

1 Stochastic Graph Exploration *

2 Aris Anagnostopoulos

3 Sapienza University of Rome

4 aris@diag.uniroma1.it

5 Ilan R. Cohen

6 CWI, Amsterdam

7 ilanrcohen@gmail.com

8 Stefano Leonardi ¹

9 Sapienza University of Rome

10 leonardi@diag.uniroma1.it

11 Jakub Łącki

12 Google Research, New York

13 jlacki@google.com

14 — Abstract —

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

18 To model this process, we introduce the *stochastic graph exploration problem*. The input is an
19 undirected graph $G = (V, E)$ with a source vertex s , stochastic edge costs drawn from a distribution
20 $\pi_e, e \in E$, and rewards on vertices of maximum value R . The goal is to find a set F of edges of total
21 cost at most B such that the subgraph of G induced by F is connected, contains s , and maximizes
22 the total reward. This problem generalizes the stochastic knapsack problem and other stochastic
23 probing problems recently studied.

24 Our focus is on the development of efficient nonadaptive strategies that are competitive against
25 the optimal adaptive strategy. A major challenge is the fact that the problem has an $\Omega(n)$ adaptivity
26 gap even on a tree of n vertices. This is in sharp contrast with $O(1)$ adaptivity gap of the stochastic
27 knapsack problem, which is a special case of our problem. We circumvent this negative result by
28 showing that $O(\log nR)$ resource augmentation suffices to obtain $O(1)$ approximation on trees and
29 $O(\log nR)$ approximation on general graphs. To achieve this result, we reduce stochastic graph
30 exploration to a memoryless process—the *minesweeper* problem—which assigns to every edge a
31 probability that the process terminates when the edge is probed. For this problem, interesting in its
32 own, we present an optimal polynomial time algorithm on trees and an $O(\log nR)$ approximation
33 for general graphs.

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

39 **2012 ACM Subject Classification** Theory of computation → Graph algorithms analysis; Theory of
40 computation → Stochastic approximation

41 **Keywords and phrases** stochastic optimization, graph exploration, approximation algorithms

42 **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".



43 **1** Introduction

44 Exploring networked data is a time consuming and expensive task which is also subject to
 45 several limitations. For example, social networks can be explored only through the use of
 46 specific APIs made available by the provider which restrict the number of nodes that can
 47 be probed and limit the number of neighbors of each node that can be discovered with one
 48 probe. The cost and the difficulty of exploring large-scale networks can be an obstacle to
 49 collecting suitable snapshots for the purpose of testing new network analysis tools. The
 50 testing is more often executed on static networks made available in public repositories [20,21]
 51 collected in the past for other purposes. It is therefore of crucial importance the development
 52 of effective and efficient methods to explore large-scale networks.

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

60 One main difficulty in the definition of an effective probing strategy is the intrinsic
 61 uncertain nature in terms of cost and probability of success of the process of discovering links
 62 in a network, especially if these links represent complex relationships between nodes. In order
 63 to confirm the existence of a link between two nodes, it may be required to execute several
 64 experiments whose outcome cannot be predicted in advance. Examples are the in-vitro
 65 reactions between proteins needed to discover protein-to-protein interaction networks [8,26]
 66 or the influence between humans in social networks.

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

76 Several recent works [19,24,25] have focused on the task of exploring real-world networks
 77 when a limited budget is available. However, these papers do not provide a comprehensive
 78 theoretical study of these problems. In this work we initiated the study of exploring an
 79 undirected network from a root node. The graph has costs on the edges and rewards on the
 80 vertices. A budget limits the total cost of the of the graph edges that are probed.

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

90 of the edges in F remains bounded by a prespecified budget B . As soon as we probe an edge
 91 such that the total cost exceeds B the process terminates.

92 In the stochastic graph exploration problem, we aim to design simple polynomial-time
 93 computable *nonadaptive strategies* with a reward as close as possible to the reward obtained
 94 by the optimal *adaptive strategy*, which decides on the next edge to be traversed after the
 95 cost of all previously traversed edges is revealed (see Section 2 for precise definitions). This
 96 is customary in a class of stochastic optimization problems [6], for which it is common to
 97 bound the *adaptivity gap* of the nonadaptive strategy.

98 The stochastic graph exploration problem generalizes some important stochastic opti-
 99 mization problems. If the graph G is a star graph, our problem models *exactly* the stochastic
 100 knapsack problem [6, 10]. Stochastic knapsack admits an $O(1)$ adaptivity gap, that is, there
 101 exists an optimal nonadaptive strategy, which approximates the optimal adaptive strategy
 102 up to a constant factor. The nonadaptive strategy is devised by exploiting a suitable LP
 103 relaxation for the problem because the standard formulation has an unbounded integrality
 104 gap defined as the worst-case ratio between the optimal integral cost and optimal fractional
 105 cost of the LP. In the LP version of the problem that is used, the costs of the edges are
 106 reduced to their truncated (by the maximum budget) expected costs and the rewards are
 107 also reduced by the probability that the cost of the item is below the maximum budget.

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

115 The stochastic graph exploration problem also models stochastic graph probing problems.
 116 Probing problems in graphs have been introduced [9, 16] because of their applications to kidney
 117 exchange and online dating. Consider a probing probability for each edge $p : E \rightarrow [0, 1]$,
 118 that is, edge e will materialize with probability $p(e)$ each time is probed, independently of
 119 the other edges and of the previous probes. The goal is to maximize the number of vertices
 120 that are connected to a source vertex s by the set F of edges that have been successfully
 121 probed when the total number of probes is limited by B . Nonadaptive strategies probe a list
 122 of edges in a sequence till success or the total budget B is reached. The stochastic graph
 123 exploration problem we study models the stochastic graph probing problem by setting the
 124 costs of the edges distributed according to $\Pr(C_e = i) = (1 - p(e))^{i-1}p(e)$, with i being the
 125 number of probes needed to discover edge e .

126 1.1 Summary of Our Results

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

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

134 One first direction we pursue to circumvent the impossibility result is to allow a limited
 135 amount of resource augmentation: instead of using budget B , we allow the algorithm to use a
 136 budget of $\beta \cdot B$, for some value of β . We call an algorithm (α, β) -approximate if it computes

131:4 Stochastic Graph Exploration

137 a strategy which uses budget $\beta \cdot B$, and obtains an expected reward of at least $1/\alpha$ times the
138 optimal reward (obtained by an adaptive algorithm). We present polynomial time computable
139 nonadaptive strategies in a graph of n vertices that are $(O(1), O(\log nR))$ -approximate for
140 trees and $(O(\log nR), O(\log nR))$ -approximate for general graphs, with R being the maximum
141 reward of a vertex.

142 The idea is to transform the stochastic exploration problem into a memoryless stochastic
143 process, which we call the *minesweeper problem*, and which may be of independent interest.
144 In the minesweeper problem, the budget and the edge costs are replaced by probabilities
145 $p(e)$, which are specified for every edge e . When an edge e is probed, the process stops with
146 probability $1 - p(e)$. Hence, the final reward of a vertex is discounted by the probability that
147 the strategy does not stop before the vertex is acquired. The minesweeper problem is, in fact,
148 a special case of stochastic graph exploration, where the support of each π_e (distribution of
149 cost of edge e) is $\{0, B + 1\}$ and the budget is B .

150 We prove that an α -approximate strategy for the minesweeper problem implies an
151 $(O(\alpha), O(\log nR))$ -approximate nonadaptive strategy for the stochastic graph exploration
152 problem. The idea of the reduction is as follows. We construct a minesweeper problem
153 instance, where $p(e) = \Pr(\pi_e < X_B)$, where X_B is random variable that follows an expo-
154 nential distribution with parameter B . We first show that, for any subset of edges F , the
155 probability that their total cost in the stochastic graph exploration is at most B is at most a
156 constant factor of the probability that minesweeper would stop on this set. On the other
157 hand, the expected additional reward that can be achieved from minesweeper after the total
158 cost becomes larger than $O(B \log nR)$ is negligible.

159 We then show how to compute in polynomial time an optimal strategy for the minesweeper
160 problem on trees and an $O(\log nR)$ -approximate strategy on general graphs. These results im-
161 ply imply an $(O(1), O(\log nR))$ -approximate strategy for trees and an $(O(\log nR), O(\log nR))$ -
162 approximate strategy for general graphs. To show the optimal result on trees we prove
163 two facts. First, the order of traversal of the edges in each subtree can be determined
164 independently. Second, we show a simple optimality condition which helps us determine how
165 many edges from each subtree should be probed before switching to a different subtree. We
166 remark that our approach is in a spirit similar to the greedy optimal strategy defined by
167 the Gittins index [11, 12] for multi-armed bandit problems. However, differently from the
168 standard setting of the Gittins index, in the minesweeper problem, a whole new set of arms
169 is made available for each node of the tree reached by the exploration process. Moreover,
170 in the minesweeper problem, the discount factor is not constant because it depends on the
171 probability assigned to the traversed edge. This approach is not viable for general graphs,
172 and we provide an approximate solution instead, by showing a reduction of minesweeper to
173 max-prize problem [7].

174 We also pursue a second direction to circumvent the lower bound on the adaptivity
175 gap for trees: we restrict the distributions by considering the case when the edge costs are
176 bounded by $\frac{\epsilon^2 B}{c \log n}$ for a suitable constant c . We show, under this condition, the existence of
177 a polynomial time computable $(1 + \epsilon, 1 + \epsilon)$ -approximate nonadaptive strategy for trees and
178 $(1 + \epsilon, 8 + \epsilon)$ -approximate nonadaptive strategy for any graph G . We note that this approach
179 can be extended to prove a result with resource augmentation similar to the one we obtained
180 through reduction to the minesweeper problem. Yet, we believe that both the minesweeper
181 problem and the reduction technique can be of independent interest.

182 Our final result is related to the problem of finding a nonadaptive probing strategy that
183 is $(o(n), O(1))$ -approximate. We leave open this challenging problem even for trees. However,
184 we establish an interesting result for the characterization of nonadaptive strategies. We prove

185 that any nonadaptive strategy that probes edges in order until it succeeds or until the budget
 186 is exceeded can be $(O(1), O(1))$ -approximated by a set strategy, which probes all edges at
 187 once and obtains a reward only if all edges of a set are successfully probed within budget.
 188 We specifically prove that the adaptivity gap of a nonadaptive strategy can be approximated
 189 up to a factor of 6 by a set strategy that uses budget $9B$. We use this result to give an
 190 algorithm for finding a strategy for trees, which is $(O(1), O(1))$ -approximate, compared to
 191 the best *nonadaptive* strategy. Surprisingly, the resulting strategy is adaptive.

192 1.2 Related Work

193 The adaptivity gap of stochastic problems has been studied for the knapsack problem [6, 10]
 194 which is a special case of the problem we study. The adaptivity gap has also been studied
 195 for budgeted multi-armed bandits [13, 14, 22] by resorting to suitable linear programming
 196 relaxation. Differently from previous work on budgeted multi-armed bandit problems, we
 197 consider the setting in which new arms appear after some arms are pulled. Stochastic probing
 198 problems have also been studied for matching [1, 4, 9] motivated from kidney exchange and
 199 for more general classes of matroid optimization problems [16, 17].

200 The stochastic graph exploration problem we introduce is also related to the *stochastic*
 201 *orienteeing* problem [5, 15]. In stochastic orienteeing, the set of traversed edges must form
 202 a path in a metric graph with deterministic costs on the edges, while the time spent on a
 203 node is a random variable, which follows an a-priori known distribution. In stochastic graph
 204 exploration, the random variables are the costs of the edges of the graph but we cannot
 205 ensure that the costs on the edges form a metric since the random variables are independent.

206 1.3 Organization of the Paper

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

214 2 Problem Definition

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

218 Let us now define the STOCHASTICEXPLORATION problem (in the following sometimes
 219 abbreviated by SGE). This problem instance is given by a tuple (G, s, C, w) , where G is an
 220 undirected graph $G = (V, E)$, $s \in V$ is a source vertex, C is a function that assigns *stochastic*
 221 edge costs to each edge, and $w : V \rightarrow \mathbb{R}_{\geq 0}$ is a function that assigns (deterministic) reward
 222 to each vertex.² And we denote R as the maximum reward of a vertex i.e. $R = \max_{v \in V} w(v)$.
 223 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.

131:6 Stochastic Graph Exploration

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

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

235 Let us now define the MINESWEEPER problem, which we often abbreviate to MS. This
 236 problem is defined by a tuple (G, s, p, w) , where G is an undirected graph, $s \in V$ is a
 237 start vertex, $p : E \rightarrow [0, 1]$ is a function that assigns to each edge e the probability that
 238 e materializes and $w : V \rightarrow \mathbb{R}_{\geq 0}$ is a function that assigns (deterministic) reward to each
 239 vertex. The only difference between MS and SGE is in how edges are probed. There are no
 240 edge costs or budget. Instead, whenever an edge e is probed, it materializes (independently
 241 of the other edges) with probability $p(e)$ and is acquired immediately. If the edge does
 242 not materialize, the process ends immediately. Note that as in SGE, probing an edge e is
 243 only allowed when $F \cup \{e\}$ is connected to s . Note that the MINESWEEPER problem is a
 244 special case of the STOCHASTICEXPLORATION problem, by letting, for each edge e , π_e be
 245 the distribution in which with probability $p(e)$ we obtain the value 0 and with probability
 246 $1 - p(e)$ the value $B + 1$.

247 We consider the following types of strategies for both problems:

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

256 For a strategy S for SGE, we denote by $r(\mathcal{I}_{\text{SGE}}, S, B)$ the expected payoff of strategy S for
 257 the SGE problem instance $\mathcal{I}_{\text{SGE}} = (G, s, C, w)$ with initial budget of B , which is the expected
 258 reward of the set of nodes in the returned solution. When $B = 1$ we sometimes omit the
 259 third argument of $r(\cdot)$. Similarly, we denote by $r_{\text{MS}}(\mathcal{I}_{\text{MS}}, S)$ the expected payoff of strategy
 260 S for the MS problem instance \mathcal{I}_{MS} . We call a strategy S *optimal* for \mathcal{I} with budget B , if
 261 for all strategies S' , $r(\mathcal{I}, S, B) \geq r(\mathcal{I}, S', B)$. Let OPT_{ad} be the optimal adaptive strategy
 262 for the SGE problem and OPT_{na} be the optimal nonadaptive strategy. We call a strategy
 263 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
 264 ALG is (α, β) -approximate if for any instance \mathcal{I} it computes a α -approximate strategy by
 265 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.

266 **3 Lower Bounds**

267 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 $r(\mathcal{I}_{LB}, OPT_{ad})/r(\mathcal{I}_{LB}, OPT_{na}) = \Omega(n)$, where OPT_{ad} and OPT_{na} denote the optimal adaptive and nonadaptive strategies.

270 The instance \mathcal{I}_{LB} is shown in Figure 1. The graph G contains the set of nodes $\{s, u_1, u_2, \dots, u_\ell, v_1, \dots, v_\ell\}$, 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.

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

276 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.

280 ▶ **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 $r(\mathcal{I}, OPT_{ad})/r(\mathcal{I}_{TR}, OPT_{na}) = \Omega(n/2^{2c})$.*

285 **4 The General Case and the Minesweeper Problem**

286 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:

293 ▶ **Observation 3.** *There exists an optimal strategy for the MINESWEEPER problem that is nonadaptive.*

295 **4.1 Reduction from STOCHASTICEXPLORATION to MINESWEEPER**

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

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

309 **► Lemma 4.** *For any $F \subseteq E$ we have that $\Pr(\mathcal{G}_F) \leq e \cdot \Pr(\mathcal{E}_F)$.*

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

314 **► 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
 315 an instance for MS as defined previously. Let OPT_{ad} denote the optimal adaptive strategy
 316 for SGE and OPT_{MS} the optimal strategy for MS. We have that*

$$317 \quad r((G, s, C, w), \text{OPT}_{ad}, 1) \leq e \cdot r_{MS}((G, s, C, w), \text{OPT}_{MS}).$$

318 Recall from Observation 3 that the optimal strategy for the MINESWEEPER problem is
 319 nonadaptive, therefore it can be specified by a list of edges that are selected sequentially
 320 until for one of them there is a failure. Let OPT_{MS} be such an optimal sequence of edges.
 321 Next we show that the sequence of edges OPT_{MS} can provide an approximate result to the
 322 STOCHASTICEXPLORATION problem if we allow for some budget augmentation.

323 **► Lemma 6.** *Consider an SGE instance $\mathcal{I}_{\text{SGE}} = (G, s, C, w)$ and let $\mathcal{I}_{\text{MS}} = (G, s, p, w)$ be
 324 an instance for MS as defined previously. Let OPT_{MS} be the optimal sequence of edges for the
 325 MINESWEEPER instance, and let S be the (nonadaptive) strategy for STOCHASTICEXPLORATION
 326 that probes the same edges, in the same order. Then we have that*

$$327 \quad r((G, s, C, w), S, 2 \ln(nR)) \geq r_{MS}((G, s, C, w), \text{OPT}_{MS}) - o(1),$$

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

329 Collecting the results of Lemmas 5 and 6 we obtain the following theorem.

330 **► Theorem 7.** *Consider an SGE instance $\mathcal{I}_{\text{SGE}} = (G, s, C, w)$ and let $\mathcal{I}_{\text{MS}} = (G, s, p, w)$ be
 331 an instance for MS as defined previously. Then*

$$332 \quad r((G, s, C, w), \text{OPT}_{na}, 2 \ln(nR)) + o(1) \geq r_{MS}((G, s, C, w), \text{OPT}_{MS}) \geq \frac{r((G, s, C, w), 1, \text{OPT}_{ad})}{e}.$$

4.2 MINESWEEPER on Trees

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

► **Theorem 8.** *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 8 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 9.** *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 $S^* = (e_1, \dots, e_k)$. 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 $w(e_1, \dots, e_i) = \sum_{j=1}^i w(e_j)$. Then S^* is a sequence that maximizes

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

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

We can write after some manipulations:

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

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

131:10 Stochastic Graph Exploration

371 For this purpose, we use the problem of *max-prize tree*. In the max-prize tree [7] we are
372 given an undirected graph $G = (V, E)$ with a source vertex $s \in V$, (deterministic) edge costs
373 $c : E \rightarrow \mathbb{R}_{\geq 0}$, deterministic rewards on the vertices $w : V \rightarrow \mathbb{R}_{\geq 0}$, and a budget $B \in \mathbb{R}$. The
374 objective is to build a subgraph $G' = (V', E')$ of G such that (1) G' is a tree, (2) $s \in V'$, and
375 (3) $\sum_{e \in E'} c(e) \leq B$, that maximizes $\sum_{v \in V'} w(v)$.

376 We use for our approximation the 8-approximation algorithm for the max-prize-tree
377 problem given by Blum et al. [7].

378 5 Approximating Set Strategy on Trees

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

381 ► **Lemma 10.** *Let $\mathcal{I} = (G, s, C, w)$ be a SGE instance, where G is a tree. Let OPT_{set} be the
382 optimal set strategy for \mathcal{I} . Then, in $O(n^4/\epsilon^2)$ time we can compute an adaptive strategy S ,
383 such that $r(\mathcal{I}, S, 1 + \epsilon) \geq r(\mathcal{I}, OPT_{set}, 1)$. Moreover, if edge costs are not stochastic, that is,
384 the support of each distribution π_e has size 1, the algorithm runs in $O(n^3/\epsilon)$ time and the
385 resulting strategy is not adaptive.*

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

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

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

407 ► **Lemma 11.** *Let A be a nonempty feasible set of edges of T and let e be the maximal edge
408 of A . Given e (and without knowing A) we can compute the set F_e of all edges f such that
409 $e \prec f$ and $A \cup \{f\}$ is a feasible set.*

410 6 Bounded Edge Costs

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

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

415 **► Theorem 12.** *Let $\mathcal{I} = (G, s, C, w)$ be an instance of SGE, where $C(e) = O(\frac{\epsilon^2}{\ln n})$ (for each
416 edge e and some $0 < \epsilon = O(1)$), $R \leq \epsilon n^{O(1)}$, and the smallest reward is 1. Then, in polynomial
417 time, we can compute a nonadaptive $(O(1), 1 + \epsilon)$ -approximate strategy for \mathcal{I} . Additionally,
418 if G is a tree, then in time $O(n^3/\epsilon)$ we can compute a nonadaptive $(1 + \epsilon, 1 + \epsilon)$ -approximate
419 strategy for \mathcal{I} .*

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

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

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

428 **► Lemma 13.** *Let $0 < \epsilon < 1/3$ and let $\mathcal{I} = (G, s, C, w)$ be an instance of SGE, in which
429 $B \geq 5c/\epsilon^2 \cdot \ln n$. Let F be a set of edges acquired by some adaptive strategy. If $\mu(F) \geq (1 + \epsilon) \cdot B$
430 then the probability that $C(F) \leq B$ is at most n^{-c} .*

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

434 **► Lemma 14.** *Let $\mathcal{I} = (G, s, C, w)$ be an instance of SGE. For any set of edges F and any
435 $\tilde{B} \geq 5c/\epsilon^2 \cdot \ln n$, if $\mu(F) = \tilde{B}$ then the probability that $C(F) \geq (1 + \epsilon)\tilde{B}$ is at most n^{-c} .*

436 **► Lemma 15.** *Let $\mathcal{I} = (G, s, C, w)$ be an instance of SGE, where $B \geq 5c/\epsilon^2 \ln n$, the
437 maximum reward R satisfies $R \leq \epsilon n^{c-1}$, and the minimum reward is 1. Let \mathcal{I}_e be obtained
438 from \mathcal{I} by replacing each edge cost with its expected value. Let OPT_{set}^ϵ be the optimal set
439 strategy using budget $(1 + \epsilon)B$ for \mathcal{I}_e and OPT_{ad} be the optimal adaptive strategy using budget
440 B for \mathcal{I} . Then, $(1 + \epsilon)r(\mathcal{I}, OPT_{set}^\epsilon, (1 + \epsilon)B) \geq r(\mathcal{I}, OPT_{ad}, B)$.*

441 Observe that finding the optimal set strategy on \mathcal{I}_e is NP-hard, as it generalizes the
442 knapsack problem. However, it becomes tractable, if we augment the budget. In particular,
443 for trees, we use the algorithm of Lemma 10, and for general graphs, in Section 4.3, we show
444 how to use the solution of the max-prize problem.

445 **7 Nonadaptive strategies**

446 In this section we consider nonadaptive strategies for the stochastic exploration problem.
447 The main result of this section is that, for the graph exploration problem, that there exists a
448 *set-strategy* with a constant budget augmentation, which is a constant competitive compared
449 to the best nonadaptive algorithm. Recall that, a *set-strategy* is to choose a set of edges
450 (without an internal order) and to try to probe all of the edge in that set. The gain of
451 strategy for a set of edges, is nonzero only if the *entire set* was successfully probed (i.e., if
452 the total cost of the set is smaller than the budget), and then it collects the rewards of all
453 the vertices connected to this set. Therefore, the expected gain of *set-strategy* given a set

131:12 Stochastic Graph Exploration

454 of edges, is the total gain of vertices spanned by these edges times the probability that the
455 total cost of these edges would not be greater than the specified budget.

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

458 7.1 Power of Budget Augmentation

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

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

464 ► **Lemma 16.** *Assume that for each edge e_i , $i \in [n]$ we have $c_i \in [0, 1]$. Then*

$$465 \quad P_n(3) \geq P_n(1) (1 - \ln(P_n(1))).$$

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

► **Corollary 17.**

$$469 \quad P_n(9) \geq P_n(1) \frac{(1 - \ln(P_n(1)))^2}{2}$$

470 7.2 List Strategy vs. Set Strategy

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

$$476 \quad r(\mathcal{I}, S_{ls}, B) = \sum_{j=1}^n w(v_j) \cdot \Pr(C^j \leq B).$$

477 Given a nonadaptive strategy $S_{ls} = \langle e_1, \dots, e_n \rangle$, consider n different set strategies S_k , for
478 $k = \{1 \dots n\}$, where $S_k = \{e_1, \dots, e_k\}$. Note that the expected payoff of S_k with budget $9 \cdot B$
479 is

$$480 \quad r(\mathcal{I}, S_k, 9B) = \Pr(C^k \leq 9B) \cdot \sum_{j=1}^k w(v_j).$$

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

► **Lemma 18.**

$$483 \quad \max_k \{r(\mathcal{I}, S_k, 9B)\} \geq 0.46 \cdot r(\mathcal{I}, S_{ls}, B).$$

484 **7.3 Algorithm for Trees**

485 By combining Lemma 18 with the algorithm of Lemma 10, we obtain the following.

486 ► **Theorem 19.** *Let $\mathcal{I} = (G, s, C, w)$ be a SGE instance, where G is a tree. Let OPT_{na} be*
 487 *the optimal nonadaptive strategy for \mathcal{I} . Then, in $O(n^4/\epsilon^2)$ time we can compute an adaptive*
 488 *strategy S , such that $r(\mathcal{I}, S, 9 + \epsilon) \geq 0.46 \cdot r(\mathcal{I}, OPT_{na}, 1)$.*

489 **8 Conclusions**

490 In this work we have introduced the stochastic exploration problem on graphs which gener-
 491 alizes the stochastic knapsack problem [6, 10]. We proved that, differently from stochastic
 492 knapsack, no $o(n)$ adaptivity gap is possible unless we allow some resource augmentation on
 493 the budget. We provided algorithms with bounded adaptivity gap and logarithmic resource
 494 augmentation by reducing stochastic exploration to a related memoryless problem—the
 495 minesweeper problem. We also considered the case of edges with small costs for which it is
 496 possible to provide an algorithm with $O(1)$ adaptivity gap and $O(1)$ resource augmentation.
 497 The most challenging problem left open from our work is the one of devising an algorithm
 498 with $O(1)$ approximation factor that uses only $O(1)$ resource augmentation for general graphs.
 499 The problem is open even for trees. We provided a set of additional results on the structure
 500 of optimal adaptive strategies and on the power of resource augmentation for set strategies
 501 with respect to list strategies that can help in addressing this problem.

502 **References**

-
- 503 1 Marek Adamczyk. Improved analysis of the greedy algorithm for stochastic matching. *Informa-*
 504 *tion 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>.
- 507 2 D Adolphson and T Ch Hu. Optimal linear ordering. *SIAM Journal on Applied Mathematics*,
 508 25(3):403–423, 1973.
- 509 3 Donald L. Adolphson. Single machine job sequencing with precedence constraints. *SIAM J.*
 510 *Comput.*, 6(1):40–54, 1977. URL: <https://doi.org/10.1137/0206002>, doi:10.1137/0206002.
- 511 4 Nikhil Bansal, Anupam Gupta, Jian Li, Julián Mestre, Viswanath Nagarajan, and Atri Rudra.
 512 When lp is the cure for your matching woes: Improved bounds for stochastic matchings.
 513 *Algorithmica*, 63(4):733–762, Aug 2012.
- 514 5 Nikhil Bansal and Viswanath Nagarajan. *On the Adaptivity Gap of Stochastic Orienteering*,
 515 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.
- 517 6 Anand Bhalgat, Ashish Goel, and Sanjeev Khanna. *Improved Approximation Re-*
 518 *sults 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.
- 521 7 Avrim Blum, Shuchi Chawla, David R. Karger, Terran Lane, Adam Meyerson, and Maria
 522 Minkoff. Approximation algorithms for orienteering and discounted-reward tsp. *SIAM Journal*
 523 *on Computing*, 37(2):653–670, 2007.
- 524 8 Michael Caldera, Pisanu Buphamalai, Felix Mueller, and Jorg Menche. Interactome-based
 525 approaches to human disease. *Current Opinion in Systems Biology*, 3:88 – 94, 2017.
- 526 9 Ning Chen, Nicole Immorlica, Anna R. Karlin, Mohammad Mahdian, and Atri Rudra.
 527 *Approximating Matches Made in Heaven*, pages 266–278. Springer Berlin Heidelberg,
 528 Berlin, Heidelberg, 2009. URL: http://dx.doi.org/10.1007/978-3-642-02927-1_23, doi:
 529 10.1007/978-3-642-02927-1_23.

- 530 **10** Brian C. Dean, Michel X. Goemans, and Jan Vondrak. Approximating the stochastic knapsack
531 problem: The benefit of adaptivity. *Mathematics of Operations Research*, 33(4):945–964, 2008.
532 doi:10.1287/moor.1080.0330.
- 533 **11** Esther Frostig and GIDEON WEISS. Four proofs of gittins multiarmed bandit theorem, 1999.
- 534 **12** Author(s) J. C. Gittins and J. C. Gittins. Bandit processes and dynamic allocation indices.
535 *Journal of the Royal Statistical Society, Series B*, pages 148–177, 1979.
- 536 **13** Sudipto Guha and Kamesh Munagala. Approximation algorithms for budgeted learning
537 problems. In *Proceedings of the Thirty-ninth Annual ACM Symposium on Theory of Computing*,
538 STOC '07, pages 104–113, New York, NY, USA, 2007. ACM. URL: [http://doi.acm.org/](http://doi.acm.org/10.1145/1250790.1250807)
539 [10.1145/1250790.1250807](http://doi.acm.org/10.1145/1250790.1250807), doi:10.1145/1250790.1250807.
- 540 **14** Anupam Gupta, Ravishankar Krishnaswamy, Marco Molinaro, and R. Ravi. Approximation
541 algorithms for correlated knapsacks and non-martingale bandits. In *IEEE 52nd Annual*
542 *Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA,*
543 *October 22-25, 2011*, pages 827–836, 2011.
- 544 **15** Anupam Gupta, Ravishankar Krishnaswamy, Viswanath Nagarajan, and R. Ravi. Running
545 errands in time: Approximation algorithms for stochastic orienteering. *Mathematics of*
546 *Operations Research*, 40(1):56–79, 2015.
- 547 **16** Anupam Gupta and Viswanath Nagarajan. *A Stochastic Probing Problem with Applications*,
548 pages 205–216. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.
- 549 **17** Anupam Gupta, Viswanath Nagarajan, and Sahil Singla. Algorithms and adaptivity gaps for
550 stochastic probing. In *Proceedings of the Twenty-seventh Annual ACM-SIAM Symposium on*
551 *Discrete Algorithms*, SODA '16, pages 1731–1747, Philadelphia, PA, USA, 2016. Society for
552 Industrial and Applied Mathematics. URL: [http://dl.acm.org/citation.cfm?id=2884435.](http://dl.acm.org/citation.cfm?id=2884435.2884555)
553 [2884435.2884555](http://dl.acm.org/citation.cfm?id=2884435.2884555).
- 554 **18** WA Horn. Single-machine job sequencing with treelike precedence ordering and linear delay
555 penalties. *SIAM Journal on Applied Mathematics*, 23(2):189–202, 1972.
- 556 **19** R. Laishram, K. Areekijseree, and S. Soundarajan. Predicted max degree sampling: Sampling
557 in directed networks to maximize node coverage through crawling. In *2017 IEEE Conference*
558 *on Computer Communications Workshops (INFOCOM WKSHPs)*, pages 940–945, May 2017.
559 doi:10.1109/INFOCOMW.2017.8116502.
- 560 **20** Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection.
561 <http://snap.stanford.edu/data>, June 2014.
- 562 **21** Jure Leskovec and Rok Sosič. Snap: A general-purpose network analysis and graph-mining
563 library. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 8(1):1, 2016.
- 564 **22** Will Ma. *Improvements and Generalizations of Stochastic Knapsack and Multi-Armed Bandit*
565 *Approximation Algorithms: Extended Abstract*, pages 1154–1163. URL: [http://epubs.siam.](http://epubs.siam.org/doi/abs/10.1137/1.9781611973402.85)
566 [org/doi/abs/10.1137/1.9781611973402.85](http://epubs.siam.org/doi/abs/10.1137/1.9781611973402.85), arXiv:[http://epubs.siam.org/doi/pdf/10.](http://epubs.siam.org/doi/pdf/10.1137/1.9781611973402.85)
567 [1137/1.9781611973402.85](http://epubs.siam.org/doi/pdf/10.1137/1.9781611973402.85), doi:10.1137/1.9781611973402.85.
- 568 **23** Colin McDiarmid. *Concentration*, pages 195–248. Springer Berlin Heidelberg, Berlin,
569 Heidelberg, 1998. URL: https://doi.org/10.1007/978-3-662-12788-9_6, doi:10.1007/
570 [978-3-662-12788-9_6](https://doi.org/10.1007/978-3-662-12788-9_6).
- 571 **24** S. Soundarajan, T. Eliassi-Rad, B. Gallagher, and A. Pinar. Maxreach: Reducing network
572 incompleteness through node probes. In *2016 IEEE/ACM International Conference on*
573 *Advances in Social Networks Analysis and Mining (ASONAM)*, pages 152–157, Aug 2016.
574 doi:10.1109/ASONAM.2016.7752227.
- 575 **25** Sucheta Soundarajan, Tina Eliassi-Rad, Brian Gallagher, and Ali Pinar. *epsilon w gxx*:
576 Adaptive edge probing for enhancing incomplete networks. In *Proceedings of the 2017 ACM*
577 *on Web Science Conference*, WebSci '17, pages 161–170, New York, NY, USA, 2017. ACM.
578 URL: <http://doi.acm.org/10.1145/3091478.3091492>, doi:10.1145/3091478.3091492.
- 579 **26** M. Vidal, ME Cusick, and AL Barabasi. Interactome networks and human disease. *Cell*,
580 144(6):986 – 98, 2011.

A

 Omitted Proofs

581

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

584 **Proof.** Let S_{AD} be the adaptive strategy that probes the edges $(s, u_\ell), (s, u_{\ell-1}), (s, u_{\ell-2}), \dots$
585 until it probes an edge (s, u_i) whose cost turns out to be 0. After that it probes the edge
586 (u_k, v_k) . Note that right before probing edge (u_k, v_k) the budget that has been used so far is

$$587 \quad 2^{-\ell} + 2^{-(\ell-1)} + \dots + 2^{-(k+1)} < 2^{-k} - 2^{-(\ell+1)}.$$

588 We have that $C(u_k, v_k) = 1 - (2^{-k} - 2^{-(\ell+1)})$ with probability 1, therefore there is sufficient
589 budget to acquire the edge (u_k, v_k) and obtain the reward of $w(v_k) = T$. The probability
590 that the cost of some edge (s, u_i) is 0 is

$$591 \quad \sum_{k=0}^{\ell-1} \left(1 - \frac{1}{\ell}\right)^k \cdot \frac{1}{\ell} = \frac{1 - \left(1 - \frac{1}{\ell}\right)^\ell}{1 - \left(1 - \frac{1}{\ell}\right)} \cdot \frac{1}{\ell} \geq 1 - \frac{1}{e}.$$

592 Therefore, the expected payoff of strategy S_{AD} is $r(\mathcal{I}_{LB}, S_{AD}, 1) \geq T \left(1 - \frac{1}{e}\right)$.

593 Consider now any nonadaptive strategy S . If the payoff of S is nonzero, it clearly attempts
594 to probe at least one edge (u_i, v_i) . Consider the first such edge (u_i, v_i) . Clearly, the strategy
595 must probe the edge (s, u_i) before attempting to probe (u_i, v_i) .

596 Thus with probability $1 - 1/\ell$ we have that $C(s, u_i) = 2^{-i}$, and in that case, just
597 before probing edge (u_i, v_i) , the leftover budget is at most $1 - 2^{-i}$. Then, since $C(u_i, v_i) =$
598 $1 - 2^{-i} + 2^{-(\ell+1)}$, the remaining budget is not sufficient, so the game terminates with total
599 payoff 0. In the remaining case we have $C(s, u_i) = 0$, which happens with probability $1/\ell$.
600 In this case the strategy acquires edge (u_i, v_i) , but after that the remaining budget is not
601 sufficient to reach another node v_j . This means that the expected payoff of strategy S is at
602 most $r(\mathcal{I}_{LB}, S, 1) \leq \frac{T}{\ell}$. The lemma follows. ◀

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

608 **Proof.** Let us define an instance $\mathcal{I} = (G, s, C, w)$, where $G = (V, E)$ is a graph such that
609 $V = \{s, u_1, \dots, u_n, v\}$, and $E = \{(s, u_1), (u_1, u_2), \dots, (u_{n-1}, u_n), (s, v)\}$. For each edge
610 $i \in \{1, \dots, 2c+1\}$, $C(u_i, u_{i+1}) = 1$ with probability $1/2$ and 0 otherwise, for $i > 2c+1$.
611 $C(u_i, u_{i+1}) = 0$ with probability 1. The reward function of u_i is 0 for $i \in \{1, \dots, 2c+1\}$ and
612 1 for $i > 2c+1$. In addition, $C(s, v) = 1$ with probability 1, and $w(v) = 1$.

613 The expected truncated cost of the edges (u_i, u_{i+1}) for $i \in \{1, \dots, 2c+1\}$ is $1/2$. Therefore,
614 an algorithm which uses the expected truncated costs will assume it cannot reach u_{2c+2} .
615 Hence it will acquire just the edge (s, v) and obtain a reward of 1. At the same time, with
616 probability 2^{-2c-2} , an algorithm can obtain the reward from all $\Omega(n)$ vertices u_i . This gives
617 a gap of $\Omega(n/2^{2c})$. ◀

618 ► **Lemma 4.** For any $F \subseteq E$ we have that $\Pr(\mathcal{G}_F) \leq e \cdot \Pr(\mathcal{E}_F)$.

619 **Proof.** First, note that, for three independent random variables X_1, X_2, X_3 distributed
620 according to an exponential distribution with the same parameter λ , we have that

$$621 \quad \begin{aligned} \Pr(X_1 > a) \cdot \Pr(X_2 > b) &= \Pr(X_3 > a) \cdot \Pr(X_3 > b) = \Pr(X_3 > a) \cdot \Pr(X_3 > a + b \mid X_3 > a) \\ &= \Pr(X_3 > a + b), \end{aligned}$$

131:16 Stochastic Graph Exploration

622 where the second equality follows from the memorylessness property of the exponential
623 distribution.

624 Let X be a random variable drawn from the exponential distribution with parameter 1.
625 Using the fact that we proved above, we have that

$$626 \quad \Pr(\mathcal{E}_F) = \prod_{e \in F} \Pr(X_e > c_e) = \Pr\left(X > \sum_{e \in F} c_e\right),$$

627 and that

$$628 \quad \Pr(\mathcal{E}_F | \mathcal{G}_F) = \Pr\left(X > \sum_{e \in F} c_e \mid \sum_{e \in F} c_e \leq 1\right) \geq \Pr(X > 1) = e^{-1}.$$

629 Therefore, we finish the proof of the lemma by combining this fact with

$$630 \quad \Pr(\mathcal{E}_F) \geq \Pr(\mathcal{E}_F, \mathcal{G}_F) = \Pr(\mathcal{E}_F | \mathcal{G}_F) \cdot \Pr(\mathcal{G}_F).$$

631 ◀

632 **► 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
633 an instance for MS as defined previously. Let OPT_{ad} denote the optimal adaptive strategy
634 for SGE and OPT_{MS} the optimal strategy for MS. We have that

$$635 \quad r((G, s, C, w), \text{OPT}_{\text{ad}}, 1) \leq e \cdot r_{\text{MS}}((G, s, C, w), \text{OPT}_{\text{MS}}).$$

636 **Proof.** Consider the optimal adaptive strategy OPT_{ad} and a strategy for minesweeper (call it
637 S) that selects to probe the same edges as OPT_{ad} . We will lower bound the probability that
638 OPT_{ad} terminates (i.e., exhausts its budget) before the strategy is applied to MS. Consider
639 a graph G , and let $g(x, G)$ be the probability that OPT_{ad} for SGE on graph G with budget
640 x does not finish after strategy S for MS. Abusing notation, define

$$641 \quad g(x, i) = \inf_{G=(V,E):|E|=i} g(x, G).$$

642 We will prove by induction on i and x that $g(x, i) \geq e^{-x}$.

643 First note that for $i = 0$ there is not any edge to probe so we are done. So consider $i \geq 1$
644 and assume that for any $x \geq 0$ and any graph $G' = (V', E')$ with $|E'| \leq i - 1$ we have that
645 $g(x, G') \geq e^{-x}$. For $x = 0$, OPT_{ad} has no budget for SGE so it cannot continue and the
646 claim is trivially true. Thus, consider the case that $i \geq 1$ and $x > 0$. Assume that we have
647 graph $G = (V, E)$ and for $e \in E$ denote by G_e the graph in which edge e has been contracted.
648 Then notice that we have that the probability that OPT_{ad} will not terminate after strategy
649 S of MS is equal to the probability that (1) OPT_{ad} finishes when probing the next edge, or
650 (2) that it does not, and neither does S but in the next steps OPT_{ad} does not terminate
651 after MS. Therefore, for any x and any graph G with i edges, we have that

$$\begin{aligned} 652 \quad g(x, G) &= \sum_{c_e \geq x} \Pr(C(e) = c_e) + \sum_{c_e < x} \Pr(C(e) = c_e) \cdot g(x - c_e, G_e) \cdot \Pr(X > c_e) \\ &\geq \sum_{c_e \geq x} \Pr(C(e) = c_e) + \sum_{c_e < x} \Pr(C(e) = c_e) \cdot g(x - c_e, i - 1) \cdot \Pr(X > c_e) \\ &\geq \sum_{c_e \geq x} \Pr(C(e) = c_e) \cdot e^{-x} + \sum_{c_e < x} \Pr(C(e) = c_e) \cdot e^{c_e - x} \cdot e^{-c_e} \geq e^{-x}, \end{aligned}$$

653 where the equality follows from the fact that in the MS problem that we defined edge e does
 654 not die with probability $\Pr(X > c_e)$, the first inequality from the fact that graph G_e has
 655 $i - 1$ edges, and the second inequality from the fact that $e^{-x} \leq 1$ and from the induction
 656 hypothesis.

657 Because this holds for all graphs G with i edges, we deduce that for any $x \geq 0$ we have
 658 that $g(x, i) \geq e^{-x}$.

659 Therefore, we have that $g(1, G) \geq 1/e$, which means that there exists a strategy for the
 660 MS problem that with probability at least $1/e$ does not stop before strategy OPT_{ad} , and
 661 therefore has expected payoff of at least $r((G, s, C, w), \text{OPT}_{\text{ad}}, 1)/e$. ◀

662 ▶ **Lemma 6.** *Consider an SGE instance $\mathcal{I}_{\text{SGE}} = (G, s, C, w)$ and let $\mathcal{I}_{\text{MS}} = (G, s, p, w)$ be
 663 an instance for MS as defined previously. Let OPT_{MS} be the optimal sequence of edges for the
 664 MINESWEEPER instance, and let S be the (nonadaptive) strategy for STOCHASTICEXPLORATION
 665 that probes the same edges, in the same order. Then we have that*

$$666 \quad r((G, s, C, w), S, 2 \ln(nR)) \geq r_{\text{MS}}((G, s, C, w), \text{OPT}_{\text{MS}}) - o(1),$$

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

668 **Proof.** Let $L = \langle e_1, \dots, e_{n-1} \rangle$, the (optimal) list of edges that OPT_{MS} probes. By definition,
 669 strategy S for STOCHASTICEXPLORATION probes the same edges. We consider an execution
 670 of MINESWEEPER using strategy OPT_{MS} and an execution of STOCHASTICEXPLORATION
 671 with strategy S . We couple the two executions, such that for each edge e the value c_e be the
 672 same in both executions.

673 Consider a materialization of the values c_{e_1} up to $c_{e_{n-1}}$ of the edges in L . Let r be the
 674 index such that

$$675 \quad \sum_{i=1}^r c_{e_i} \leq 2 \ln(nR)$$

676 and

$$677 \quad \sum_{i=1}^{r+1} c_{e_i} > 2 \ln(nR).$$

678 Then notice that the revenue of strategy S for STOCHASTICEXPLORATION is a value W ,
 679 which is the sum of the rewards of the nodes reachable from the source node s using edges
 680 e_1, \dots, e_r . We now show that the expected value of OPT_{MS} is $W + o(1)$. Recall that when
 681 OPT_{MS} probes an edge e_i , the probability that it materializes is $p(e_i) = \Pr(X_{e_i} > c_{e_i})$,
 682 with X_{e_i} being an exponentially distributed random variable with parameter 1. If OPT_{MS}
 683 succeeds up to edge e_r , it acquires reward up to W . Otherwise, the probability that OPT_{MS}
 684 succeeds for more than r edges is at most

$$685 \quad \prod_{i=1}^{r+1} p(e_i) = \prod_{i=1}^{r+1} e^{-c_{e_i}} = e^{-\sum_{i=1}^{r+1} c_{e_i}} < e^{-2 \ln(nR)} = (nR)^{-2}.$$

686 The maximum reward that it can acquire is upper bounded by nR , therefore, the expected
 687 reward of OPT_{MS} (given the values c_{e_i}) is

$$688 \quad W + \frac{nR}{(nR)^2} = W + (nR)^{-1}.$$

131:18 Stochastic Graph Exploration

689 Therefore, for each materialization of the edge costs, we have that the expected reward
 690 of OPT_{MS} on MINESWEEPER is at most $(nR)^{-1}$ higher than the reward of strategy S for
 691 $\text{STOCHASTICEXPLORATION}$, implying that the expected reward of S is at most $(nR)^{-1}$ less
 692 than the expected reward of OPT_{MS} . \blacktriangleleft

693 **► Theorem 9.** *Consider the instance $\mathcal{I} = (G, s, p, w)$ of the minesweeper problem, where
 694 $G = (V, E)$ is an undirected graph. An $O(\log nR)$ -approximate strategy can be computed in
 695 polynomial time.*

696 **Proof.** Assume that the optimal solution is the sequence of edges $S^* = (e_1, \dots, e_k)$. Define
 697 $\mathcal{M}(E')$ to be the event that all the edges in the set E' materialize. Also let $w(e_1, \dots, e_i) =$
 698 $\sum_{j=1}^i w(e_j)$. Then S^* is a sequence that maximizes

$$699 \quad O^* = \sum_{i=1}^k \Pr(\mathcal{M}(\{e_1, \dots, e_i\}), \neg \mathcal{M}(\{e_{i+1}\})) \cdot w(e_1, \dots, e_i).$$

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

702 We can write

$$\begin{aligned} 703 \quad O^* &= \sum_{\ell=0}^{\log nR} \sum_{i \in I(\ell)} \Pr(\mathcal{M}(\{e_1, \dots, e_j\}), \neg \mathcal{M}(\{e_{i+1}\})) \cdot w(e_1, \dots, e_j) \\ &\leq \sum_{\ell=0}^{\log nR} 2w(e_1, \dots, e_{\iota(\ell)}) \cdot \sum_{i \in I(\ell)} \Pr(\mathcal{M}(\{e_1, \dots, e_j\}), \neg \mathcal{M}(\{e_{i+1}\})) \\ &\leq \sum_{\ell=0}^{\log nR} 2w(e_1, \dots, e_{\iota(\ell)}) \cdot \Pr(\mathcal{M}(\{e_1, \dots, e_{\iota(\ell)}\})). \end{aligned}$$

704 Note that the optimal sequence of edges S^* cannot contain a cycle because for each node
 705 $v \in V$ we either reach it through a path and we collect value $w(v)$, or we die before and the
 706 process stops. Furthermore, we collect value $w(v)$ only for nodes reachable from s through
 707 edges of S^* that have materialized. Therefore, the edges in S^* must form a tree.

708 Let $\tilde{E} \subset E$ be the set of edges that defines a tree that contains s and maximizes

$$709 \quad w(\tilde{E}) \cdot \Pr(\mathcal{M}(\tilde{E})).$$

710 For each $\ell = 0, 1, \dots, \log nR$ we have that

$$711 \quad w(\tilde{E}) \cdot \Pr(\mathcal{M}(\tilde{E})) \geq w(e_1, \dots, e_{\iota(\ell)}) \cdot \Pr(\mathcal{M}(e_1, \dots, e_{\iota(\ell)}))$$

712 and we obtain that

$$713 \quad O^* \leq 2 \log(nR) \cdot w(\tilde{E}) \cdot \Pr(\mathcal{M}(\tilde{E})). \quad (1)$$

714 Therefore, our goal becomes that of finding that set of edges \tilde{E} that forms a tree and
 715 maximizes

$$716 \quad w(\tilde{E}) \cdot \Pr(\mathcal{M}(\tilde{E})).$$

717 We next show how to find a tree that approximates this quantity. Consider the graph
 718 with the same vertex and edge set as G , with the same rewards on the vertices $w(\cdot)$, and
 719 with edge costs $\psi : E \rightarrow \mathbb{R}_{\geq 0}$, defined as $\psi(e) = -\log p(e)$. Note that

$$720 \quad \Pr(\mathcal{M}(\tilde{E})) = \prod_{e \in \tilde{E}} p(e) = 2^{-\sum_{e \in \tilde{E}} \psi(e)}.$$

721 Let $\Psi_{\min} = \min_{e \in E} \psi(e)$, and $\Psi_{\max} = \sum_{e \in E} \psi(e)$. For $\ell \in \{\Psi_{\min}, \Psi_{\min} + 1, \dots, \Psi_{\max}\}$,
 722 let T^ℓ be the tree T that (1) has cost $\sum_{e \in T^\ell} \psi(e) \leq \ell$, (2) contains node s , and (3) maximizes
 723 $w(T)$, defined as $\sum_{v \in T} w(v)$, where we say that $v \in T$ if node v belongs to the tree T .

724 This is precisely the max-prize-tree problem, and by the discussion just before the proof
 725 we can deduce that for each $\ell \in \{\Psi_{\min}, \Psi_{\min} + 1, \dots, \Psi_{\max}\}$ we can compute a tree \hat{T}^ℓ with
 726 cost at most ℓ and total weight at least $w(T^\ell)/\gamma$. Note that $\Psi_{\max} \leq -n \log(\max_e p(e))$, so
 727 we solve only a polynomial (in the input length) number of max-prize-tree problems.

728 Let

$$729 \ell^* = \arg \max_{\ell} \{2^{-\ell} \cdot w(T^\ell)\}.$$

730 We have $\sum_{e \in \hat{T}^{\ell^*}} \psi(e) \leq \ell^*$, which implies that $\Pr(\mathcal{M}(\hat{T}^{\ell^*})) = \prod_{e \in \hat{T}^{\ell^*}} p(e) \geq 2^{-\ell^*}$. Define

$$731 \tilde{\ell} = \lfloor -\log \Pr(\mathcal{M}(\tilde{E})) \rfloor$$

732 and notice that $\tilde{E} = T^{\tilde{\ell}}$. Also, notice that we have

$$733 2^{-\tilde{\ell}} \geq \frac{1}{2} \Pr(\mathcal{M}(\tilde{E})).$$

734 Putting everything together, we obtain that, for the solution \hat{T}^{ℓ^*} that we compute, we have

$$735 \Pr(\mathcal{M}(\hat{T}^{\ell^*})) \cdot w(\hat{T}^{\ell^*}) \geq 2^{-\ell^*} \cdot w(\hat{T}^{\ell^*}) \geq \frac{2^{-\ell^*}}{\gamma} \cdot w(T^{\ell^*}) \geq \frac{2^{-\tilde{\ell}}}{\gamma} \cdot w(T^{\tilde{\ell}}) \geq \frac{1}{2\gamma} \Pr(\mathcal{M}(\tilde{E})) \cdot w(\tilde{E}).$$

736 Using this in Equation (1), we conclude that

$$737 \Pr(\mathcal{M}(\hat{T}^{\ell^*})) \cdot w(\hat{T}^{\ell^*}) \geq \frac{1}{4\gamma \log(nR)} \cdot O^*.$$

738

739 **► Lemma 10.** Let $\mathcal{I} = (G, s, C, w)$ be a SGE instance, where G is a tree. Let OPT_{set} be the
 740 optimal set strategy for \mathcal{I} . Then, in $O(n^4/\epsilon^2)$ time we can compute an adaptive strategy S ,
 741 such that $r(\mathcal{I}, S, 1 + \epsilon) \geq r(\mathcal{I}, OPT_{set}, 1)$. Moreover, if edge costs are not stochastic, that is,
 742 the support of each distribution π_e has size 1, the algorithm runs in $O(n^3/\epsilon)$ time and the
 743 resulting strategy is not adaptive. ◀

744 **Proof.** We now describe the algorithm computing the strategy. First, we quantize the edge
 745 costs, by rounding them up to multiples of ϵ/n . Namely, each time a cost of c is incurred, we
 746 actually deduct $\lceil c/(\epsilon/n) \rceil (\epsilon/n)$ from the budget. Because each strategy acquires at most n
 747 edges, this turns a strategy using a budget of 1 into a strategy using budget of $1 + \epsilon$. Hence,
 748 by using budget of $1 + \epsilon$, we can assume that the cost of each edge is a multiple of ϵ/n . In
 749 particular, at every step the remaining budget in has one of $(1 + \epsilon)/(\epsilon/n) = O(n/\epsilon)$ distinct
 750 values.

751 The dynamic programming uses an array $D(e, b)$, indexed by an edge e and the remaining
 752 budget b . Note that the array has size $O(n^2/\epsilon)$. The value $D(e, b)$ denotes what is the
 753 maximum expected reward that a strategy can get in the remaining part of the game if the
 754 first edge to be probed is e and the remaining budget is b . Of course the next edge to be
 755 probed after e has to belong to the set F_e (see Lemma 11). Denote by r_e the reward that we
 756 obtain from acquiring the edge e . We use the following recursive formula:

$$757 D(e, b) = \sum_{i=0}^{b/(\epsilon/n)} \Pr(C(e) = i(\epsilon/n)) \left(r_e + \max_{f \in F_e} D(f, b - i(\epsilon/n)) \right).$$

131:20 Stochastic Graph Exploration

758 Observe that there are $O(n^2/\epsilon)$ values of $D(e, b)$ to compute and evaluating each of
 759 them requires $O(n^2/\epsilon)$ time. Hence, the dynamic programming requires $O(n^4/\epsilon^2)$ time. The
 760 expected payoff of the strategy can be obtained by taking maximum of $D(e, 1 + \epsilon)$ over all
 761 edges e incident to s . The recursive formula directly translates to an adaptive algorithm.
 762 After acquiring the edge e , if the remaining budget is $b - i(\epsilon/n)$, the next edge to probe
 763 is $\arg \max_{f \in F_e} D(f, b - i(\epsilon/n))$. Note that this can only be evaluated once we know the
 764 remaining budget and thus the obtained strategy is adaptive.

765 Clearly, the obtained strategy is the optimal adaptive strategy, among strategies that
 766 probe edges according to the ordering \prec . Because, without loss of generality, we can assume
 767 that each set strategy probes edges according to this order, we immediately get that the
 768 $r(\mathcal{I}, 1 + \epsilon, S) \geq r(\mathcal{I}, 1, \text{OPT}_{\text{set}})$.

769 Finally, let us consider the case of non-stochastic edge costs. Observe that the sum in
 770 the formula for $D(e, b)$, has only single summand, which improves the running time by a
 771 factor of n/ϵ . Moreover, the choices of the algorithm can be simulated beforehand, so the
 772 final strategy is nonadaptive. ◀

773 ► **Lemma 11.** *Let A be a nonempty feasible set of edges of T and let e be the maximal edge
 774 of A . Given e (and without knowing A) we can compute the set F_e of all edges f such that
 775 $e \prec f$ and $A \cup \{f\}$ is a feasible set.*

776 **Proof.** Consider the path S that starts at the root of T and ends with e . Observe that every
 777 edge whose one endpoint is on the path and that comes after e in the order \prec belongs to the
 778 set F_e . Indeed, by adding each such edge to F we obtain a feasible set. It is also easy to see
 779 that adding any other edge larger than e to A would not yield a feasible set. ◀

780 We use the following theorem in the proof of the next lemma.

► **Theorem 20** ([23]). *Let X be a martingale such that $\sum_{i=1}^t \text{Var}[X_i | X_{i-1}] \leq V$ and
 $X_i - X_{i-1} \leq M$. Then,*

$$\Pr(X_t - \mathbf{E}[X] \leq -\lambda) \leq \exp\left(\frac{-\lambda^2}{2 \cdot V + M\lambda/3}\right).$$

781 ► **Lemma 10.** *Let $\mathcal{I} = (G, s, C, w)$ be a SGE instance, where G is a tree. Let OPT_{set} be the
 782 optimal set strategy for \mathcal{I} . Then, in $O(n^4/\epsilon^2)$ time we can compute an adaptive strategy S ,
 783 such that $r(\mathcal{I}, S, 1 + \epsilon) \geq r(\mathcal{I}, \text{OPT}_{\text{set}}, 1)$. Moreover, if edge costs are not stochastic, that is,
 784 the support of each distribution π_e has size 1, the algorithm runs in $O(n^3/\epsilon)$ time and the
 785 resulting strategy is not adaptive.*

786 **Proof.** We now describe the algorithm computing the strategy. First, we quantize the edge
 787 costs, by rounding them up to multiples of ϵ/n . Namely, each time a cost of c is incurred, we
 788 actually deduct $\lceil c/(\epsilon/n) \rceil (\epsilon/n)$ from the budget. Because each strategy acquires at most n
 789 edges, this turns a strategy using a budget of 1 into a strategy using budget of $1 + \epsilon$. Hence,
 790 by using budget of $1 + \epsilon$, we can assume that the cost of each edge is a multiple of ϵ/n . In
 791 particular, at every step the remaining budget in has one of $(1 + \epsilon)/(\epsilon/n) = O(n/\epsilon)$ distinct
 792 values.

793 The dynamic programming uses an array $D(e, b)$, indexed by an edge e and the remaining
 794 budget b . Note that the array has size $O(n^2/\epsilon)$. The value $D(e, b)$ denotes what is the
 795 maximum expected reward that a strategy can get in the remaining part of the game if the
 796 first edge to be probed is e and the remaining budget is b . Of course the next edge to be

797 probed after e has to belong to the set F_e (see Lemma 11). Denote by r_e the reward that we
798 obtain from acquiring the edge e . We use the following recursive formula:

$$799 \quad D(e, b) = \sum_{i=0}^{b/(\epsilon/n)} \Pr(C(e) = i(\epsilon/n)) \left(r_e + \max_{f \in F_e} D(f, b - i(\epsilon/n)) \right).$$

800 Observe that there are $O(n^2/\epsilon)$ values of $D(e, b)$ to compute and evaluating each of
801 them requires $O(n^2/\epsilon)$ time. Hence, the dynamic programming requires $O(n^4/\epsilon^2)$ time. The
802 expected payoff of the strategy can be obtained by taking maximum of $D(e, 1 + \epsilon)$ over all
803 edges e incident to s . The recursive formula directly translates to an adaptive algorithm.
804 After acquiring the edge e , if the remaining budget is $b - i(\epsilon/n)$, the next edge to probe
805 is $\arg \max_{f \in F_e} D(f, b - i(\epsilon/n))$. Note that this can only be evaluated once we know the
806 remaining budget and thus the obtained strategy is adaptive.

807 Clearly, the obtained strategy is the optimal adaptive strategy, among strategies that
808 probe edges according to the ordering \prec . Because, without loss of generality, we can assume
809 that each set strategy probes edges according to this order, we immediately get that the
810 $r(\mathcal{I}, 1 + \epsilon, S) \geq r(\mathcal{I}, 1, \text{OPT}_{\text{set}})$.

811 Finally, let us consider the case of non-stochastic edge costs. Observe that the sum in
812 the formula for $D(e, b)$, has only single summand, which improves the running time by a
813 factor of n/ϵ . Moreover, the choices of the algorithm can be simulated beforehand, so the
814 final strategy is nonadaptive. ◀

815 ▶ **Lemma 11.** *Let A be a nonempty feasible set of edges of T and let e be the maximal edge
816 of A . Given e (and without knowing A) we can compute the set F_e of all edges f such that
817 $e \prec f$ and $A \cup \{f\}$ is a feasible set.*

818 **Proof.** Consider the path S that starts at the root of T and ends with e . Observe that every
819 edge whose one endpoint is on the path and that comes after e in the order \prec belongs to the
820 set F_e . Indeed, by adding each such edge to F we obtain a feasible set. It is also easy to see
821 that adding any other edge larger than e to A would not yield a feasible set. ◀

822 We use the following theorem in the proof of the next lemma.

▶ **Theorem 21** ([23]). *Let X be a martingale such that $\sum_{i=1}^t \mathbf{Var}[X_i | X_{i-1}] \leq V$ and
 $X_i - X_{i-1} \leq M$. Then,*

$$\Pr(X_t - \mathbf{E}[X] \leq -\lambda) \leq \exp\left(\frac{-\lambda^2}{2 \cdot V + M\lambda/3}\right).$$

823 ▶ **Lemma 13.** *Let $0 < \epsilon < 1/3$ and let $\mathcal{I} = (G, s, C, w)$ be an instance of SGE, in which
824 $B \geq 5c/\epsilon^2 \cdot \ln n$. Let F be a set of edges acquired by some adaptive strategy. If $\mu(F) \geq (1+\epsilon) \cdot B$
825 then the probability that $C(F) \leq B$ is at most n^{-c} .*

826 **Proof.** In order to bound the probability, it is crucial to exploit the property of irrevocable
827 decisions, which forces the adaptive strategy to keep an item even if its size turns out to be
828 very large. Therefore, we would use the *martingale* framework. For an adaptive strategy let
829 F_t denote the set of the first t items chosen by the strategy. Note that no further items are
830 added to F_t once the cost of the edges in F_t exceeds the budget. Define $X_t = \sum_{e \in F_t} (c(e) - \mu_e)$.
831 It is easy to verify that X_t is a martingale, since $\mathbf{E}[X_{t+1} | X_t] = X_t$. Note that $X_0 = 0$ by
832 definition. Next, we bound the variance of

$$\mathbf{Var}[X_t | X_{t-1}] = \mathbf{Var}[(C(e) - \mu_e)] = \mathbf{E}[C(e)^2] - \mathbf{E}[C(e)]^2 \leq \mathbf{E}[C(e)^2] \leq \mathbf{E}[C(e)],$$

131:22 Stochastic Graph Exploration

833 where the last inequality holds since the cost of each edge is bounded by 1.

834 We now apply Theorem 21 with $X_t = C(F_t) - \mu(F_t)$, $M = 1$ and $\lambda = \epsilon B$. Observe that
 835 $\sum_{i=1}^t \mathbf{Var}[X_i | X_{i-1}] \leq (1 + \epsilon) \cdot B$. Hence, we have

$$\begin{aligned}
 836 \quad \Pr(C(F_t) \leq B) &\leq \Pr(C(F_t) \leq \mu(F_t) - \epsilon B) \leq \Pr(C(F_t) - \mu(F_t) \leq -\epsilon B) \\
 837 &\leq \exp\left(\frac{-\epsilon^2 \cdot B^2}{2(1 + \epsilon)B + \epsilon \cdot B/3}\right) \leq \exp\left(\frac{-\epsilon^2 \cdot B}{2(1 + \epsilon) + \epsilon/3}\right) \\
 838 &\leq \exp\left(\frac{-4c \log n}{2(1 + \epsilon) + \epsilon/3}\right) \leq n^{-c}. \\
 839
 \end{aligned}$$

840 ◀

841 ▶ **Lemma 22** (Chernoff bound). *Let X be the sum of binary independent random variables*
 842 *X_1, X_2, \dots, X_n where $X_i \in [0, 1]$. Then, for any $\delta > 1$, $\Pr(X > (1 + \delta) \cdot \mathbf{E}[X]) \leq \exp(-\delta \cdot$*
 843 *$\mathbf{E}[X]/3)$.*

844 ▶ **Lemma 14.** *Let $\mathcal{I} = (G, s, C, w)$ be an instance of SGE. For any set of edges F and any*
 845 *$\tilde{B} \geq 5c/\epsilon^2 \cdot \ln n$, if $\mu(F) = \tilde{B}$ then the probability that $C(F) \geq (1 + \epsilon)\tilde{B}$ is at most n^{-c} .*

Proof. We apply Chernoff bound by setting $X_i = C(e)$. We have $X = C(S)$ and

$$\Pr(C(F) > (1 + \epsilon) \cdot \tilde{B}) \leq \exp\left(\frac{-\epsilon \cdot \tilde{B}}{3}\right) \leq n^{-c}.$$

846 ◀

847 ▶ **Lemma 15.** *Let $\mathcal{I} = (G, s, C, w)$ be an instance of SGE, where $B \geq 5c/\epsilon^2 \ln n$, the*
 848 *maximum reward R satisfies $R \leq \epsilon n^{c-1}$, and the minimum reward is 1. Let \mathcal{I}_e be obtained*
 849 *from \mathcal{I} by replacing each edge cost with its expected value. Let $\text{OPT}_{\text{set}}^\epsilon$ be the optimal set*
 850 *strategy using budget $(1 + \epsilon)B$ for \mathcal{I}_e and OPT_{ad} be the optimal adaptive strategy using budget*
 851 *B for \mathcal{I} . Then, $(1 + \epsilon)r(\mathcal{I}, \text{OPT}_{\text{set}}^\epsilon, (1 + \epsilon)B) \geq r(\mathcal{I}, \text{OPT}_{\text{ad}}, B)$.*

852 **Proof.** First, we bound $r(\mathcal{I}, \text{OPT}_{\text{ad}}, B)$. This payoff is the sum of two terms: the expected
 853 payoff when the expected cost of the acquired edges is at most $(1 + \epsilon) \cdot B$ and the payoff when
 854 the expected cost of edges is greater than $(1 + \epsilon)B$. In the following we use $r(\mathcal{I}, S, B | \mathcal{E})$ to
 855 denote the expected payoff of strategy S , conditioned on the event \mathcal{E} .

856 Let F be the set of edges acquired by OPT_{ad} . Observe that $r(\mathcal{I}, \text{OPT}_{\text{ad}}, B | \mu(F) \leq$
 857 $(1 + \epsilon)B) \leq r(\mathcal{I}_e, \text{OPT}_{\text{set}}^\epsilon, (1 + \epsilon)B)$. This follows from the fact that the conditioning on
 858 $\mu(F) \leq (1 + \epsilon)B$ limits the possible sets of edges that the adaptive strategy can acquire to
 859 the sets of edges considered by the set strategy. We now obtain

$$\begin{aligned}
 860 \quad r(\mathcal{I}, \text{OPT}_{\text{ad}}, B) &\leq r(\mathcal{I}, \text{OPT}_{\text{ad}}, B | \mu(F) \leq (1 + \epsilon)B) + r(\mathcal{I}, \text{OPT}_{\text{ad}} | \mu(F) > (1 + \epsilon)B) \\
 861 &\leq r(\mathcal{I}, \text{OPT}_{\text{ad}}, B | \mu(F) \leq (1 + \epsilon)B) + n \cdot R \cdot n^{-c} \\
 862 &\leq r(\mathcal{I}_e, \text{OPT}_{\text{set}}^\epsilon, (1 + \epsilon)B) + \epsilon \\
 863 &\leq (1 + \epsilon)r(\mathcal{I}_e, \text{OPT}_{\text{set}}^\epsilon, (1 + \epsilon)B),
 \end{aligned}$$

864 where the second inequality follows from Lemma 13 and the last one from the fact that the
 865 set strategy obtains at least a reward of 1. ◀

866 ▶ **Theorem 12.** *Let $\mathcal{I} = (G, s, C, w)$ be an instance of SGE, where $C(e) = O(\frac{\epsilon^2}{\ln n})$ (for each*
 867 *edge e and some $0 < \epsilon = O(1)$), $R \leq \epsilon n^{O(1)}$, and the smallest reward is 1. Then, in polynomial*
 868 *time, we can compute a nonadaptive $(O(1), 1 + \epsilon)$ -approximate strategy for \mathcal{I} . Additionally,*
 869 *if G is a tree, then in time $O(n^3/\epsilon)$ we can compute a nonadaptive $(1 + \epsilon, 1 + \epsilon)$ -approximate*
 870 *strategy for \mathcal{I} .*

871 **Proof.** By combining Lemma 15 with Lemma 10 we obtain an algorithm which can compute
 872 a nonadaptive $(1 + \epsilon, (1 + \epsilon)^2)$ -approximate algorithm in $O(n^3/\epsilon)$ time. By tuning constants,
 873 the approximation can be improved to $(1 + \epsilon, 1 + \epsilon)$.

874 On the other hand, for general graphs we combine Lemma 15 with the 8-approximate
 875 algorithm for the max-prize problem (see Section 4.3) and obtain a $(8(1 + \epsilon), 1 + \epsilon)$ strategy,
 876 which by tuning constants can be improved to $(8 + \epsilon, 1 + \epsilon)$. ◀

877 ▶ **Lemma 16.** *Assume that for each edge e_i , $i \in [n]$ we have $c_i \in [0, 1]$. Then*

$$878 \quad P_n(3) \geq P_n(1) (1 - \ln(P_n(1))).$$

879 **Proof.** First we show that, for all $j \in [n]$ we have that:

$$880 \quad \Pr(C^n \leq 3 \mid C^{j-1} \leq 1, C^j > 1) \geq \Pr(C^n \leq 1 \mid C^j \leq 1). \quad (2)$$

881 The events $C^j \leq 2$ and $C_{j+1}^n \leq 1$ imply $C^n \leq 3$. The event $C^{j-1} \leq 1$ implies that $C^j \leq 2$
 882 because $c_j \leq 1$. For $i \neq j$, c_i and c_j are independent random variables, thus C^j and C_{j+1}^n
 883 are also independent. Using these observations we obtain

$$\begin{aligned} \Pr(C^n \leq 3 \mid C^{j-1} \leq 1, C^j > 1) &\geq \Pr(C^j \leq 2, C_{j+1}^n \leq 1 \mid C^{j-1} \leq 1, C^j > 1) \\ &= \Pr(C_{j+1}^n \leq 1 \mid C^{j-1} \leq 1, C^j > 1) \\ &= \Pr(C_{j+1}^n \leq 1) \\ &= \Pr(C_{j+1}^n \leq 1) \cdot \sum_{t \leq 1} \Pr(C^j = t \mid C^j \leq 1) \\ &= \sum_{t \leq 1} \Pr(C_{j+1}^n \leq 1) \cdot \Pr(C^j = t \mid C^j \leq 1) \\ &\geq \sum_{t \leq 1} \Pr(C_{j+1}^n \leq 1 - t) \cdot \Pr(C^j = t \mid C^j \leq 1) \\ &= \sum_{t \leq 1} \Pr(C_{j+1}^n \leq 1 - t \mid C^j \leq 1) \cdot \Pr(C^j = t \mid C^j \leq 1) \\ 884 &= \Pr(C^n \leq 1 \mid C^j \leq 1). \end{aligned}$$

885 To simplify the notation in the following, we define $N_1 = \Pr(C^1 \leq 1)$ and for $r = 2, \dots, n$,
 886 $N_r = \Pr(C^n \leq 1 \mid C^{r-1} \leq 1)$.

887 We will next make use of the following equalities.

$$888 \quad \Pr(C^n \leq 1) = \prod_{i=1}^n N_i,$$

$$889 \quad \Pr(C^n \leq 1 \mid C^j \leq 1) = \prod_{i=j+1}^n N_i,$$

$$890 \quad \Pr(C^n > 1, C^{r-1} \leq 1) = (1 - N_r) \cdot \prod_{i=1}^{r-1} N_i.$$

892 We partition the event $C^n \leq 3$ into disjoint events according to the first index i (if any)
 893 such that $C^i > 1$, we get

$$\begin{aligned}
\Pr(C^n \leq 3) &= \Pr(C^n \leq 3, C^n \leq 1) + \sum_{j=1}^n \Pr(C^n \leq 3, C^j > 1, C^{j-1} \leq 1) \\
&= \Pr(C^n \leq 1) + \sum_{j=1}^n \Pr(C^n \leq 3 \mid C^j > 1, C^{j-1} \leq 1) \cdot \Pr(C^j > 1, C^{j-1} \leq 1) \\
&= \Pr(C^n \leq 1) + \sum_{j=1}^n \left(\Pr(C^n \leq 3 \mid C^j > 1, C^{j-1} \leq 1) \cdot (1 - N_j) \cdot \prod_{i=1}^{j-1} N_i \right) \\
&\geq \Pr(C^n \leq 1) + \sum_{j=1}^n \left(\Pr(C^n \leq 1 \mid C^j \leq 1) \cdot (1 - N_j) \cdot \prod_{i=1}^{j-1} N_i \right) \\
&= \Pr(C^n \leq 1) + \sum_{j=1}^n \left(\prod_{i=j+1}^n N_i \cdot (1 - N_j) \cdot \prod_{i=1}^{j-1} N_i \right) \\
&= \Pr(C^n \leq 1) + \sum_{j=1}^n \frac{1 - N_j}{N_j} \cdot \left(\prod_{i=1}^n N_i \right) \\
&= \Pr(C^n \leq 1) + \sum_{j=1}^n \frac{1 - N_j}{N_j} \cdot \Pr(C^n \leq 1) = \Pr(C^n \leq 1) \cdot \left(1 + \sum_{j=1}^n \frac{1 - N_j}{N_j} \right) \\
&\geq \Pr(C^n \leq 1) \cdot \left(1 - \sum_{j=1}^n \ln N_j \right) = \Pr(C^n \leq 1) \cdot \left(1 - \ln \left(\prod_{j=1}^n N_j \right) \right) \\
&= \Pr(C^n \leq 1) \cdot (1 - \ln(\Pr(C^n \leq 1))),
\end{aligned}$$

894

895 where the first inequality follows from Equation (2) and the second from $e^y \geq y + 1$, which
896 for $y = -\ln x$ gives $\frac{1-x}{x} \geq -\ln x$. ◀

► **Corollary 17.**

$$897 \quad P_n(9) \geq P_n(1) \frac{(1 - \ln(P_n(1)))^2}{2}$$

898 **Proof.** Let for $i \in [n]$ define $\tilde{c}_i = c_i/3$, and define analogously \tilde{C}^n . Notice that $\tilde{c}_i \in [0, 1/3] \subset$
899 $[0, 1]$. We apply Lemma 16 twice (first for \tilde{C}^n and then for C^n and we obtain By rescaling
900 the costs by factor of $1/3$ (denoted with a tilde), we get that $\tilde{c}_i \in [0, 1/3]$, and therefore
901 $\tilde{c}_i \in [0, 1]$ and we apply Lemma 16 we have

$$\begin{aligned}
P_n(9) &= \Pr(C^n \leq 9) \\
&= \Pr(\tilde{C}^n \leq 3) \\
&\geq \Pr(\tilde{C}^n \leq 1) \cdot (1 - \ln(\Pr(\tilde{C}^n \leq 1))) \\
&\geq P_n(3) \cdot (1 - \ln(P_n(3))) \\
&\geq P_n(1) \cdot (1 - \ln(P_n(1))) \cdot (1 - \ln(P_n(1) \cdot (1 - \ln(P_n(1)))))) \\
&= P_n(1) \cdot (1 - \ln(P_n(1))) \cdot (1 - \ln(P_n(1)) - \ln(1 - \ln(P_n(1)))) \\
902 &\geq P_n(1) \cdot \frac{(1 - \ln(P_n(1)))^2}{2},
\end{aligned}$$

903 where the last inequality uses the fact that $x/2 \geq \ln(x)$, where we apply $x = 1 - \ln(P_n(1))$. ◀

► **Lemma 18.**

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

Proof. First note that $\Pr(C^j \leq 1)$ is a nonincreasing function of j . We partition the edge set into classes A_0, A_1, \dots, A_ℓ such that $j \in A_i$ if $e^{-i-1} < \Pr(C^j \leq B) \leq e^{-i}$. Let $w(A_i) = \sum_{j \in A_i} w(v_j)$ and $T(r) = \sum_{i=0}^r w(A_i)$. Then

$$r(\mathcal{I}, S_{ls}, B) \leq \sum_{i=1}^{\ell} e^{-i} w(A_i) \leq \sum_{i=1}^{\ell} e^{-i} T(i).$$

Let $k^* = \arg \max_k \{r(\mathcal{I}, S_k, 9B)\}$. For all i , let $a(i) = \max\{j : j \in A_i\}$, and $M = r(\mathcal{I}, S_{k^*}, 9B)$. Then

$$M = r(\mathcal{I}, S_{k^*}, 9B) \geq T(i) \cdot \Pr(C^{a(i)} \leq 9B) \geq T(i) \cdot \frac{(i+2)^2}{2e^{i+1}}.$$

Here, the first inequality follows from the fact that k^* corresponds to the maximum value. For the second inequality, note that $a(i) \in A_i$ and, hence, $\Pr(C^{a(i)} \leq B) \geq e^{-i-1}$. By Corollary 17, we obtain that $\Pr(C^{a(i)} \leq 9B) \geq e^{-i-1} \cdot (1 - \ln(e^{-i-1}))^2 / 2 = \frac{(i+2)^2}{2e^{i+1}}$.

Therefore, we have that $T(i) \leq \frac{2 \cdot M e^{i+1}}{(i+2)^2}$, and

$$\begin{aligned} r(\mathcal{I}, S_{ls}, B) &\leq \sum_{i=1}^{\ell} e^{-i} \cdot T(i) \leq \sum_{i=1}^{\ell} e^{-i} \frac{2 \cdot M e^{i+1}}{(i+2)^2} \leq 2Me \cdot \sum_{i=1}^{\ell} \frac{1}{(i+2)^2} \\ &\leq r(\mathcal{I}, S_{k^*}, 9B) / 0.46. \end{aligned}$$

◀

B MINESWEEPER on Trees

In this section we describe the optimal algorithm for MINESWEEPER problem, in the case when the input graph is a tree.

B.1 Scheduling with Tree-Like Precedence Constraints

We now introduce and solve a problem of n jobs with tree-like precedence constraints on a single machine. Some special cases of the problem we consider here have been solved before [2, 3, 18] and actually the main idea of the algorithm is the same as in the previous works.

Let A be a set of n jobs. An *ordering* is a sequence of length n consisting of distinct elements of A (a permutation of A). We denote by $Ord(A)$ the set of all orderings of A and by $Seq(A)$ we denote the set consisting of all sequences that contain distinct elements of A . Moreover, if X and Y are sequences we use $X \cdot Y$ to denote their concatenation.

The input to the scheduling problem we consider is a tuple (A, n, T, c) , where A is a set of n jobs, T is a rooted tree, whose vertex set is A , and $c : Ord(A) \rightarrow \mathbb{R}$ is the cost function.

Let $par_T(a_i)$ be the parent of a_i in T . We say that an ordering a_1, \dots, a_n is *valid* (w.r.t. T) iff a_1 is the root of T and for each $2 \leq i \leq n$ we have that $par_T(a_i) \in \{a_1, \dots, a_{i-1}\}$. The goal is to compute a valid ordering a_1, \dots, a_n such that the cost $c(a_1, \dots, a_n)$ is minimized.

We call every such sequence an *optimal ordering*.

131:26 Stochastic Graph Exploration

936 Let $S = W \cdot X \cdot Y \cdot Z$ be a valid ordering. We say that X and Y are *swappable* (with
 937 respect to S) iff $W \cdot Y \cdot X \cdot Z$ is a valid ordering. Observe that X and Y are swappable when
 938 for each $x \in X$ and $y \in Y$, x and y are not in ancestor–descendant (or descendant–ancestor)
 939 relation in T .

940 ► **Definition 23.** Consider a scheduling problem (A, n, T, c) . Let $W \cdot X \cdot Y \cdot Z \in \text{Ord}(A)$.
 941 We say that a function $u : \text{Seq}(A) \rightarrow \mathbb{R}$ is a utility function for (A, n, T, c) when

$$942 \quad c(W \cdot X \cdot Y \cdot Z) \leq c(W \cdot Y \cdot X \cdot Z) \Leftrightarrow u(X) \geq u(Y).$$

943 ► **Corollary 24.** Let S be an optimal ordering. Assume that $S = W \cdot X \cdot Y \cdot Z$, where X and
 944 Y are swappable. Then $u(X) \geq u(Y)$.

945 **Proof.** Assuming $u(X) < u(Y)$ immediately gives a contradiction to the assumption that S
 946 is optimal. ◀

947 In the remaining of this section we show an efficient algorithm, which, given (A, n, T, c)
 948 and a utility function, computes an optimal ordering. Note that the previous works only
 949 considered concrete cost functions (which, although not stated explicitly, did admit utility
 950 functions).

951 ► **Lemma 25.** Let a_1 be a job with a single child a_2 in T . Moreover, assume that $u(a_1) \leq$
 952 $u(a_2)$. Then, a_2 is scheduled immediately after a_1 in some optimal ordering.

953 **Proof.** Let S be an optimal ordering. Clearly, a_1 has to appear before a_2 , so S is of the
 954 form $S = X \cdot a_1 \cdot Y \cdot a_2 \cdot Z$ for some $X, Y, Z \in \text{Seq}(A)$. Since a_2 is the only child of a_1 , we
 955 know that a_1 and Y are swappable. Similarly, Y and a_2 are swappable.

956 Thus, the order $S' = X \cdot Y \cdot a_1 \cdot a_2 \cdot Z$ is valid. Moreover, by Corollary 24 applied to S
 957 we get $u(Y) \geq u(a_2)$, which implies $u(Y) \geq u(a_1)$. Thus, by Definition 23, $c(S') \leq c(S)$.

958 Hence, S' is an optimal ordering, in which a_2 is scheduled immediately after a_1 . ◀

959 Consider an algorithm that computes a valid ordering a_1, \dots, a_n greedily. It runs in n
 960 steps. In the i th step it selects a job a_i that is either the root of T or which has the maximum
 961 utility $u(a_i)$ among jobs whose parents were already added to the ordering. Note that the
 962 resulting ordering is not necessarily unique, as the algorithm may have to choose between
 963 jobs of the same utility. We call every ordering that may be produced by the algorithm
 964 a *greedy ordering* of A . Let $a \in A$ and let D_a be the set of jobs (vertices) that are *proper*
 965 descendants of a in T (hence, $a \notin D_a$). The greedy algorithm described above can also be
 966 used to order jobs in D_a , and we extend the definition of greedy orderings to such sets.

967 ► **Lemma 26.** Let $a \in A$ be a job. Assume that for each proper descendant b of a such that
 968 $\text{par}_T(b) \neq a$ we have $u(b) < u(\text{par}_T(b))$. Let S_a be some greedy ordering of the set D_a of
 969 proper descendants of a . Then, there exists an optimal ordering of A containing S_a as a
 970 subsequence.

971 **Proof.** Let k denote the number of descendants of a in T and let a'_1, \dots, a'_k be their greedy
 972 ordering. Moreover, let a_1, \dots, a_k be the ordering of all proper descendants of a in some
 973 optimal ordering S . Thus, $S = X_1 \cdot a_1 \cdot X_2 \cdot \dots \cdot X_k \cdot a_k \cdot X_{k+1}$ for some $X_1, \dots, X_{k+1} \in \text{Seq}(A)$.
 974 By definition, each descendant of a'_i is some other a'_j , which means that each a'_i is swappable
 975 with X_ℓ for $\ell \geq 2$. Therefore, $S' = X_1 \cdot a'_1 \cdot X_2 \cdot \dots \cdot X_k \cdot a'_k \cdot X_{k+1}$ is a valid ordering. To
 976 complete the proof, it remains to show that $c(S') = c(S)$.

977 First, consider $S = X_1 \cdot a_1 \cdot X_2 \cdot \dots \cdot X_k \cdot a_k \cdot X_{k+1}$. For $2 \leq i \leq k$ we have that X_i
 978 is swappable with a_i and for each $1 \leq i \leq k$, a_i is swappable with X_{i+1} . By applying
 979 Corollary 24 we conclude that $u(a_i) \geq u(a_{i+1})$ for $1 \leq i < k$.

980 From the assumption that $u(b) < u(\text{par}_T(b))$, combined with the algorithm for computing
 981 greedy orderings, we also have $u(a'_i) \geq u(a'_{i+1})$. Hence, in both S and S' the jobs are sorted
 982 in nonincreasing order of utilities. This means that S' can be obtained from S by permuting
 983 jobs that have equal utilities.

984 Recall that $S' = X_1 \cdot a'_1 \cdot X_2 \cdot \dots \cdot X_k \cdot a'_k \cdot X_{k+1}$. To complete the proof, we show that
 985 if $u(a'_i) = u(a'_{i+1})$, then the two jobs can be swapped and the resulting ordering is still
 986 valid and has the same cost. The assumption $u(a'_i) = u(a'_{i+1})$ combined with the fact that
 987 $u(a'_i) \geq u(X_{i+1}) \geq u(a'_{i+1})$, yields $u(a'_i) = u(X_{i+1}) = u(a'_{i+1})$.

988 We now swap a'_i with a'_{i+1} by performing 3 swaps of adjacent elements a'_i , X_{i+1} , and a'_{i+1} .
 989 Since the utilities of each these elements are equal, swapping does not influence the cost of
 990 the ordering. We only have to show that each time we swap a swappable pair. Clearly, both
 991 a'_i and a'_{i+1} can be swapped with X_{i+1} . In addition a'_i and a'_{i+1} can be swapped with each
 992 other. The fact that $u(a'_i) = u(a'_{i+1})$ combined with the assumption that $u(b) < u(\text{par}_T(b))$
 993 implies that neither of these elements is an ancestor of the other. The lemma follows. ◀

994 Lemmas 26 and 25 motivate two reductions, which we can apply to our problem, without
 995 changing the cost of the optimal solution.

996 We first describe a *merging reduction*, based on Lemma 26. Let $a \in A$ be a job satisfying
 997 the assumptions of Lemma 26. The merging reduction consists in replacing the subtree of a
 998 by a path containing the jobs ordered in a greedy way.

999 ▶ **Corollary 27.** *Consider a scheduling problem (A, n, T, c) . Let $a \in A$ be a job satisfying the*
 1000 *assumptions of Lemma 26 and T_a be the tree obtained from T by performing a merging step*
 1001 *on a . Then, every optimal ordering for (A, n, T_a, c) is an optimal ordering for (A, n, T, c) .*

1002 **Proof.** Observe that the precedence constraints defined by T_a can only be stricter than the
 1003 ones given by T . Thus, every optimal ordering for (A, n, T_a, c) is valid for (A, n, T, c) . From
 1004 Lemma 26 it follows that the costs of optimal orderings in both problems are equal. ◀

1005 Now consider two jobs a_1 and a_2 satisfying the assumptions of Lemma 25. A *contracting*
 1006 *reduction* contracts two jobs a_1 and a_2 into one job a_3 of utility $u(a_1 \cdot a_2)$. More generally,
 1007 we define $u(X \cdot a_3 \cdot Y) := u(X \cdot a_1 \cdot a_2 \cdot Y)$. By Lemma 25 and Corollary 27 both reductions
 1008 do not change the cost of the optimal ordering. In order to compute the optimal ordering,
 1009 we apply the reductions repeatedly.

1010 ▶ **Lemma 28.** *Consider a job $a \in A$ and the subtree T_a of T rooted in a . If the contracting*
 1011 *reduction cannot be applied within T_a , then either T_a imposes a linear order of jobs or the*
 1012 *merge reduction can be applied within T_a .*

1013 **Proof.** Consider the lowest node x of T_a that has at least two children. Note that such
 1014 node exists if T_a is not a single path (which would impose a linear order on the jobs).
 1015 From the choice of x we have that each subtree rooted at x is a path. Moreover, since the
 1016 contracting reduction cannot be applied, we immediately have that the merge reduction can
 1017 be applied. ◀

1018 It follows that, by applying the reductions, we obtain a tree T' that imposes a linear order
 1019 of jobs. It remains to show that the final ordering can be computed efficiently. To that end,
 1020 we need to assume that the utility function can be computed efficiently. In our algorithm,

1021 we compute the utility of sequences consisting of a single job and merge jobs during a merge
 1022 reduction. When this happens, two jobs a_1 and a_2 are replaced with one, which is equivalent
 1023 to executing a_1 immediately followed by a_2 . Jobs a_1 and a_2 are discarded, which implies
 1024 that this step can be executed at most $n - 1$ times, where n is the initial number of jobs.
 1025 We say that a utility function can be *maintained efficiently* if we can compute the utility of
 1026 each individual job in $O(1)$ time, including jobs that are created during a merge reduction.

1027 ► **Theorem 29.** *Let (A, n, T, c) be an instance of the scheduling problem. Assume that there*
 1028 *exists a utility function u for (A, n, T, c) that can be maintained efficiently. Then, the optimal*
 1029 *ordering for (A, n, T, c) can be computed in $O(n \log n)$ time.*

1030 **Proof.** The algorithm is a recursive function that, given a node of T , replaces the subtree
 1031 rooted at T with a path (subtree imposing a linear order of jobs) without affecting the cost
 1032 of the optimal ordering. It turns out that we can reuse the existing algorithm from [2] that
 1033 proceeds analogously (the only difference is that it uses some particular utility function).
 1034 The algorithm is a simple implementation of the recursive function, using leftist trees. As
 1035 shown in [2] it runs in $O(n \log n)$ time. ◀

1036 B.2 Optimal Algorithm for MINESWEEPER on Trees

1037 ► **Theorem 8.** *Consider the instance $\mathcal{I} = (T, s, p, w)$ of the minesweeper problem, where*
 1038 *T is a tree. The optimal strategy, OPT_{MS} , for MINESWEEPER on T can be computed in*
 1039 *$O(n \log n)$ time, where n is the number of vertices of T .*

1040 **Proof.** Observe that the minesweeper problem can be viewed as a job scheduling problem with
 1041 precedence constraints, where the jobs to be scheduled are nodes of T . In the minesweeper
 1042 problem, the goal is to produce an ordering of edges, but in the case of trees this is the same
 1043 as ordering the vertices. Namely, having ordered the vertices, we can produce the ordering of
 1044 edges by taking the edges connecting each vertex to the parent (and ignoring the first vertex
 1045 in the sequence).

1046 Let $\{a_i; 1 \leq i \leq n\}$ be the set of vertices of T , $a_1 = s$ being the root of T . To simplify
 1047 notation, in the following, for $1 \leq i \leq n$ we denote by $p(a_i)$ the probability that the edge
 1048 $(par_T(a_i), a_i)$ materializes (we also set $p(a_1) = 1$). An ordering $a_1 \cdot a_2 \cdot \dots \cdot a_n$ of vertices of
 1049 T defines a (nonadaptive) strategy for the minesweeper problem.

1050 For a sequence of vertices $b_1 \cdot \dots \cdot b_k$ in T define

$$1051 \quad W(b_1, \dots, b_k) = \sum_{i=1}^k \left(w(b_i) \prod_{j=1}^i p(b_j) \right).$$

1052 Then notice that $W(a_1 \cdot \dots \cdot a_n)$ is the expected payoff of the nonadaptive strategy that
 1053 probes the edges according to order $a_1 \cdot \dots \cdot a_n$: Executing this strategy we collect a weight
 1054 $w(a_i)$ iff the process does not stop before reaching a_i , that is, iff all the edges $(par_T(a_j), a_j)$
 1055 for $j = 2, \dots, i$ materialize.

1056 To compute the optimal ordering of vertices of T , we use Theorem 29. To apply it, it
 1057 suffices to show that the job-scheduling problem we obtain admits a utility function and that
 1058 the function can be maintained efficiently.

1059 Consider an ordering $S = X_1 \cdot X_2 \cdot X_3 \cdot X_4$ and let P_i be the probability that all parent
 1060 edges of vertices of X_i materialize. That is, for $X_i = b_1, \dots, b_k$, $P_i = \prod_{j=1}^k p(b_j)$.

1061 Notice that we can write

$$1062 \quad W(S) = W(X_1 \cdot X_2 \cdot X_3 \cdot X_4) = W(X_1) + P_1 W(X_2) + P_1 P_2 W(X_3) + P_1 P_2 P_3 W(X_4).$$

1063 Furthermore, we have that the following inequalities are equivalent:

$$\begin{aligned}
 1064 \quad & W(X_1 \cdot X_2 \cdot X_3 \cdot X_4) \leq W(X_1 \cdot X_3 \cdot X_2 \cdot X_4) \\
 1065 \quad & P_1 W(X_2) + P_1 P_2 W(X_3) \leq P_1 W(X_3) + P_1 P_3 W(X_2) \\
 1066 \quad & W(X_2) + P_2 W(X_3) \leq W(X_3) + P_3 W(X_2) \\
 1067 \quad & W(X_2)(1 - P_3) \leq W(X_3)(1 - P_2) \\
 1068 \quad & \frac{W(X_2)}{1 - P_2} \leq \frac{W(X_3)}{1 - P_3} \\
 1069 \quad &
 \end{aligned}$$

1070 Hence, $u(X_i) = W(X_i)/(1 - P_i)$ is a utility function for this problem. Note that if $P_i = 1$,
 1071 then we can set $u(X_i)$ to be equal to a fixed value M that is larger than any other utility
 1072 (this value can be the same, regardless of X_i). It is easy to see that this utility function can
 1073 be maintained efficiently by storing P_i and $W(X_i)$ for each job X_i that we obtain as a result
 1074 of merging jobs. Thus, we can now apply Theorem 29 to complete the proof. ◀