Stochastic Graph Exploration *

Aris Anagnostopoulos
Sapienza University of Rome
aris@diag.uniroma1.it

Ilan R. Cohen
CWI, Amsterdam
ilanrcohen@gmail.com

Stefano Leonardi [1]
Sapienza University of Rome
leonardi@diag.uniroma1.it

Jakub Łącki
Google Research, New York
jlacki@google.com

Abstract

Exploring large-scale networks is a time consuming and expensive task which is usually operated in a complex and uncertain environment. A crucial aspect of network exploration is the development of suitable strategies that decide which nodes and edges to probe at each stage of the process.

To model this process, we introduce the *stochastic graph exploration problem*. The input is an undirected graph $G = (V, E)$ with a source vertex s, stochastic edge costs drawn from a distribution $\pi_e, e \in E$, and rewards on vertices of maximum value R. The goal is to find a set F of edges of total cost at most B such that the subgraph of G induced by F is connected, contains s, and maximizes the total reward. This problem generalizes the stochastic knapsack problem and other stochastic probing problems recently studied.

Our focus is on the development of efficient nonadaptive strategies that are competitive against the optimal adaptive strategy. A major challenge is the fact that the problem has an $\Omega(n)$ adaptivity gap even on a tree of n vertices. This is in sharp contrast with $O(1)$ adaptivity gap of the stochastic knapsack problem, which is a special case of our problem. We circumvent this negative result by showing that $O(\log nR)$ resource augmentation suffices to obtain $O(1)$ approximation on trees and $O(\log nR)$ approximation on general graphs. To achieve this result, we reduce stochastic graph exploration to a memoryless process—the *minesweeper* problem—which assigns to every edge a probability that the process terminates when the edge is probed. For this problem, interesting in its own, we present an optimal polynomial time algorithm on trees and an $O(\log nR)$ approximation for general graphs.

We study also the problem in which the maximum cost of an edge is a logarithmic fraction of the budget. We show that under this condition, there exist polynomial-time oblivious strategies that use $1 + \epsilon$ budget, whose adaptivity gaps on trees and general graphs are $1 + \epsilon$ and $8 + \epsilon$, respectively. Finally, we provide additional results on the structure and the complexity of nonadaptive and adaptive strategies.

2012 ACM Subject Classification Theory of computation → Graph algorithms analysis; Theory of computation → Stochastic approximation

Keywords and phrases stochastic optimization, graph exploration, approximation algorithms

Digital Object Identifier 10.4230/LIPIcs.ICALP.2019.131

* Partially done while the authors were visiting the Algorithms and Uncertainty program at the Simons Institute for the Theory of Computing, Berkeley. The authors would like to thank Yossi Azar, Anupam Gupta, Peter Auer and C. Seshadhri for fruitful discussions.

[1] Partially supported by ERC Advanced Grant 788893 AMDROMA "Algorithmic and Mechanism Design Research in Online Markets".

1 Introduction

Exploring networked data is a time consuming and expensive task which is also subject to several limitations. For example, social networks can be explored only through the use of specific APIs made available by the provider which restrict the number of nodes that can be probed and limit the number of neighbors of each node that can be discovered with one probe. The cost and the difficulty of exploring large-scale networks can be an obstacle to collecting suitable snapshots for the purpose of testing new network analysis tools. The testing is more often executed on static networks made available in public repositories [17, 18] collected in the past for other purposes. It is therefore of crucial importance the development of effective and efficient methods to explore large-scale networks.

The core of a network exploration method is the definition of a probing strategy that decides which nodes or edges to probe at each stage of the process. Both the edge-probe and the node-probe models are useful in this setting. In the case of the exploration of social networks, a node-probing strategy allows to gain knowledge on a subset of the neighbors of the probed node. In the case of the exploration of the Twitter graph, an edge-probing strategy allows to gain information on those tweets of a user that are retweeted from his followers.

One main difficulty in the definition of an effective probing strategy is the intrinsic uncertain nature in terms of cost and probability of success of the process of discovering links in a network, especially if these links represent complex relationships between nodes. In order to confirm the existence of a link between two nodes, it may be required to execute several experiments whose outcome cannot be predicted in advance. Examples are the in-vitro reactions between proteins needed to discover protein-to-protein interaction networks [6, 22] or the influence between humans in social networks.

The second main difficulty stems from the adaptive nature of the optimal probing strategy that needs to be updated from time to time while new parts of the network are discovered. Adaptive strategies are computationally expensive, given that they must be continuously updated. In the case of large network exploration, the communication cost of adaptive strategies is also high since many machines are usually working in parallel at the exploration process, and the updated strategy must be communicated to the machines participating in the process. We are therefore interested in devising *nonadaptive probing strategies* that are simple and that define the sequence of probes in advance before the process is started. The obvious drawback is that nonadaptive probing strategies may be suboptimal.

Several recent works [16, 20, 21] have focused on the task of exploring real-world networks when a limited budget is available. However, these papers do not provide a comprehensive theoretical study of these problems. In this work we initiated the study of exploring an undirected network from a root node. The graph has costs on the edges and rewards on the vertices. A budget limits the total cost of the of the graph edges that are probed.

More formally, the input of the *stochastic graph exploration problem* is an undirected graph $G = (V, E)$ with a source vertex $s \in V$, *stochastic* edge costs $C : E \to \mathbb{R}_{\geq 0}$ distributed according to π_e, $e \in E$, and deterministic rewards of vertices $w : V \to \mathbb{R}_{\geq 0}$. (The model can be easily extended to rewards distributed according to independent random variables.) During the graph-exploration process we construct a set of edges $F \subseteq E$ that we probe and we traverse. All vertices of the subgraph of G spanned by F must be connected to s via the edges of F. We probe edges one by one and we add them to F. The actual cost of an edge e, drawn from the distribution π_e, is revealed only when the edge is traversed. The goal is to maximize the total reward from the vertices spanned by the edge set F while the total cost of the edges in F remains bounded by a prespecified budget B. As soon as we probe an edge such that the total cost exceeds B the process terminates.

In the stochastic graph exploration problem, we aim to design simple polynomial-time computable *nonadaptive strategies* with a reward as close as possible to the reward obtained by the optimal *adaptive strategy*, which decides on the next edge to be traversed after the cost of all previously traversed edges is revealed (see Section 2 for precise definitions). This is customary in a class of stochastic optimization problems [4], for which it is common to bound the *adaptivity gap* of the nonadaptive strategy.

The stochastic graph exploration problem generalizes some important stochastic optimization problems. If the graph G is a star graph, our problem models *exactly* the stochastic knapsack problem [4, 8]. Stochastic knapsack admits an $O(1)$ adaptivity gap, that is, there exists an optimal nonadaptive strategy, which approximates the optimal adaptive strategy up to a constant factor. The nonadaptive strategy is devised by exploiting a suitable LP relaxation for the problem because the standard formulation has an unbounded integrality gap defined as the worst-case ratio between the optimal integral cost and optimal fractional cost of the LP. In the LP version of the problem that is used, the costs of the edges are reduced to their truncated (by the maximum budget) expected costs and the rewards are also reduced by the probability that the cost of the item is below the maximum budget.

If the network we need to explore is a tree, the stochastic graph problem is a stochastic knapsack problem with precedence constraints: only a subset of items are available in the beginning and adding each item to the knapsack will make some new items—the direct descendants of the explored node—available. Unfortunately, as opposed to the knapsack problem, the adaptivity gap of the stochastic graph exploration problem that we consider is unbounded even on a tree network and therefore the LP-based approach of stochastic knapsack cannot directly be extended.

The stochastic graph exploration problem also models stochastic graph probing problems. Probing problems in graphs have been introduced [7, 14] because of their applications to kidney exchange and online dating. Consider a probing probability for each edge $p : E \to [0, 1]$, that is, edge e will materialize with probability $p(e)$ each time is probed, independently of the other edges and of the previous probes. The goal is to maximize the number of vertices that are connected to a source vertex s by the set F of edges that have been successfully probed when the total number of probes is limited by B. Nonadaptive strategies probe a list of edges in a sequence till success or the total budget B is reached. The stochastic graph exploration problem we study models the stochastic graph probing problem by setting the costs of the edges distributed according to $\mathbf{Pr}(C_e = i) = (1 - p(e))^{i-1}p(e)$, with i being the number of probes needed to discover edge e.

1.1 Summary of Our Results

Our main contribution is the definition of the stochastic graph exploration problem and the study of the adaptivity gap of nonadaptive probing strategies. Here is a summary of our results:

Our first result is an $\Omega(n)$ adaptivity gap for the stochastic graph exploration problem even on a spider graph, which is a tree containing a single node of degree more than two. (Observe that the problem for a simple path is easy because the optimal strategy will traverse sequentially the edges of the path starting from the root.)

One first direction we pursue to circumvent the impossibility result is to allow a limited amount of resource augmentation: instead of using budget B, we allow the algorithm to use a budget of $\beta \cdot B$, for some value of β. We call an algorithm (α, β)-approximate if it computes a strategy which uses budget $\beta \cdot B$, and obtains an expected reward of at least $1/\alpha$ times the optimal reward (obtained by an adaptive algorithm). We present polynomial time computable nonadaptive strategies in a graph of n vertices that are $(O(1), O(\log nR))$-approximate for trees and $(O(\log nR), O(\log nR))$-approximate for general graphs, with R being the maximum reward of a vertex.

The idea is to transform the stochastic exploration problem into a memoryless stochastic process, which we call the *minesweeper problem*, and which may be of independent interest. In the minesweeper problem, the budget and the edge costs are replaced by probabilities $p(e)$, which are specified for every edge e. When an edge e is probed, the process stops with probability $1 - p(e)$. Hence, the final reward of a vertex is discounted by the probability that the strategy does not stop before the vertex is acquired. The minesweeper problem is, in fact, a special case of stochastic graph exploration, where the support of each π_e (distribution of cost of edge e) is $\{0, B + 1\}$ and the budget is B.

We prove that an α-approximate strategy for the minesweeper problem implies an $(O(\alpha), O(\log nR))$-approximate nonadaptive strategy for the stochastic graph exploration problem. The idea of the reduction is as follows. We construct a minesweeper problem instance, where $p(e) = \mathbf{Pr}(\pi_e < X_B)$, where X_B is random variable that follows an exponential distribution with parameter B. We first show that, for any subset of edges F, the probability that their total cost in the stochastic graph exploration is at most B is at most a constant factor of the probability that minesweeper would stop on this set. On the other hand, the expected additional reward that can be achieved from minesweeper after the total cost becomes larger than $O(B \log nR)$ is negligible.

We then show how to compute in polynomial time an optimal strategy for the minesweeper problem on trees and an $O(\log nR)$-approximate strategy on general graphs. These results imply imply an $(O(1), O(\log nR))$-approximate strategy for trees and an $(O(\log nR), O(\log nR))$-approximate strategy for general graphs. To show the optimal result on trees we prove two facts. First, the order of traversal of the edges in each subtree can be determined independently. Second, we show a simple optimality condition which helps us determine how many edges from each subtree should be probed before switching to a different subtree. We remark that our approach is in a spirit similar to the greedy optimal strategy defined by the Gittins index [9, 10] for multi-armed bandit problems. However, differently from the standard setting of the Gittins index, in the minesweeper problem, a whole new set of arms is made available for each node of the tree reached by the exploration process. Moreover, in the minesweeper problem, the discount factor is not constant because it depends on the probability assigned to the traversed edge. This approach is not viable for general graphs, and we provide an approximate solution instead, by showing a reduction of minesweeper to max-prize problem [5].

We also pursue a second direction to circumvent the lower bound on the adaptivity gap for trees: we restrict the distributions by considering the case when the edge costs are bounded by $\frac{\epsilon^2 B}{c \log n}$ for a suitable constant c. We show, under this condition, the existence of a polynomial time computable $(1 + \epsilon, 1 + \epsilon)$-approximate nonadaptive strategy for trees and $(1 + \epsilon, 8 + \epsilon)$-approximate nonadaptive strategy for any graph G. We note that this approach can be extended to prove a result with resource augmentation similar to the one we obtained through reduction to the minesweeper problem. Yet, we believe that both the minesweeper problem and the reduction technique can be of independent interest.

Our final result is related to the problem of finding a nonadaptive probing strategy that is $(o(n), O(1))$-approximate. We leave open this challenging problem even for trees. However, we establish an interesting result for the characterization of nonadaptive strategies. We prove that any nonadaptive strategy that probes edges in order until it succeeds or until the budget is exceeded can be $(O(1), O(1))$-approximated by a set strategy, which probes all edges at once and obtains a reward only if all edges of a set are successfully probed within budget. We specifically prove that the adaptivity gap of a nonadaptive strategy can be approximated up to a factor of 6 by a set strategy that uses budget $9B$. We use this result to give an algorithm for finding a strategy for trees, which is $(O(1), O(1))$-approximate, compared to the best *nonadaptive* strategy. Surprisingly, the resulting strategy is adaptive.

1.2 Related Work

The adaptivity gap of stochastic problems has been studied for the knapsack problem [4, 8] which is a special case of the problem we study. The adaptivity gap has also been studied for budgeted multi-armed bandits [11, 12, 19] by resorting to suitable linear programming relaxation. Differently from previous work on budgeted multi-armed bandit problems, we consider the setting in which new arms appear after some arms are pulled. Stochastic probing problems have also been studied for matching [1, 2, 7] motivated from kidney exchange and for more general classes of matroid optimization problems [14, 15].

The stochastic graph exploration problem we introduce is also related to the *stochastic orienteering* problem [3, 13]. In stochastic orienteering, the set of traversed edges must form a path in a metric graph with deterministic costs on the edges, while the time spent on a node is a random variable, which follows an a-priori known distribution. In stochastic graph exploration, the random variables are the costs of the edges of the graph but we cannot ensure that the costs on the edges form a metric since the random variables are independent.

1.3 Organization of the Paper

In Section 2 we formally define our problems. In Section 3 we show the lower bounds on the adaptivity gap for stochastic graph exploration. In Section 4 we show our reduction to the minesweeper problem and our results for stochastic graph exploration with resource augmentation. In Section 5 we present a near-optimal set strategy for trees. In Section 6 we present our results for the case of edges of small costs and, finally, in Section 7 we study the power of resource augmentation for relating the cost of nonadaptive strategies to the cost of optimal set strategies.

2 Problem Definition

We start by an auxiliary definition. Let $G = (V, E)$, with $|V| = n$, be an undirected graph and $s \in V$. We say that a set $F \subseteq E$ is *connected to* s if F induces a connected subgraph of G and s is the endpoint of at least one $e \in F$.

Let us now define the StochasticExploration problem (in the following sometimes abbreviated by SGE). This problem instance is given by a tuple (G, s, C, w), where G is an undirected graph $G = (V, E)$, $s \in V$ is a source vertex, C is a function that assigns *stochastic* edge costs to each edge, and $w : V \to \mathbb{R}_{\geq 0}$ is a function that assigns (deterministic) reward to each vertex.[2] And we denote R as the maximum reward of a vertex i.e. $R = \max_{v \in V} w(v)$. Formally, for each $e \in E$, $C(e)$ is a random variable distributed according to π_e that takes

[2] The results hold also if the rewards are random variables that are independent of each other and the edge costs. It suffices to replace each reward with its expected value.

values in $\mathbb{R}_{\geq 0}$, all random variables $C(e)$ being jointly independent. For an edge (u, v) we will often denote $C(u, v) = C((u, v))$.

Consider the following single-player game. The player has an initial budget of B ($B = 1$ if not specified) and maintains an initially empty subset F of E, which we call the set of *acquired edges*. In each step the player can choose an edge $e \in E \setminus F$ and *probe* it (if $F = E$, the game finishes). Probing an edge e is only allowed when $F \cup \{e\}$ is connected to s. When e is probed, the actual cost $C(e)$ of e, drawn from the distribution π_e, is revealed. If the cost e is not greater than the remaining budget, e is *acquired* (added to F) and $C(e)$ is subtracted from the budget. If $C(e)$ exceeds the remaining budget, the game finishes. The goal of the player is to maximize the final *payoff* of F, which is the total reward of all vertices in the subgraph of G induced by F.

Let us now define the MineSweeper problem, which we often abbreviate to MS. This problem is defined by a tuple (G, s, p, w), where G is an undirected graph, $s \in V$ is a start vertex, $p : E \to [0, 1]$ is a function that assigns to each edge e the probability that e materializes and $w : V \to \mathbb{R}_{\geq 0}$ is a function that assigns (deterministic) reward to each vertex. The only difference between MS and SGE is in how edges are probed. There are no edge costs or budget. Instead, whenever an edge e is probed, it materializes (independently of the other edges) with probability $p(e)$ and is acquired immediately. If the edge does not materialize, the process ends immediately. Note that as in SGE, probing an edge e is only allowed when $F \cup \{e\}$ is connected to s. Note that the MineSweeper problem is a special case of the StochasticExploration problem, by letting, for each edge e, π_e be the distribution in which with probability $p(e)$ we obtain the value 0 and with probability $1 - p(e)$ the value $B + 1$.

We consider the following types of strategies for both problems:

- An *adaptive* strategy is a mapping from the set of already acquired edges (and the remaining budget, in the case of SGE) to the next edge to be probed.
- A *nonadaptive* strategy, also called a *list* strategy, is described by a sequence $e_1, \ldots, e_k$ consisting of distinct elements of E, such that for each $1 \leq i \leq k$, the set $\{e_1, \ldots, e_i\}$ is connected to s. In this strategy, the edges are simply probed according to their order in the sequence.
- A *set* strategy is a nonadaptive strategy with the additional restriction that it does not obtain any payoff if it does not acquire all edges from the list.[3]

For a strategy S for SGE, we denote by $r(\mathcal{I}_{\text{SGE}}, S, B)$ the expected payoff of strategy S for the SGE problem instance $\mathcal{I}_{\text{SGE}} = (G, s, C, w)$ with initial budget of B, which is the expected reward of the set of nodes in the returned solution. When $B = 1$ we sometimes omit the third argument of $r(\cdot)$. Similarly, we denote by $r_{\text{MS}}(\mathcal{I}_{\text{MS}}, S)$ the expected payoff of strategy S for the MS problem instance $\mathcal{I}_{\text{MS}}$. We call a strategy S *optimal* for $\mathcal{I}$ with budget B, if for all strategies S', $r(\mathcal{I}, S, B) \geq r(\mathcal{I}, S', B)$. Let OPT_{ad} be the optimal adaptive strategy for the SGE problem and OPT_{na} be the optimal nonadaptive strategy. We call a strategy S α-approximate, if for each instance $\mathcal{I}$, $r(\mathcal{I}, S) \geq 1/\alpha \cdot r(\mathcal{I}, \text{OPT}_{\text{ad}})$. Finally, an algorithm ALG is (α, β)-approximate if for any instance $\mathcal{I}$ it computes a α-approximate strategy by using a β factor resource augmentation, i.e. $r(\mathcal{I}, \text{ALG}(\mathcal{I}), \beta \cdot B) \geq 1/\alpha \cdot r(\mathcal{I}, \text{OPT}_{\text{ad}}, B)$.

[3] Note that we abuse earlier definitions slightly for the sake of simplicity.

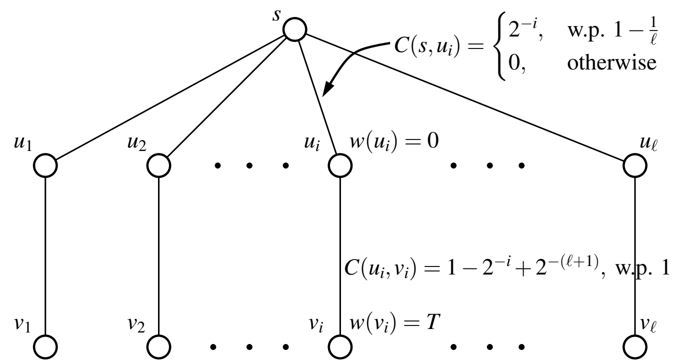

Figure 1 An instance in which the optimal adaptive strategy obtains a payoff which is $\Omega(n)$ larger than the payoff of the optimal nonadaptive strategy.

3 Lower Bounds

In this section we prove a lower bound on the *adaptivity gap* of StochasticExploration. Namely, we show an instance $\mathcal{I}_{LB} = (G, s, C, w)$ such that $r(\mathcal{I}_{LB}, \mathrm{OPT}_{\mathrm{ad}})/r(\mathcal{I}_{LB}, \mathrm{OPT}_{\mathrm{na}}) = \Omega(n)$, where $\mathrm{OPT}_{\mathrm{ad}}$ and $\mathrm{OPT}_{\mathrm{na}}$ denote the optimal adaptive and nonadaptive strategies.

The instance $\mathcal{I}_{LB}$ is shown in Figure 1. The graph G contains the set of nodes $\{s, u_1, u_2, \ldots, u_\ell, v_1, \ldots, v_\ell\}$, and the set of edges (s, u_i) and (u_i, v_i) for each $i \in [\ell]$. For each $i \in [\ell]$ we set $w(u_i) = 0$, $w(v_i) = T$, $C(s, u_i) = 2^{-i}$ with probability $1 - 1/l$ and 0 otherwise, and $C(u_i, v_i) = 1 - 2^{-i} + 2^{-(\ell+1)}$ with probability 1.

▶ **Lemma 1.** *Let OPT_{ad} and OPT_{na} denote the optimal adaptive and nonadaptive strategies for instance $\mathcal{I}_{LB}$. Then, $r(\mathcal{I}_{LB}, OPT_{ad})/r(\mathcal{I}_{LB}, OPT_{na}) = \Omega(n)$.*

One natural approach for StochasticExploration instance is to replace the stochastic edge costs with the truncated expected costs, that is, set the cost of an edge e to $\mathbb{E}[\min\{1, C(e)\}]$. However as this following example illustrates this approach does not lead to a good solution, even if constant budget augmentation is allowed.

▶ **Lemma 2.** *Let OPT_{ad} denote the optimal adaptive strategy for an instance $\mathcal{I}$ and let n be the number of vertices in the instance. Let OPT_{na} be the optimal nonadaptive strategy computed on instance $\mathcal{I}_{TR}$ obtained from $\mathcal{I}$ by setting edge costs $\mathbb{E}[\min\{1, C(e)\}]$, $e \in E$. Assume the nonadaptive algorithm is allowed to use a budget of $1 < c < n/10$. Then, there exists an instance $\mathcal{I}$ such that $r(\mathcal{I}, OPT_{ad})/r(\mathcal{I}_{TR}, OPT_{na}) = \Omega(n/2^{2c})$.*

4 The General Case and the Minesweeper Problem

In this section we describe algorithms for solving StochasticExploration, which use logarithmic budget augmentation. We first show how to reduce an instance of SGE to MineSweeper and then present solutions for MineSweeper on trees and general graphs. During the description of the reduction we also introduce the logarithmic budget augmentation. First, we observe that in the MineSweeper problem we do not have budget so there is no history that an algorithm may have to remember, except for the edges that it has probed (and succeeded). This implies the following:

▷ Observation 1. There exists an optimal strategy for the MineSweeper problem that is nonadaptive.

4.1 Reduction from StochasticExploration to MineSweeper

In this section we show how, given an instance $\mathcal{I}_{\text{SGE}} = (G, s, C, w)$ of StochasticExploration, we transform it to an instance $\mathcal{I}_{\text{MS}} = (G, s, p, w)$ of MineSweeper. The graph and the rewards remain the same; the challenge is to define the correct edge probability function $p(\cdot)$ for MS and relate it to the cost function $C(\cdot)$ of SGE. For each edge e' we transform the cost distribution $C(e')$ to the probability that the edge materializes, $p(e')$ (a scalar). Let $X_{e'}$ be a random variable distributed according to the exponential distribution with parameter 1, let $c_{e'}$ be the cost, which is distributed according to $C(e')$, and we set $p(e') = \mathbf{Pr}(X_{e'} > c_{e'})$. Next we show how this choice couples the two problems.

First, we show that for any subset of edges F the probability that their total cost in SGE is at most 1 is at most a factor e times of the probability that all the edges in F materialize, and therefore MS does not stop on this set. Let $\mathcal{E}_F$ be the event that all the edges in F materialize and $\mathcal{G}_F$ the event that $\sum_{e' \in F} c_{e'} \leq 1$. The following lemma makes use of properties of the exponential distribution.

▶ **Lemma 3.** *For any $F \subseteq E$ we have that $\mathbf{Pr}(\mathcal{G}_F) \leq e \cdot \mathbf{Pr}(\mathcal{E}_F)$.*

This lemma allows us to prove the following lemma, which gives a strategy for MS that is competitive with the optimal adaptive strategy for SGE. The idea behind the proof is to define a strategy for MS in such a way that we can couple the execution of the two strategies in the corresponding problems.

▶ **Lemma 4.** *Consider an* SGE *instance $\mathcal{I}_{\text{SGE}} = (G, s, C, w)$ and let $\mathcal{I}_{\text{MS}} = (G, s, p, w)$ be an instance for* MS *as defined previously. Let OPT_{ad} denote the optimal adaptive strategy for* SGE *and OPT_{MS} the optimal strategy for* MS*. We have that*

$$r((G, s, C, w), OPT_{ad}, 1) \leq e \cdot r_{MS}((G, s, C, w), OPT_{MS}).$$

Recall from Observation 1 that the optimal strategy for the MineSweeper problem is nonadaptive, therefore it can be specified by a list of edges that are selected sequentially until for one of them there is a failure. Let OPT_{MS} be such an optimal sequence of edges. Next we show that the sequence of edges OPT_{MS} can provide an approximate result to the StochasticExploration problem if we allow for some budget augmentation.

▶ **Lemma 5.** *Consider an* SGE *instance $\mathcal{I}_{\text{SGE}} = (G, s, C, w)$ and let $\mathcal{I}_{\text{MS}} = (G, s, p, w)$ be an instance for* MS *as defined previously. Let OPT_{MS} be the optimal sequence of edges for the* MineSweeper *instance, and let S be the (nonadaptive) strategy for* StochasticExploration *that probes the same edges, in the same order. Then we have that*

$$r((G, s, C, w), S, 2\ln(nR)) \geq r_{MS}((G, s, C, w), OPT_{MS}) - o(1),$$

where $R = \max_{v \in V} w(v)$.

Collecting the results of Lemmas 4 and 5 we obtain the following theorem.

▶ **Theorem 6.** *Consider an* SGE *instance $\mathcal{I}_{\text{SGE}} = (G, s, C, w)$ and let $\mathcal{I}_{\text{MS}} = (G, s, p, w)$ be an instance for* MS *as defined previously. Then*

$$r((G, s, C, w), OPT_{na}, 2\ln(nR)) + o(1) \geq r_{MS}((G, s, C, w), OPT_{MS}) \geq \frac{r((G, s, C, w), 1, OPT_{ad})}{e}.$$

4.2 MineSweeper on Trees

We show that the minesweeper problem on trees can be solved optimally in near-linear time.

▶ **Theorem 7.** *Consider the instance $\mathcal{I} = (T, s, p, w)$ of the minesweeper problem, where T is a tree. The optimal strategy, OPT_{MS}, for* MineSweeper *on T can be computed in $O(n \log n)$ time, where n is the number of vertices of T.*

The algorithm is surprisingly simple and based on a greedy approach. We define the utility of an edge to be the expected payoff from probing it, divided by the probability that the edge does not materialize. The algorithm is based on two observations. First, we observe that if there is a node x in the graph with a single child y and the utility of the edge xy is larger than the utility of the edge connecting x and its parent, then without loss of optimality we can assume that the edge xy is probed right after the edge connecting x and its parent, so we can merge these two edges into a single one. Second, if there is a node x, such that one can probe all edges in the subtree of x in the order of decreasing utilities (and not violate the constraint that an edge can be probed only after one of its endpoints has been acquired) then one can replace the entire subtree of x with a line, which is a subtree imposing the concrete order of probing edges. It turns out that by using both these rules one can find the optimal order of probing edges efficiently.

We obtain the algorithm by generalizing some existing results from the area of scheduling. At the same time our analysis is arguably simpler. We give the proof of Theorem 7 in the full version of the paper.

4.3 MineSweeper on general graphs

In this section we present an algorithmic solution to Minesweeper for general graphs, which provides a bicriteria approximation for our problem. We prove the following theorem.

▶ **Theorem 8.** *Consider the instance $\mathcal{I} = (G, s, p, w)$ of the minesweeper problem, where $G = (V, E)$ is an undirected graph. An $O(\log nR)$-approximate strategy can be computed in polynomial time.*

In the following we provide a sketch of the proof. Assume that the optimal solution is the sequence of edges $S^* = (e_1, \ldots, e_k)$. We first observe that the edges in S^* must form a tree. Define $\mathcal{M}(E')$ to be the event that all the edges in the set E' materialize. Also let $w(e_1, \ldots, e_i) = \sum_{j=1}^{i} w(e_i)$. Then S^* is a sequence that maximizes

$$O^* = \sum_{i=1}^{k} \mathbf{Pr}(\mathcal{M}(\{e_1, \ldots, e_i\}), \neg\mathcal{M}(\{e_{i+1}\})) \cdot w(e_1, \ldots, e_i)\,.$$

For $\ell = 0, 1, \ldots, \ln nR$, define $I(\ell)$ to be all values j such that $w(e_1, \ldots, e_j) \in [2^\ell, 2^{\ell+1} - 1]$, and $\iota(\ell)$ to be the smallest such j.

We can write after some manipulations:

$$O^* \leq \sum_{\ell=0}^{\ln nR} 2w(e_1, \ldots, e_{\iota(\ell)}) \cdot \mathbf{Pr}\big(\mathcal{M}(\{e_1, \ldots, e_{\iota(\ell)}\})\big) \leq 2\ln(nR) \cdot w(\tilde{E}) \cdot \mathbf{Pr}\big(\mathcal{M}(\tilde{E})\big)\,,$$

with $\tilde{E} \subset E$ being the set of edges that defines a tree that contains s and maximizes $w(\tilde{E}) \cdot \mathbf{Pr}\big(\mathcal{M}(\tilde{E})\big)$. Therefore, our goal becomes that of finding that set of edges $\tilde{E}$ that forms a tree and maximizes $w(\tilde{E}) \cdot \mathbf{Pr}\big(\mathcal{M}(\tilde{E})\big)$.

For this purpose, we use the problem of *max-prize tree*. In the max-prize tree [5] we are given an undirected graph $G = (V, E)$ with a source vertex $s \in V$, (deterministic) edge costs $c : E \to \mathbb{R}_{\geq 0}$, deterministic rewards on the vertices $w : V \to \mathbb{R}_{\geq 0}$, and a budget $B \in \mathbb{R}$. The objective is to build a subgraph $G' = (V', E')$ of G such that (1) G' is a tree, (2) $s \in V'$, and (3) $\sum_{e \in E'} c(e) \leq B$, that maximizes $\sum_{v \in V'} w(v)$.

We use for our approximation the 8-approximation algorithm for the max-prize–tree problem given by Blum et al. [5].

5 Approximating Set Strategy on Trees

In this section we show an algorithm for computing a strategy for trees, which is $(1, 1+\epsilon)$-approximate compared to the optimal set strategy. The strategy itself is adaptive.

▶ **Lemma 9.** *Let $\mathcal{I} = (G, s, C, w)$ be a* SGE *instance, where G is a tree. Let OPT_{set} be the optimal set strategy for $\mathcal{I}$. Then, in $O(n^4/\epsilon^2)$ time we can compute an adaptive strategy S, such that $r(\mathcal{I}, S, 1+\epsilon) \geq r(\mathcal{I}, OPT_{set}, 1)$. Moreover, if edge costs are not stochastic, that is, the support of each distribution π_e has size 1, the algorithm runs in $O(n^3/\epsilon)$ time and the resulting strategy is not adaptive.*

We briefly describe the ideas behind the algorithm. Consider the instance $\mathcal{I} = (T, s, C, w)$, where T is a tree. We root the tree at s and assume an order on the children of each node. Consider the sequence $P = \langle e_1, \ldots, e_n \rangle$ of the tree edges built with the following recursive algorithm. Given a node of T, iterate through its descendant edges (according to their order) and for each such edge output it and recur on the other endpoint. This traverses the tree in a preorder fashion. We define $\prec$ to be the linear order on the edges of T induced by this traversal. In the following, we assume that the edges are ordered according to $\prec$, for example, by a maximal element of a set of edges, we mean the edges that is largest according to $\prec$.

We say that a subset A of edges of T is *feasible*, if each edge $e \in A$ is either incident to the root of T, or the parent edge of e also belongs to A. Observe that given sufficient budget, a strategy can acquire any feasible set of edges of T. This follows from the fact that for each edge e of T, its parent comes before it in P. Our algorithm will probe some feasible set of edges according to the order $\prec$, that is, after probing an edge e it will not probe any edge f such that $f \prec e$.

The algorithm for computing our strategy is based on dynamic programming. A simple and inefficient approach is to use an exponential number of states. Namely, each state can be characterized by the set of edges acquired so far, denoted by A, and the remaining budget, which we discretize to a multiple of ϵ/n. Knowing the set A allows us to find all such edges e that $A \cup \{e\}$ is a feasible set and e comes after the maximal element of A in the order $\prec$. The key idea is that we can improve the number of states to polynomial, by taking advantage of the following property of the ordering $\prec$.

▶ **Lemma 10.** *Let A be a nonempty feasible set of edges of T and let e be the maximal edge of A. Given e (and without knowing A) we can compute the set F_e of all edges f such that $e \prec f$ and $A \cup \{f\}$ is a feasible set.*

6 Bounded Edge Costs

In this section, we deal with the special case of StochasticExploration, where the cost of each edge is bounded by $O(\frac{\epsilon^2}{\ln n})$ and the ratio between the smallest and largest reward R

is polynomial in n. We prove that in this setting a $(O(1), 1+\epsilon)$ strategy for SGE can be computed in polynomial time.

▶ **Theorem 11.** *Let $\mathcal{I} = (G, s, C, w)$ be an instance of* SGE*, where $C(e) = O(\frac{\epsilon^2}{\ln n})$ (for each edge e and some $0 < \epsilon = O(1)$), $R \leq \epsilon n^{O(1)}$, and the smallest reward is 1. Then, in polynomial time, we can compute a nonadaptive $(O(1), 1+\epsilon)$-approximate strategy for $\mathcal{I}$. Additionally, if G is a tree, then in time $O(n^3/\epsilon)$ we can compute a nonadaptive $(1+\epsilon, 1+\epsilon)$-approximate strategy for $\mathcal{I}$.*

To prove the theorem, we consider the following strategy. We replace the stochastic edge costs with their expected values (i.e., the edge cost distributions in the modified instance have size 1). Then, we show that the optimal set strategy using budget augmented by a factor of $1+\epsilon$ gives a $(1+\epsilon)$-approximate solution.

For ease of notation, we scale the edge costs and the budgets by a factor of $\Theta(\epsilon^2/\ln n)$, so that the edge costs are bounded by 1 and the available budget is $B = O(\epsilon^2/\ln n)$.

First, we bound the payoff of an adaptive strategy when the expected cost of its acquired edges is more than $B \cdot (1+\epsilon)$. Let $\mu_e = \mathbf{E}[C(e)]$, and $\mu(F) = \sum_{e \in F} \mu_e$.

▶ **Lemma 12.** *Let $0 < \epsilon < 1/3$ and let $\mathcal{I} = (G, s, C, w)$ be an instance of* SGE*, in which $B \geq 5c/\epsilon^2 \cdot \ln n$. Let F be a set of edges acquired by some adaptive strategy. If $\mu(F) \geq (1+\epsilon) \cdot B$ then the probability that $C(F) \leq B$ is at most n^{-c}.*

Next, we show that if the expected cost of some set of edges is close to the budget, then this cost is highly concentrated around the expected value. This enables us to give a set strategy with small budget augmentation.

▶ **Lemma 13.** *Let $\mathcal{I} = (G, s, C, w)$ be an instance of* SGE*. For any set of edges F and any $\tilde{B} \geq 5c/\epsilon^2 \cdot \ln n$, if $\mu(F) = \tilde{B}$ then the probability that $C(F) \geq (1+\epsilon)\tilde{B}$ is at most n^{-c}.*

▶ **Lemma 14.** *Let $\mathcal{I} = (G, s, C, w)$ be an instance of* SGE*, where $B \geq 5c/\epsilon^2 \ln n$, the maximum reward R satisfies $R \leq \epsilon n^{c-1}$, and the minimum reward is 1. Let $\mathcal{I}_e$ be obtained from $\mathcal{I}$ by replacing each edge cost with its expected value. Let OPT^{ϵ}_{set} be the optimal set strategy using budget $(1+\epsilon)B$ for $\mathcal{I}_e$ and OPT_{ad} be the optimal adaptive strategy using budget B for $\mathcal{I}$. Then, $(1+\epsilon)r(\mathcal{I}, OPT^{\epsilon}_{set}, (1+\epsilon)B) \geq r(\mathcal{I}, OPT_{ad}, B)$.*

Observe that finding the optimal set strategy on $\mathcal{I}_e$ is NP-hard, as it generalizes the knapsack problem. However, it becomes tractable, if we augment the budget. In particular, for trees, we use the algorithm of Lemma 9, and for general graphs, in Section 4.3, we show how to use the solution of the max-prize problem.

7 Nonadaptive strategies

In this section we consider nonadaptive strategies for the stochastic exploration problem. The main result of this section is that, for the graph exploration problem, that there exists a *set-strategy* with a constant budget augmentation, which is a constant competitive compared to the best nonadaptive algorithm. Recall that, a *set-strategy* is to choose a set of edges (without an internal order) and to try to probe all of the edge in that set. The gain of strategy for a set of edges, is nonzero only if the *entire set* was successfully probed (i.e., if the total cost of the set is smaller than the budget), and then it collects the rewards of all the vertices connected to this set. Therefore, the expected gain of *set-strategy* given a set of edges, is the total gain of vertices spanned by these edges times the probability that the total cost of these edges would not be greater than the specified budget.

First, we are able to show how much is the increment in the probability to successfully probe a set, when using a constant budget augmentation.

7.1 Power of Budget Augmentation

Let $S = \{e_1, e_2, \ldots, e_n\}$ be a set of edges and let $c_i \triangleq C(e_i)$. Define $C_k^n = \sum_{i=k}^n c_i$ the realized cost of the subset of the edges $\{e_k, \ldots e_n\}$ and, for ease of notation, let $C^j = C_1^j$. For any $i \in [n]$ let $P_i(a)$ be the probability that the sum of cost of the edges $\{e_1, \ldots e_i\}$ is at most a, that is, $P_i(a) = \mathbf{Pr}\big(C^i \le a\big)$.

The next lemma will allow us to take advantage of budget augmentation.

▶ **Lemma 15.** *Assume that for each edge e_i, $i \in [n]$ we have $c_i \in [0, 1]$. Then*

$$P_n(3) \ge P_n(1)\,(1 - \ln(P_n(1)))\,.$$

Interestingly, the multiplicative factor increases as the probability to succeed with the original budget decreases. We will use this fact, but to compare to a list-strategy we need stronger guarantees, we simply use the above lemma twice and deduce the following.

▶ **Corollary 16.**

$$P_n(9) \ge P_n(1)\frac{(1 - \ln(P_n(1)))^2}{2}$$

7.2 List Strategy vs. Set Strategy

Now, we are ready to prove the main claim of this section, that we are able to compare the strategies using a budget augmentation. Consider an SGE problem instance $\mathcal{I} = (G, s, C, w)$. Let $\mathrm{S}_{ls} = \langle e_1, \ldots, e_n \rangle$ be a nonadaptive strategy (a feasible sequence of edges) and let v_i denote the vertex whose reward is obtained when e_i is acquired. The expected payoff of probing the list with budget $B(\ge 1)$ is by linearity of expectation:

$$r(\mathcal{I}, \mathrm{S}_{ls}, B) = \sum_{j=1}^{n} w(v_j) \cdot \mathbf{Pr}\big(C^j \le B\big)\,.$$

Given a nonadaptive strategy $\mathrm{S}_{ls} = \langle e_1, \ldots, e_n \rangle$, consider n different set strategies S_k, for $k = \{1 \ldots n\}$, where $S_k = \{e_1, \ldots e_k\}$. Note that the expected payoff of S_k with budget $9 \cdot B$ is

$$r(\mathcal{I}, S_k, 9B) = \mathbf{Pr}\big(C^k \le 9B\big) \cdot \sum_{j=1}^{k} w(v_j).$$

Finally, we show that there exists $k \in \{1, \ldots, n\}$ such that the set strategy S_k with budget $9B$ obtains a constant fraction of strategy S_{ls}.

▶ **Lemma 17.**

$$\max_k \{r(\mathcal{I}, S_k, 9B)\} \ge 0.46 \cdot r(\mathcal{I}, S_{ls}, B).$$

7.3 Algorithm for Trees

By combining Lemma 17 with the algorithm of Lemma 9, we obtain the following.

▶ **Theorem 18.** *Let $\mathcal{I} = (G, s, C, w)$ be a* SGE *instance, where G is a tree. Let OPT_{na} be the optimal nonadaptive strategy for $\mathcal{I}$. Then, in $O(n^4/\epsilon^2)$ time we can compute an adaptive strategy S, such that $r(\mathcal{I}, S, 9 + \epsilon) \geq 0.46 \cdot r(\mathcal{I}, OPT_{na}, 1)$.*

8 Conclusions

In this work we have introduced the stochastic exploration problem on graphs which generalizes the stochastic knapsack problem [4, 8]. We proved that, differently from stochastic knapsack, no $o(n)$ adaptivity gap is possible unless we allow some resource augmentation on the budget. We provided algorithms with bounded adaptivity gap and logarithmic resource augmentation by reducing stochastic exploration to a related memoryless problem—the minesweeper problem. We also considered the case of edges with small costs for which it is possible to provide an algorithm with $O(1)$ adaptivity gap and $O(1)$ resource augmentation. The most challenging problem left open from our work is the one of devising an algorithm with $O(1)$ approximation factor that uses only $O(1)$ resource augmentation for general graphs. The problem is open even for trees. We provided a set of additional results on the structure of optimal adaptive strategies and on the power of resource augmentation for set strategies with respect to list strategies that can help in addressing this problem.

References

1 Marek Adamczyk. Improved analysis of the greedy algorithm for stochastic matching. *Information Processing Letters*, 111(15):731 – 737, 2011. URL: `http://www.sciencedirect.com/science/article/pii/S002001901100127X`, `doi:http://dx.doi.org/10.1016/j.ipl.2011.05.007`.

2 Nikhil Bansal, Anupam Gupta, Jian Li, Julián Mestre, Viswanath Nagarajan, and Atri Rudra. When lp is the cure for your matching woes: Improved bounds for stochastic matchings. *Algorithmica*, 63(4):733–762, Aug 2012.

3 Nikhil Bansal and Viswanath Nagarajan. *On the Adaptivity Gap of Stochastic Orienteering*, pages 114–125. Springer International Publishing, Cham, 2014. URL: `http://dx.doi.org/10.1007/978-3-319-07557-0_10`, `doi:10.1007/978-3-319-07557-0_10`.

4 Anand Bhalgat, Ashish Goel, and Sanjeev Khanna. *Improved Approximation Results for Stochastic Knapsack Problems*, pages 1647–1665. URL: `http://epubs.siam.org/doi/abs/10.1137/1.9781611973082.127`, `arXiv:http://epubs.siam.org/doi/pdf/10.1137/1.9781611973082.127`, `doi:10.1137/1.9781611973082.127`.

5 Avrim Blum, Shuchi Chawla, David R. Karger, Terran Lane, Adam Meyerson, and Maria Minkoff. Approximation algorithms for orienteering and discounted-reward tsp. *SIAM Journal on Computing*, 37(2):653–670, 2007.

6 Michael Caldera, Pisanu Buphamalai, Felix Mueller, and Jorg Menche. Interactome-based approaches to human disease. *Current Opinion in Systems Biology*, 3:88 – 94, 2017.

7 Ning Chen, Nicole Immorlica, Anna R. Karlin, Mohammad Mahdian, and Atri Rudra. *Approximating Matches Made in Heaven*, pages 266–278. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. URL: `http://dx.doi.org/10.1007/978-3-642-02927-1_23`, `doi:10.1007/978-3-642-02927-1_23`.

8 Brian C. Dean, Michel X. Goemans, and Jan Vondrak. Approximating the stochastic knapsack problem: The benefit of adaptivity. *Mathematics of Operations Research*, 33(4):945–964, 2008. `doi:10.1287/moor.1080.0330`.

9 Esther Frostig and GIDEON WEISS. Four proofs of gittins multiarmed bandit theorem, 1999.

10 Author(s) J. C. Gittins and J. C. Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society, Series B*, pages 148–177, 1979.

11 Sudipto Guha and Kamesh Munagala. Approximation algorithms for budgeted learning problems. In *Proceedings of the Thirty-ninth Annual ACM Symposium on Theory of Computing*, STOC '07, pages 104–113, New York, NY, USA, 2007. ACM. URL: `http://doi.acm.org/10.1145/1250790.1250807`, `doi:10.1145/1250790.1250807`.

12 Anupam Gupta, Ravishankar Krishnaswamy, Marco Molinaro, and R. Ravi. Approximation algorithms for correlated knapsacks and non-martingale bandits. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 827–836, 2011.

13 Anupam Gupta, Ravishankar Krishnaswamy, Viswanath Nagarajan, and R. Ravi. Running errands in time: Approximation algorithms for stochastic orienteering. *Mathematics of Operations Research*, 40(1):56–79, 2015.

14 Anupam Gupta and Viswanath Nagarajan. *A Stochastic Probing Problem with Applications*, pages 205–216. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.

15 Anupam Gupta, Viswanath Nagarajan, and Sahil Singla. Algorithms and adaptivity gaps for stochastic probing. In *Proceedings of the Twenty-seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '16, pages 1731–1747, Philadelphia, PA, USA, 2016. Society for Industrial and Applied Mathematics. URL: `http://dl.acm.org/citation.cfm?id=2884435.2884555`.

16 R. Laishram, K. Areekijseree, and S. Soundarajan. Predicted max degree sampling: Sampling in directed networks to maximize node coverage through crawling. In *2017 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pages 940–945, May 2017. `doi:10.1109/INFCOMW.2017.8116502`.

17 Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection. `http://snap.stanford.edu/data`, June 2014.

18 Jure Leskovec and Rok Sosič. Snap: A general-purpose network analysis and graph-mining library. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 8(1):1, 2016.

19 Will Ma. *Improvements and Generalizations of Stochastic Knapsack and Multi-Armed Bandit Approximation Algorithms: Extended Abstract*, pages 1154–1163. URL: `http://epubs.siam.org/doi/abs/10.1137/1.9781611973402.85`, `arXiv:http://epubs.siam.org/doi/pdf/10.1137/1.9781611973402.85`, `doi:10.1137/1.9781611973402.85`.

20 S. Soundarajan, T. Eliassi-Rad, B. Gallagher, and A. Pinar. Maxreach: Reducing network incompleteness through node probes. In *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 152–157, Aug 2016. `doi:10.1109/ASONAM.2016.7752227`.

21 Sucheta Soundarajan, Tina Eliassi-Rad, Brian Gallagher, and Ali Pinar. *epsilon* w gxx: Adaptive edge probing for enhancing incomplete networks. In *Proceedings of the 2017 ACM on Web Science Conference*, WebSci '17, pages 161–170, New York, NY, USA, 2017. ACM. URL: `http://doi.acm.org/10.1145/3091478.3091492`, `doi:10.1145/3091478.3091492`.

22 M. Vidal, ME Cusick, and AL Barabasi. Interactome networks and human disease. *Cell.*, 144(6):986 – 98, 2011.