{\mu}_{i,u,l,t} - \mu_{i,u,l} + \rho_{i,u,b,t} - \frac{M_{i,u,b}}{2|\mathcal{E}'|} \\
 &\leq \min\{2\rho_{i,u,l,t} + \rho_{i,u,b,t}, 1\} - \frac{M_{i,u,b}}{2|\mathcal{E}'|} \\
 &\leq \min\{3\rho_{i,u,b,t}, 1\} - \frac{M_{i,u,b}}{2|\mathcal{E}'|} \\
 &\leq \min\left\{3\sqrt{\frac{3 \ln T}{2T_{i,u,b,t-1}}}, 1\right\} - \frac{M_{i,u,b}}{2|\mathcal{E}'|}.
 \end{aligned}$$

where the first inequality is due to $\mu_{i,u,j} \leq \mu_{i,u,b}$ for $j \leq b$, the second inequality is due to the event \mathcal{N}_t and the third inequality holds because $T_{i,u,j,t-1} \geq T_{i,u,b,t-1}$ for $j \leq b$.

Now consider the following cases:

Case 1. If $T_{i,u,b,t-1} \leq \ell_T(M_{i,u,b})$, we have $\bar{\mu}_{i,u,b,t} - \mu_{i,u,b} + \rho_{i,u,b,t} - \frac{M_{i,u,b}}{2|\mathcal{E}'|} \leq \min\left\{3\sqrt{\frac{3 \ln T}{2T_{i,u,b,t-1}}}, 1\right\} \leq \frac{1}{2}\kappa_T(M_{i,u,b}, T_{i,u,b,t-1})$.

Case 2. If $T_{i,u,b,t-1} \geq \ell_T(M_{i,u,b}) + 1$, then $3\sqrt{\frac{3 \ln T}{2T_{i,u,b,t-1}}} \leq \frac{M_{i,u,b}}{2|\mathcal{E}'|}$, so $\bar{\mu}_{i,u,b,t} - \mu_{i,u,b} + \rho_{i,u,b,t} - \frac{M_{i,u,b}}{2|\mathcal{E}'|} \leq 0 = \frac{1}{2}\kappa_T(M_{i,u,b}, T_{i,u,b,t-1})$.

Combine these two cases and inequality (31), we have the following inequality, which concludes the lemma.

$$\Delta_{\mathbf{k}_t} \leq \sum_{(i,u,b) \in \mathcal{S}_t} \kappa_T(M_{i,u,b}, T_{i,u,b,t-1}).$$

■

Theorem 3. Algorithm 3 has the following distribution-

dependent $(1 - e^{-\eta}, 1)$ approximation regret,

$$\text{Reg}_{\mu, \sigma}(T) \leq \sum_{(i,u,b) \in \mathcal{A}} \frac{108m|\mathcal{V}|\ln T}{\Delta_{\min}^{i,u,b}} + 2|\mathcal{A}| + \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max}.$$

Proof of Theorem 3. By above lemmas, we have

$$\begin{aligned}
 &\sum_{t=1}^T \mathbb{I}(\{\Delta_{\mathbf{k}_t} \geq M_{\mathbf{k}_t}\} \wedge \mathcal{N}_t) \Delta_{\mathbf{k}_t} \\
 &\leq \sum_{t=1}^T \sum_{(i,u,b) \in \mathcal{S}_t} \kappa_T(M_{i,u,b}, T_{i,u,b,t-1}) \\
 &= \sum_{(i,u,b) \in \mathcal{A}} \sum_{t' \in \{t | t \in [T], \mathbf{k}_t, i=b\}} \kappa_T(M_{i,u,b}, T_{i,u,b,t'-1}) \\
 &\leq \sum_{(i,u,b) \in \mathcal{A}} \sum_{s=0}^{T_{i,u,b,T}} \kappa_T(M_{i,u,b}, s) \\
 &\leq 2|\mathcal{A}| + \sum_{(i,u,b) \in \mathcal{A}} \sum_{s=1}^{\ell_T(M_{i,u,b})} 3\sqrt{6\frac{\ln T}{s}} \\
 &\leq 2|\mathcal{A}| + \sum_{(i,u,b) \in \mathcal{A}} \int_{s=0}^{\ell_T(M_{i,u,b})} 3\sqrt{6\frac{\ln T}{s}} ds \\
 &\leq 2|\mathcal{A}| + \sum_{(i,u,b) \in \mathcal{A}} 6\sqrt{6 \ln T \ell_T(M_{i,u,b})} \\
 &\leq 2|\mathcal{A}| + \sum_{(i,u,b) \in \mathcal{A}} \frac{108|\mathcal{E}'| \ln T}{M_{i,u,b}}.
 \end{aligned}$$

For distribution dependent bound, let us set $M_{i,u,b} = \Delta_{\min}^{i,u,b}$ and thus the event $\{\Delta_{\mathbf{k}_t} \geq M_{\mathbf{k}_t}\}$ always holds, i.e., $\mathbb{I}\{\Delta_{\mathbf{k}_t} < M_{\mathbf{k}_t}\} = 0$. We have

$$\begin{aligned}
 &\sum_{t=1}^T \mathbb{I}(\mathcal{N}_t \wedge \{\Delta_{\mathbf{k}_t} \geq M_{\mathbf{k}_t}\}) \Delta_{\mathbf{k}_t} \\
 &\leq 2|\mathcal{A}| + \sum_{(i,u,b) \in \mathcal{A}} \frac{108|\mathcal{E}'| \ln T}{\Delta_{\min}^{i,u,b}}.
 \end{aligned}$$

By Lem. 11, $\Pr\{-\mathcal{N}_t\} \leq 2|\mathcal{A}|t^{-2}$, then

$$\mathbb{E} \left[\sum_{t=1}^T \mathbb{I}(-\mathcal{N}_t) \Delta_{\mathbf{k}_t} \right] \leq \sum_{t=1}^T 2|\mathcal{A}|t^{-2} \Delta_{\max} \leq \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max}.$$

By our choice of $M_{i,u,b}$,

$$\sum_{t=1}^T \mathbb{I}(\Delta_{\mathbf{k}_t} < M_{\mathbf{k}_t}) \Delta_{\mathbf{k}_t} = 0.$$

By the definition of the regret,

$$\begin{aligned}
 \text{Reg}_\mu(T) &= \mathbb{E} \left[\sum_{t=1}^T \Delta_{\mathbf{k}_t} \right] \\
 &\leq \mathbb{E} \left[\sum_{t=1}^T \left(\mathbb{I}(\neg \mathcal{N}_t) + \mathbb{I}(\Delta_{\mathbf{k}_t} < M_{\mathbf{k}_t}) \right. \right. \\
 &\quad \left. \left. + \mathbb{I}(\mathcal{N}_t \wedge \{\Delta_{\mathbf{k}_t} \geq M_{\mathbf{k}_t}\}) \right) \Delta_{\mathbf{k}_t} \right] \\
 &\leq \sum_{(i,u,b) \in \mathcal{A}} \frac{108|\mathcal{E}'| \ln T}{\Delta_{\min}^{i,u,b}} + 2|\mathcal{A}| + \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max} \\
 &= \sum_{(i,u,b) \in \mathcal{A}} \frac{108m|\mathcal{V}| \ln T}{\Delta_{\min}^{i,u,b}} + 2|\mathcal{A}| + \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max}.
 \end{aligned}$$

For distribution independent bound, we take $M_{i,u,b} = M = \sqrt{108|\mathcal{E}'||\mathcal{A}| \ln T/T}$, then we have $\sum_{t=1}^T \mathbb{I}\{\Delta_{\mathbf{k}_t} \leq M_{\mathbf{k}_t}\} \leq TM$.

Then

$$\begin{aligned}
 \text{Reg}_\mu(T) &\leq \sum_{(i,u,b) \in \mathcal{A}} \frac{108|\mathcal{E}'| \ln T}{M_{i,u,b}} + 2|\mathcal{A}| + \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max} + TM \\
 &\leq \frac{108|\mathcal{E}'||\mathcal{A}| \ln T}{M} + 2|\mathcal{A}| + \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max} + TM \\
 &= 2\sqrt{108|\mathcal{E}'||\mathcal{A}|T \ln T} + \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max} + 2|\mathcal{A}| \\
 &\leq 21\sqrt{m|\mathcal{V}||\mathcal{A}|T \ln T} + \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max} + 2|\mathcal{A}|.
 \end{aligned}$$

■

F.3. Extension for Random Node Weights

In this section, we extend the deterministic node weight σ_u for node $u \in \mathcal{V}$ to a random weight. Specifically, we denote $\tilde{\sigma}_u \in [0, 1]$ as the random weight when we first visit u (after the first visit we will not get any reward so on so forth), and denote $\sigma_u := \mathbb{E}[\tilde{\sigma}_u]$.

For the offline setting, the reward function in Eq.(7) remains the same since $\tilde{\sigma}_u$ is independent of all other random variables. For the online setting, since the reward function remains unchanged, property 1 and property 2 still hold. However, we need to change the way we maintain the optimistic node weight $\bar{\sigma}_u$.

We present our CUCB-MAX-R algorithm in Alg. 8. Compared with Alg. 3, the major difference is that $\bar{\sigma}_u$ now represent the upper confidence bound (UCB) value of σ_u . In line 2, we denote T'_v as the number of times node v is played so far and $\hat{\sigma}_v$ as the empirical mean of v 's weight. In line 9, ρ'_v is the confidence radius and $\bar{\sigma}_v$ is the UCB value of node

Algorithm 8 CUCB-MAX-R: Combinatorial Upper Confidence Bound algorithm for the overlapping MuLaNE, with Random node weights

Input: Budget B , number of layers m , number of nodes $|\mathcal{V}|$, constraints c , offline oracle BEG.

- 1: For each arm $(i, u, b) \in \mathcal{A}$, $T_{i,u,b} \leftarrow 0$, $\hat{\mu}_{i,u,b} \leftarrow 0$.
- 2: For each node $v \in \mathcal{V}$, $T'_v \leftarrow 0$, $\hat{\sigma}_v \leftarrow 0$.
- 3: **for** $t = 1, 2, 3, \dots, T$ **do**
- 4: **for** $(i, u, b) \in \mathcal{A}$ **do**
- 5: $\rho_{i,u,b} \leftarrow \sqrt{3 \ln t / (2T_{i,u,b})}$.
- 6: $\tilde{\mu}_{i,u,b} \leftarrow \min\{\hat{\mu}_{i,u,b} + \rho_{i,u,b}, 1\}$.
- 7: **end for**
- 8: **for** $v \in \mathcal{V}$ **do**
- 9: $\rho'_v \leftarrow \sqrt{3 \ln t / (2T'_v)}$.
- 10: $\hat{\sigma}_v \leftarrow \min\{\hat{\sigma}_v + \rho'_v, 1\}$.
- 11: **end for**
- 12: For $(i, u, b) \in \mathcal{A}$, $\bar{\mu}_{i,u,b} \leftarrow \max_{j \in [b]} \tilde{\mu}_{i,u,j}$.
- 13: $\mathbf{k} \leftarrow \text{BEG}((\bar{\mu}_{i,u,b})_{(i,u,b) \in \mathcal{A}}, (\bar{\sigma}_v)_{v \in \mathcal{V}}, B, c)$.
- 14: Apply budget allocation \mathbf{k} , which gives trajectories $\mathbf{X} := (X_{i,1}, \dots, X_{i,k_i})_{i \in [m]}$ as feedbacks.
- 15: For any visited node $v \in \bigcup_{i \in [m]} \{X_{i,1}, \dots, X_{i,k_i}\}$, receive its random node weight $\tilde{\sigma}_v$, update $T'_v \leftarrow T'_v + 1$, $\hat{\sigma}_v \leftarrow \hat{\sigma}_v + (\tilde{\sigma}_v - \hat{\sigma}_v) / T'_v$.
- 16: For any $(i, u, b) \in \tau := \{(i, u, b) \in \mathcal{A} | b \leq k_i\}$, $Y_{i,u,b} \leftarrow 1$ if $u \in \{X_{i,1}, \dots, X_{i,b}\}$ and 0 otherwise.
- 17: For $(i, u, b) \in \tau$, update $T_{i,u,b}$ and $\hat{\mu}_{i,u,b}$:
 $T_{i,u,b} \leftarrow T_{i,u,b} + 1$, $\hat{\mu}_{i,u,b} \leftarrow \hat{\mu}_{i,u,b} + (Y_{i,u,b} - \hat{\mu}_{i,u,b}) / T_{i,u,b}$.
- 18: **end for**

weight σ_u . In line 15, we update the empirical mean $\hat{\sigma}_v$ for any visited v .

Regret analysis. Let \mathcal{A} denote the set containing all base arms for visiting probabilities, i.e., $\mathcal{A} = \{(i, u, b) | i \in [m], u \in \mathcal{V}, b \in [c_i]\}$. With a bit of abuse of notation, we also use \mathcal{V} to denote the set of base arms corresponding to the node weights. Therefore, $\mathcal{A} \cup \mathcal{V}$ is the set of all base arms, which is different from the Section F.2 that only has \mathcal{A} as base arms and does not have base arms for random weights. Following the notations of (Wang & Chen, 2017), we define $X_a \in [0, 1]$ be the random outcome of base arm $a \in \mathcal{A} \cup \mathcal{V}$ and define D as the unknown joint distribution of random outcomes $\mathbf{X} = (X_a)_{a \in \mathcal{A} \cup \mathcal{V}}$ with unknown mean $\mu_a = \mathbb{E}_{\mathbf{X} \sim D}[X_a]$. Let \mathbf{k} be any feasible budget allocation, we denote $p_a^{D,\mathbf{k}}$ the probability that action \mathbf{k} triggers arm a (i.e. outcome X_a is observed) when the unknown environment distribution is D . Specifically, $p_a^{D,\mathbf{k}} = 1$ for $a \in \{(i, u, b) | i \in [m], u \in \mathcal{V}, b = k_i\}$, $p_a^{D,\mathbf{k}} = 1 - \prod_{i \in [m]} (1 - \mu_{i,a,k_i})$ for $a \in \mathcal{V}$ and $p_i^{D,\mathbf{k}} = 0$ for the rest of base arms. Therefore, we can rewrite the

inequality in property 2 as

$$\begin{aligned}
 & |r_{\mu, \sigma}(\mathbf{k}) - r_{\mu', \sigma'}(\mathbf{k})| \\
 & \leq \sum_{i \in [m], u \in \mathcal{V}, b = k_i} (\sigma'_u |\mu_{i,u,b} - \mu'_{i,u,b}|) \\
 & + \sum_{u \in \mathcal{V}} |\sigma_u - \sigma'_u| \left(1 - \prod_{i \in [m]} (1 - \mu_{i,u,b}) \right) \\
 & \leq \sum_{a \in \mathcal{A}} p_a^{D, \mathbf{k}} |\mu_a - \mu'_a|, \tag{32}
 \end{aligned}$$

where we abuse the notation to denote μ_a as σ_a and μ'_a as σ'_a if $a \in \mathcal{A}_2$ in the last inequality. We can observe that the above rewritten property 2 can be viewed as a special case of the 1-norm TPM bounded smoothness condition in (Wang & Chen, 2017).

Now we can apply the similar proof in Section B.3 in (Wang & Chen, 2017), together with how we deal with the $|\bar{\mu}_{i,u,b,t} - \mu_{i,u,b}|$ in Section F.2.2 to get a distribution-dependent regret bound.

Let $K = m|\mathcal{V}| + |\mathcal{V}|$ be the maximum number of triggered arms, $|\mathcal{A} \cup \mathcal{V}| = Bm|\mathcal{V}| + |\mathcal{V}|$ be the number of base arms and $\lceil x \rceil_0 = \max\{\lceil x \rceil, 0\}$. For any feasible budget allocation \mathbf{k} and any base arm $a \in \mathcal{A} \cup \mathcal{V}$, the gap $\Delta_{\mathbf{k}} = \max\{0, \xi r_{\mu, \sigma}(\mathbf{k}^*) - r_{\mu, \sigma}(\mathbf{k})\}$ and we define $\Delta_{\min}^a = \inf_{\mathbf{k}: p_a^{D, \mathbf{k}} > 0, \Delta_{\mathbf{k}} > 0} \Delta_{\mathbf{k}}$, $\Delta_{\max}^a = \sup_{\mathbf{k}: p_a^{D, \mathbf{k}} > 0, \Delta_{\mathbf{k}} > 0} \Delta_{\mathbf{k}}$. As convention, if there is no action \mathbf{k} such that $p_a^{D, \mathbf{k}} > 0$ and $\Delta_{\mathbf{k}} > 0$, we define $\Delta_{\min}^a = +\infty$, $\Delta_{\max}^a = 0$. We define $\min_{a \in \mathcal{A}} \Delta_{\min}^a$ and $\max_{a \in \mathcal{A}} \Delta_{\max}^a$.

Theorem 7. *Algorithm 8 has the following distribution-dependent $(1 - e^{-\eta}, 1)$ approximation regret,*

$$\begin{aligned}
 \text{Reg}_{\mu, \sigma}(T) & \leq \sum_{a \in \mathcal{A} \cup \mathcal{V}} \frac{576K \ln T}{\Delta_{\min}^a} \\
 & + 4|\mathcal{A} \cup \mathcal{V}| + \sum_{a \in \mathcal{A} \cup \mathcal{V}} \left(\lceil \log_2 \frac{2K}{\Delta_{\min}^a} \rceil_0 + 2 \right) \frac{\pi^2}{6} \Delta_{\max}^a,
 \end{aligned}$$

We can also have a distribution-independent bound,

$$\begin{aligned}
 \text{Reg}_{\mu, \sigma}(T) & \leq 12\sqrt{|\mathcal{A} \cup \mathcal{V}|KT \ln T} \\
 & + 2|\mathcal{A} \cup \mathcal{V}| + \left(\lceil \log_2 \frac{T}{18 \ln T} \rceil_0 + 2 \right) \frac{\pi^2}{6} \Delta_{\max},
 \end{aligned}$$

Remark. Compared with the regret bound in Thm. 3 in the main text, we can see the base arms now become $\mathcal{A} \cup \mathcal{V}$ instead of \mathcal{A} , which is worse than the regret bound of the fixed unknown weight setting. This is a trade-off since we need more estimation to handle the uncertain random node weight, which incurs additional regrets. It also shows our non-trivial analysis of the fixed unknown weights carefully merge the error term into base arms \mathcal{A} and incurs only constant-factor of extra regrets.

G. Online Learning for Non-overlapping MuLaNE

G.1. Online Learning Algorithm for Non-overlapping MuLaNE

For non-overlapping MuLaNE, we estimate layer-level marginal gains as our parameters, rather than probabilities $P_{i,u}(b)$ and node weights σ . By doing so, we can largely simplify our analysis since we no longer need to handle the unknown weights and the extra constraint of $P_{i,u}(b)$ as in the overlapping setting. Concretely, we maintain a set of base arms $\mathcal{A} = \{(i, b) | i \in [m], b \in [c_i]\}$, and let $|\mathcal{A}| = \sum_{i \in [m]} c_i$ be the total number of these arms. For each base arm $(i, b) \in \mathcal{A}$, denote $\mu_{i,b}$ as the truth value of each base arm, i.e., $\mu_{i,b} = \sum_{u \in \mathcal{V}} \sigma_u (P_{i,u}(b) - P_{i,u}(b-1))$. Based on the definition of base arms, we can see that we do not need to explicitly estimate the node weights σ or probabilities $P_{i,u}(b)$, so we will use $r_{\mu}(\mathbf{k})$ to denote the reward function $r_{\mathcal{G}, \alpha, \sigma}(\mathbf{k})$.

We present our algorithm in Alg. 9, which has three differences compared with the Alg. 3 for overlapping MuLaNE. First, the expected reward of action \mathbf{k} is $r_{\mu}(\mathbf{k}) = \sum_{i=1}^m \sum_{b=1}^{k_i} \mu_{i,b}$, which is a linear reward function and satisfies two new Monotonicity and 1-Norm Bounded Smoothness conditions (see Condition 1 and 2). Second, we use the Dynamic Programming method (Alg. 7) as our Oracle, which does not require the parameters $\mu_{i,b}$ to be increasing. Moreover, the offline oracle always outputs *optimal* solutions. Third, we use a new random variable $Y_{i,b}$, which indicates the gain of weights which depends on whether a new node is visited by random walker from the i -th layer exactly at the b -th step, i.e., $Y_{i,b} = \sum_{u \in \cup_{j=1}^b \{X_{i,j}\}} \sigma_u - \sum_{u \in \cup_{j=1}^{b-1} \{X_{i,j}\}} \sigma_u$ for base arms (i, b) such that $b \leq k_i$.

Condition 1 (Monotonicity). *The reward $r_{\mu}(\mathbf{k})$ is monotonically increasing, i.e., for any budget allocation \mathbf{k} and any two vectors $\mu = (\mu_{i,b})_{(i,b) \in \mathcal{A}}$, $\mu' = (\mu'_{i,b})_{(i,b) \in \mathcal{A}}$, we have $r_{\mu}(\mathbf{k}) \leq r_{\mu'}(\mathbf{k})$, if $\mu_{i,b} \leq \mu'_{i,b}$, for $(i, b) \in \mathcal{A}$.*

Condition 2 (1-Norm Bounded Smoothness). *The reward function $r_{\mu}(\mathbf{k})$ satisfies the 1-norm bounded smoothness, i.e., for any budget allocation \mathbf{k} and any two vectors $\mu = (\mu_{i,b})_{(i,b) \in \mathcal{A}}$, $\mu' = (\mu'_{i,b})_{(i,b) \in \mathcal{A}}$, we have $|r_{\mu}(\mathbf{k}) - r_{\mu'}(\mathbf{k})| \leq \sum_{b \leq k_i} |\mu_{i,b} - \mu'_{i,b}|$.*

We define the reward gap $\Delta_{\mathbf{k}} = r_{\mu}(\mathbf{k}^*) - r_{\mu}(\mathbf{k})$ for all feasible action \mathbf{k} satisfying $\sum_{i=1}^m k_i = B$, $0 \leq k_i \leq c_i$. For each base arm (i, b) , we define $\Delta_{\min}^{i,b} = \min_{\Delta_{\mathbf{k}} > 0, k_i \geq b} \Delta_{\mathbf{k}}$ and $\Delta_{\max}^{i,b} = \max_{\Delta_{\mathbf{k}} > 0, k_i \geq b} \Delta_{\mathbf{k}}$. As a convention, if there is no action with $k_i \geq b$ such that $\Delta_{\mathbf{k}} > 0$, we define $\Delta_{\min}^{i,b} = \infty$ and $\Delta_{\max}^{i,b} = 0$. Also, we define $\Delta_{\min} = \min_{(i,b) \in \mathcal{A}} \Delta_{\min}^{i,b}$ and $\Delta_{\max} = \max_{(i,b) \in \mathcal{A}} \Delta_{\max}^{i,b}$.

Theorem 8. *CUCB-MG has the following distribution-*

Algorithm 9 CUCB-MG: Combinatorial Upper Confidence Bound (CUCB) algorithm for non-overlapping MuLaNE, using Marginal Gains as the base arms

Input: Budget B , number of layers m , offline oracle DP.
 1: For each $(i, b) \in \mathcal{A}$, $T_{i,b} \leftarrow 0$, $\hat{\mu}_{i,b} \leftarrow 1$.
 2: **for** $t = 1, 2, 3, \dots, T$ **do**
 3: **for** $(i, b) \in \mathcal{A}$ **do**
 4: $\rho_{i,b} \leftarrow \sqrt{3 \ln t / (2T_{i,b})}$.
 5: $\hat{\mu}_{i,b} = \min\{\hat{\mu}_{i,b} + \rho_{i,b}, 1\}$.
 6: **end for**
 7: $\mathbf{k} \leftarrow \text{DP}((\bar{\mu}_{i,b})_{i,b}, B, c)$
 8: Apply budget allocation \mathbf{k} , which gives trajectories $\mathbf{X} := (X_{i,1}, \dots, X_{i,k_i})_{i \in [m]}$ as feedbacks.
 9: For $(i, b) \in \tau := \{(i, b) \in \mathcal{A} | b \leq k_i\}$,
 $Y_{i,b} = \sum_{u \in \cup_{j=1}^b \{X_{i,j}\}} \sigma_u - \sum_{u \in \cup_{j=1}^{b-1} \{X_{i,j}\}} \sigma_u$.
 10: For $(i, b) \in \tau$, update $T_{i,b}$ and $\hat{\mu}_{i,b}$:
 11: $T_{i,b} = T_{i,b} + 1$, $\hat{\mu}_{i,b} = \hat{\mu}_{i,b} + (Y_{i,b} - \hat{\mu}_{i,b}) / T_{i,b}$.
 12: **end for**

dependent regret bound,

$$\text{Reg}_\mu(T) \leq \sum_{(i,b) \in \mathcal{A}} \frac{48B \ln T}{\Delta_{\min}^{i,b}} + 2|\mathcal{A}| + \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max}.$$

G.2. Regret Analysis for Non-overlapping MuLaNE

In this section, we give the regret analysis for CUCB-MG, which applies the standard CUCB algorithm (Wang & Chen, 2017) to this setting.

G.2.1. PROOF DETAILS

Let \mathcal{A} denote the set containing all base arms, i.e., $\mathcal{A} = \{(i, b) | i \in [m], b \in [c_i]\}$. We add subscript t to denote the value of a variable at the end of t . For example, $T_{i,b,t}$ denotes the total times of arm $(i, b) \in \mathcal{A}$ is played at the end of round t . Let us first introduce a definition of an unlikely event that $\hat{\mu}_{i,b,t-1}$ is not accurate as expected.

Definition 3. We say that the sampling is nice at the beginning of round t , if for every arm $(i, b) \in \mathcal{A}$, $|\hat{\mu}_{i,b,t-1} - \mu_{i,b}| < \rho_{i,b,t}$, where $\rho_{i,b,t} = \sqrt{\frac{3 \ln t}{2T_{i,b,t-1}}}$. Let \mathcal{N}_t be such event.

Lemma 13. For each round $t \geq 1$, $\Pr\{\neg \mathcal{N}_t\} \leq 2|\mathcal{A}|t^{-2}$

Proof. For each round $t \geq 1$, we have

$$\begin{aligned} & \Pr\{\neg \mathcal{N}_t\} \\ &= \Pr\{\exists (i, b) \in \mathcal{A}, |\hat{\mu}_{i,b,t-1} - \mu_{i,b}| \geq \sqrt{\frac{3 \ln t}{2T_{i,b,t-1}}}\} \\ &\leq \sum_{(i,b) \in \mathcal{A}} \Pr\{|\hat{\mu}_{i,b,t-1} - \mu_{i,b}| \geq \sqrt{\frac{3 \ln t}{2T_{i,b,t-1}}}\} \\ &= \sum_{(i,b) \in \mathcal{A}} \sum_{k=1}^{t-1} \Pr\{T_{i,b,t-1} = k \\ &\quad \wedge |\hat{\mu}_{i,b,t-1} - \mu_{i,b}| \geq \sqrt{\frac{3 \ln t}{2T_{i,b,t-1}}}\} \\ &\leq \sum_{(i,b) \in \mathcal{A}} \sum_{k=1}^{(t-1)} \frac{2}{t^3} \leq 2|\mathcal{A}|t^{-2}. \end{aligned}$$

When $T_{i,b,t-1} = k$, $\hat{\mu}_{i,b,t-1}$ is the average of k i.i.d. random variables $Y_{i,b}^{[1]}, \dots, Y_{i,b}^{[k]}$, where $Y_{i,b}^{[j]}$ is the Bernoulli random variable of arm (i, b) when it is played for the j -th time during the execution. That is, $\hat{\mu}_{i,b,t-1} = \sum_{j=1}^k Y_{i,b}^{[j]} / k$. Note that the independence of $Y_{i,b}^{[1]}, \dots, Y_{i,b}^{[k]}$ comes from the fact we select a new starting position following the starting distribution α_i at the beginning of each round t . The last inequality uses the Hoeffding Inequality (Fact 1). ■

Then we can use monotonicity and 1-norm bounded smoothness property (Condition 1 and Condition 2) to bound the reward gap $\Delta_{\mathbf{k}_t} = r_\mu(\mathbf{k}^*) - r_\mu(\mathbf{k}_t)$ between the optimal action \mathbf{k}^* and the action $\mathbf{k}_t = (k_{t,1}, \dots, k_{t,m})$ selected by our algorithm A . To achieve this, we introduce a positive real number $M_{i,b}$ for each arm (i, b) and define $M_{\mathbf{k}_t} = \max_{(i,b) \in S_t} M_{i,b}$, where $S_t = \{(i, b) \in \mathcal{A} | b \leq k_{t,i}\}$ and $|S_t| = B$. Define,

$$\kappa_T(M, s) = \begin{cases} 2, & \text{if } s = 0, \\ 2\sqrt{\frac{6 \ln T}{s}}, & \text{if } 1 \leq s \leq \ell_T(M), \\ 0, & \text{if } s \geq \ell_T(M) + 1, \end{cases}$$

where

$$\ell_T(M) = \lfloor \frac{24B^2 \ln T}{M^2} \rfloor.$$

Lemma 14. If $\{\Delta_{\mathbf{k}_t} \geq M_{\mathbf{k}_t}\}$ and \mathcal{N}_t holds, we have

$$\Delta_{\mathbf{k}_t} \leq \sum_{(i,b) \in S_t} \kappa_T(M_{i,b}, T_{i,b,t-1}). \quad (33)$$

Proof. First, we can observe $\bar{\mu}_{i,b,t} \geq \mu_{i,b}$, for $(i, b) \in \mathcal{A}$, when \mathcal{N}_t holds. This comes from the fact $\bar{\mu}_{i,b,t} = \min\{\hat{\mu}_{i,b,t-1} + \rho_{i,b,t}, 1\}$ and $|\hat{\mu}_{i,b,t-1} - \mu_{i,b}| < \rho_{i,b,t}$.

Notice that the right hand of Inequality (33) is non-negative, it is trivially satisfied when $\Delta_{\mathbf{k}_t} = 0$. We only need to consider $\Delta_{\mathbf{k}_t} > 0$. By \mathcal{N}_t and Condition 1, We have the following inequalities,

$$r_{\bar{\mu}_t}(\mathbf{k}_t) \geq r_{\bar{\mu}_t}(\mathbf{k}^*) \geq r_{\mu}(\mathbf{k}^*) = \Delta_{\mathbf{k}_t} + r_{\mu}(\mathbf{k}_t).$$

The first inequality is due to the *oracle* outputs the optimal solution given parameters $\bar{\mu}_t$, the second inequality is due to Condition 1 (monotonicity).

Then, by Condition 2, we have

$$\Delta_{\mathbf{k}_t} \leq r_{\bar{\mu}_t}(\mathbf{k}_t) - r_{\mu}(\mathbf{k}_t) \leq \sum_{(i,b) \in S_t} |\bar{\mu}_{i,b,t} - \mu_{i,b}|.$$

Next, we can bound $\Delta_{\mathbf{k}_t}$ by bounding $\bar{\mu}_{i,b,t} - \mu_{i,b}$ by applying a transformation. As $\{\Delta_{\mathbf{k}_t} \geq M_{\mathbf{k}_t}\}$ holds by assumption, so $\sum_{(i,b) \in S_t} |\bar{\mu}_{i,b,t} - \mu_{i,b}| \geq \Delta_{\mathbf{k}_t} \geq M_{\mathbf{k}_t}$. We have

$$\begin{aligned} \Delta_{\mathbf{k}_t} &\leq \sum_{(i,b) \in S_t} |\bar{\mu}_{i,b,t} - \mu_{i,b}| \\ &\leq -M_{\mathbf{k}_t} + 2 \sum_{(i,b) \in S_t} |\bar{\mu}_{i,b,t} - \mu_{i,b}| \\ &= 2 \sum_{(i,b) \in S_t} (|\bar{\mu}_{i,b,t} - \mu_{i,b}| - \frac{M_{\mathbf{k}_t}}{2B}) \\ &\leq 2 \sum_{(i,b) \in S_t} (|\bar{\mu}_{i,b,t} - \mu_{i,b}| - \frac{M_{i,b}}{2B}). \end{aligned} \quad (34)$$

By \mathcal{N}_t , we have:

$$\begin{aligned} |\bar{\mu}_{i,b,t} - \mu_{i,b}| - \frac{M_{i,b}}{2B} &\leq \min\{2\rho_{i,b,t}, 1\} - \frac{M_{i,b}}{2B} \\ &\leq \min\left\{2\sqrt{\frac{3 \ln T}{2T_{i,b,t-1}}}, 1\right\} - \frac{M_{i,b}}{2B}. \end{aligned}$$

Now consider the following cases:

Case 1. If $T_{i,b,t-1} \leq \ell_T(M_{i,b})$, we have $\bar{\mu}_{i,b,t} - \mu_{i,b} - \frac{M_{i,b}}{2B} \leq \min\left\{2\sqrt{\frac{3 \ln T}{2T_{i,b,t-1}}}, 1\right\} \leq \frac{1}{2}\kappa_T(M_{i,b}, T_{i,b,t-1})$.

Case 2. If $T_{i,b,t-1} \geq \ell_T(M_{i,b,t}) + 1$, then $2\sqrt{\frac{3 \ln T}{2T_{i,b,t-1}}} \leq \frac{M_{i,b}}{2B}$, so $\bar{\mu}_{i,b,t} - \mu_{i,b} - \frac{M_{i,b}}{2B} \leq 0 = \frac{1}{2}\kappa_T(M_{i,b}, T_{i,b,t-1})$.

Combine these two cases and inequality (34), we have the following inequality, which concludes the lemma.

$$\Delta_{\mathbf{k}_t} \leq \sum_{(i,b) \in S_t} \kappa_T(M_{i,b}, T_{i,b,t-1}).$$

■

Theorem 8. *CUCB-MG has the following distribution-*

dependent regret bound,

$$\text{Reg}_{\mu}(T) \leq \sum_{(i,b) \in \mathcal{A}} \frac{48B \ln T}{\Delta_{\min}^{i,b}} + 2|\mathcal{A}| + \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max}.$$

Proof of Theorem 8. By above lemmas, we have

$$\begin{aligned} &\sum_{t=1}^T \mathbb{I}(\{\Delta_{\mathbf{k}_t} \geq M_{\mathbf{k}_t}\} \wedge \mathcal{N}_t) \Delta_{\mathbf{k}_t} \\ &\leq \sum_{t=1}^T \sum_{(i,b): (i,b) \in S_t} \kappa_T(M_{i,b}, T_{i,b,t-1}) \\ &= \sum_{(i,b) \in \mathcal{A}} \sum_{t' \in \{t | t \in [T], k_{t,i} \geq b\}} \kappa_T(M_{i,b}, T_{i,b,t'-1}) \\ &\leq \sum_{(i,b) \in \mathcal{A}} \sum_{s=0}^{T_{i,b,T}} \kappa_T(M_{i,b}, s) \\ &\leq 2|\mathcal{A}| + \sum_{(i,b) \in \mathcal{A}} \sum_{s=1}^{\ell_T(M_{i,b})} 2\sqrt{6\frac{\ln T}{s}} \\ &\leq 2|\mathcal{A}| + \sum_{(i,b) \in \mathcal{A}} \int_{s=0}^{\ell_T(M_{i,b})} 2\sqrt{6\frac{\ln T}{s}} ds \\ &\leq 2|\mathcal{A}| + \sum_{(i,b) \in \mathcal{A}} 4\sqrt{6 \ln T \ell_T(M_{i,b})} \\ &\leq 2|\mathcal{A}| + \sum_{(i,b) \in \mathcal{A}} \frac{48B \ln T}{M_{i,b}}. \end{aligned}$$

Let us set $M_{i,b} = \Delta_{\min}^{i,b}$ and thus the event $\{\Delta_{\mathbf{k}_t} \geq M_{\mathbf{k}_t}\}$ always holds. We have

$$\begin{aligned} &\sum_{t=1}^T \mathbb{I}(\mathcal{N}_t \wedge \{\Delta_{\mathbf{k}_t} \geq M_{\mathbf{k}_t}\}) \Delta_{\mathbf{k}_t} \\ &\leq 2|\mathcal{A}| + \sum_{(i,b) \in \mathcal{A}} \frac{48B \ln T}{\Delta_{\min}^{i,b}}. \end{aligned}$$

By Lem. 13, $Pr\{\neg \mathcal{N}_t\} \leq 2|\mathcal{A}|t^{-2}$, then we have

$$\mathbb{E} \left[\sum_{t=1}^T \mathbb{I}(\neg \mathcal{N}_t) \Delta_{\mathbf{k}_t} \right] \leq \sum_{t=1}^T 2|\mathcal{A}|t^{-2} \Delta_{\max} \leq \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max}.$$

By our choice of $M_{i,b}$,

$$\sum_{t=1}^T \mathbb{I}(\Delta_{\mathbf{k}_t} < M_{\mathbf{k}_t}) \Delta_{\mathbf{k}_t} = 0.$$

By the definition of the regret,

$$\begin{aligned}
 \text{Reg}_\mu(T) &= \mathbb{E} \left[\sum_{t=1}^T \Delta_{\mathbf{k}_t} \right] \\
 &\leq \mathbb{E} \left[\sum_{t=1}^T \left(\mathbb{I}(\neg \mathcal{N}_t) + \mathbb{I}(\Delta_{\mathbf{k}_t} < M_{\mathbf{k}_t}) \right. \right. \\
 &\quad \left. \left. + \mathbb{I}(\mathcal{N}_t \wedge \{\Delta_{\mathbf{k}_t} \geq M_{\mathbf{k}_t}\}) \right) \Delta_{\mathbf{k}_t} \right] \\
 &\leq \sum_{(i,b) \in \mathcal{A}} \frac{48B \ln T}{\Delta_{\min}^{i,b}} + 2|\mathcal{A}| + \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max}.
 \end{aligned}$$

For distribution independent bound, we take $M_{i,b} = M = \sqrt{48B|\mathcal{A}| \ln T/T}$, then we have $\sum_{t=1}^T \mathbb{I}\{\Delta_{\mathbf{k}_t} \leq M_{\mathbf{k}_t}\} \leq TM$.

Then

$$\begin{aligned}
 \text{Reg}_\mu(T) &\leq \sum_{(i,b) \in \mathcal{A}} \frac{48B \ln T}{M_{i,b}} + 2|\mathcal{A}| + \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max} + TM \\
 &\leq \frac{48B|\mathcal{A}| \ln T}{M} + 2|\mathcal{A}| + \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max} + TM \\
 &= 2\sqrt{48B|\mathcal{A}|T \ln T} + \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max} + 2|\mathcal{A}| \\
 &\leq 14\sqrt{B|\mathcal{A}|T \ln T} + \frac{\pi^2}{3} |\mathcal{A}| \Delta_{\max} + 2|\mathcal{A}|.
 \end{aligned}$$

■

H. Supplemental Experiments

H.1. Experimental Results for Stationary Distributions

In this section, we show experimental results with explorers starting from stationary distributions in Fig. 3. Specifically, each explorer W_i starts from the node u in layer L_i with probability proportional to the out-degree of u . For offline overlapping case, the total weights of unique nodes visited for BEG and MG are the same, and outperforms two baselines, which accords with our theoretical analysis. Similar to the Sec. 6, BEG and MG are empirically close to the optimal solution. For offline non-overlapping case, DP and MG give us the same *optimal* solution. For the online learning algorithms, note that we still use BEG rather than MG as the offline oracle, this is due to the UCB values do not have the diminishing return properties. The same applies for the non-overlapping case, where DP (instead of MG) is used as the offline oracle. The trend of the curves and the analysis are similar to that in the Sec. 6.

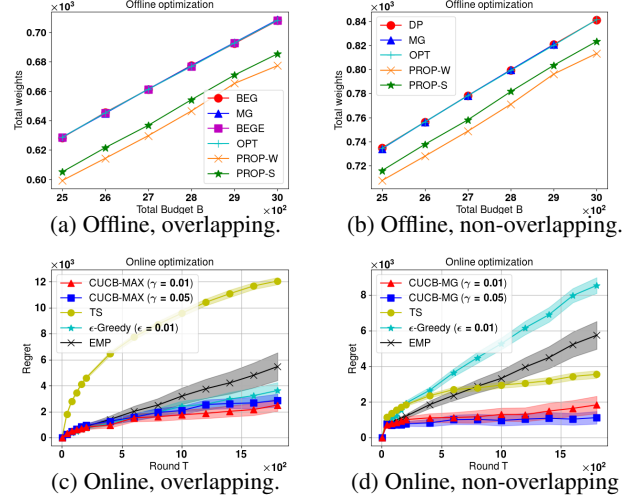


Figure 3. Explorers start with stationary distributions. Above: total weights of unique nodes visited given by different offline algorithms. Below: regret for different online algorithms when $B = 3000$.

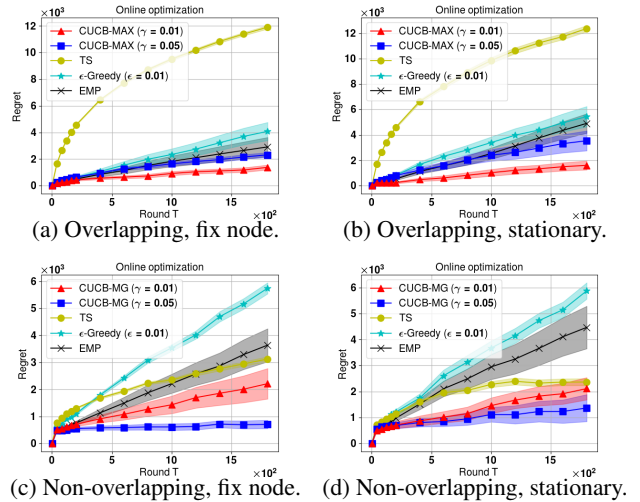


Figure 4. Regret for different online algorithms, when total budget $B = 2000$.

H.2. Experimental Results for Online Learning Algorithms with Different Budgets B

As shown in Fig. 4 and Fig. 5, although there are some fluctuations, the trend of the curves and the analysis are almost the same as that when $B = 3000$, which shows our experimental results are consistent for different budgets.

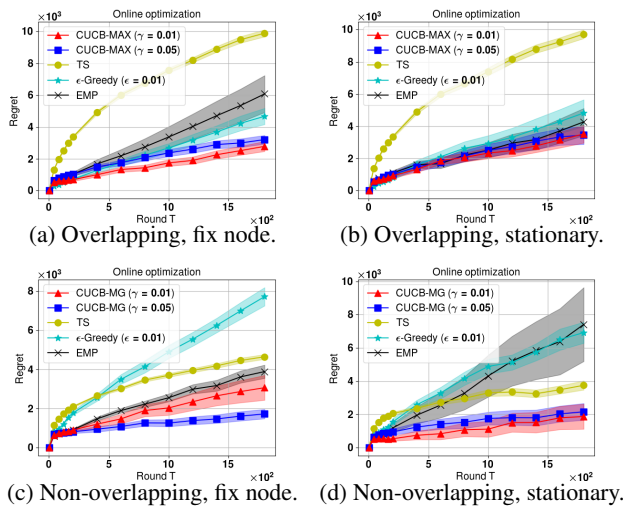


Figure 5. Regret for different online algorithms, when total budget $B = 4000$.