



HAL
open science

Searching a Tree with Permanently Noisy Advice

Lucas Boczkowski, Amos Korman, Yoav Rodeh

► **To cite this version:**

Lucas Boczkowski, Amos Korman, Yoav Rodeh. Searching a Tree with Permanently Noisy Advice. ESA 2018 - 26th Annual European Symposium on Algorithms, Aug 2018, Helsinki, Finland. pp.1-32, 10.4230/LIPIcs.ESA.2018.54 . hal-01958133v1

HAL Id: hal-01958133

<https://hal.science/hal-01958133v1>

Submitted on 17 Dec 2018 (v1), last revised 20 Jul 2021 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

1 Searching a Tree with Permanently Noisy Advice

2 **Lucas Boczkowski**

3 CNRS, IRIF, Univ. Paris Diderot, Paris, France

4 lucas.boczkowski@irif.fr

5 **Amos Korman**

6 CNRS, IRIF, Univ. Paris Diderot, Paris, France

7 amos.korman@irif.fr

8 **Yoav Rodeh**

9 Dep. of Physics of Complex Systems. Weizmann Institute, Rehovot, Israel

10 yoav.rodeh@gmail.com

11 — Abstract —

12 We consider a search problem on trees using unreliable guiding instructions. Specifically, an
13 agent starts a search at the root of a tree aiming to find a treasure hidden at one of the nodes
14 by an adversary. Each visited node holds information, called *advice*, regarding the most prom-
15 ising neighbor to continue the search. However, the memory holding this information may be
16 unreliable. Modeling this scenario, we focus on a probabilistic setting. That is, the advice at a
17 node is a pointer to one of its neighbors. With probability q each node is *faulty*, independently
18 of other nodes, in which case its advice points at an arbitrary neighbor, chosen uniformly at
19 random. Otherwise, the node is *sound* and points at the correct neighbor. Crucially, the advice
20 is *permanent*, in the sense that querying a node several times would yield the same answer. We
21 evaluate efficiency by two measures: The *move complexity* denotes the expected number of edge
22 traversals, and the *query complexity* denotes the expected number of queries.

23 Let Δ denote the maximal degree. Roughly speaking, the main message of this paper is that
24 a phase transition occurs when the *noise parameter* q is roughly $1/\sqrt{\Delta}$. More precisely, we prove
25 that above the threshold, every search algorithm has query complexity (and move complexity)
26 which is both exponential in the depth d of the treasure and polynomial in the number of nodes n .
27 Conversely, below the threshold, there exists an algorithm with move complexity $\mathcal{O}(d\sqrt{\Delta})$, and
28 an algorithm with query complexity $\mathcal{O}(\sqrt{\Delta} \log \Delta \log^2 n)$. Moreover, for the case of regular trees,
29 we obtain an algorithm with query complexity $\mathcal{O}(\sqrt{\Delta} \log n \log \log n)$. For q that is below but
30 close to the threshold, the bound for the move complexity is tight, and the bounds for the query
31 complexity are not far from the lower bound of $\Omega(\sqrt{\Delta} \log_{\Delta} n)$.

32 In addition, we also consider a *semi-adversarial* variant, in which faulty nodes are still chosen
33 at random, but an adversary chooses (beforehand) the advice of such nodes. For this variant,
34 the threshold for efficient moving algorithms happens when the noise parameter is roughly $1/\Delta$.
35 In fact, above this threshold a simple protocol that follows each advice with a fixed probability
36 already achieves optimal move complexity.

37 **2012 ACM Subject Classification** Theory of Computation → Probabilistic Computation, Data-
38 base Theory, Theory of Randomized Search Heuristics

39 **Keywords and phrases** Data structures - Graph search - Average Case Analysis - Permanents
40 Faults

41 **Digital Object Identifier** 10.4230/LIPIcs.ESA.2018.54

42 **Funding** This work has received funding from the European Research Council (ERC) under
43 the European Union's Horizon 2020 research and innovation programme (grant agreement No
44 648032).



© Lucas Boczkowski, Amos Korman and Yoav Rodeh;
licensed under Creative Commons License CC-BY

26th Annual European Symposium on Algorithms (ESA 2018).

Editors: Yossi Azar, Hannah Bast, and Grzegorz Herman; Article No. 54; pp. 54:1–54:32

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

45 **1** Introduction

46 This paper considers a basic search problem on trees, in which the goal is to find a treasure
 47 that is placed at one of the nodes by an adversary. Each node of the tree holds information,
 48 called *advice*, regarding which of its neighbors is closer to the treasure, and the search may
 49 consult the advice at some nodes in order to accelerate the search. Crucially, we assume that
 50 advice at nodes may be faulty with some probability. Many works consider noisy queries
 51 in the context of search, but it is typically assumed that queries can be resampled (see
 52 e.g., [11, 18, 4, 10]). In contrast, we assume that each location is associated with a single
 53 *permanent* advice. That is, faults are in the physical memory associated with the node, and
 54 hence querying the node again would yield the same answer. This difference is dramatic,
 55 as the search under our model does not allow for simple amplification procedures (similar
 56 to [6] albeit in the context of sorting). Searching in contexts of permanently faulty nodes
 57 has been studied in a number of works [7, 12, 15, 16, 17], but only assuming that the faulty
 58 nodes are chosen by an adversary. The difference between such worst case scenarios and the
 59 probabilistic version studied here is again significant, both in terms of results and techniques
 60 (see more details in Section 1.3).

61 1.1 The Noisy Advice Model

62 We start with some notation. Further notations are given in Section 1.4. Let T be an n -node
 63 tree¹ rooted at some arbitrary node σ . We consider an agent that is initially located at the
 64 root σ of T , aiming to find a node τ , called the *treasure*, which is chosen by an adversary.
 65 The distance $d(u, v)$ is the number of edges on the path between u and v . The depth of a
 66 node u is $d(u) = d(\sigma, u)$. Let $d = d(\tau)$ denote the depth of τ , and let the depth D of the
 67 tree be the maximal depth of a node. Finally, let Δ_u denote the degree of node u and let Δ
 68 denote the maximal degree in the tree.

69 Each node $u \neq \tau$ is assumed to be provided with an *advice*, termed $\text{adv}(u)$, which
 70 provides information regarding the direction of the treasure. Specifically, $\text{adv}(u)$ is a pointer
 71 to one of u 's neighbors. It is called *correct* if the pointed neighbor is one step closer to the
 72 treasure than u is. Each vertex $u \neq \tau$ is *faulty* with probability q_u (the meaning of being
 73 faulty will soon be explained). Otherwise, u is considered *sound*, in which case its advice
 74 is correct. We call q_u the *noise parameter* of u , and define the *general noise parameter* as
 75 $q = \max\{q_u \mid u \in T\}$.

76 We consider two models for faulty nodes. The main model assumes that the advice at
 77 a faulty node points to one of its neighbors chosen uniformly at random (and so possibly
 78 pointing at the correct one). We also consider an adversarial variant, called the *semi-*
 79 *adversarial model*, where this neighbor is chosen by an oblivious adversary. That is, an
 80 adversary specifies for each node what advice it would have assuming it is faulty. Then,
 81 faulty nodes are still chosen randomly as in the main model, but their advice is specified by
 82 the adversary.

83 The agent can move by traversing edges of the tree. At any time, the agent can query its
 84 hosting node in order to “see” the corresponding advice and to detect whether the treasure
 85 is present there. The protocol terminates when the agent queries the treasure. We evaluate
 86 a search algorithm A by two measures: The move complexity, termed $\mathcal{M}(A)$, is the expected

¹ We present the model for trees, but it should be clear that it can be similarly defined for arbitrary graphs (see also Section 5).

number of edge traversals, and the query complexity, termed $\mathcal{Q}(A)$, is the expected number of queries². Expectation is taken over both the randomness involved in sampling advice and the possible probabilistic choices made by A . We note that when considering walking algorithms, we assume that the agent does not know the structure of the tree in advance, and discovers it as it moves. Conversely, when focusing on minimizing the query complexity only, we assume that the tree structure is known to the algorithm.

The noise parameters $(q_u)_{u \in T}$ govern the accuracy of the environment. On the one extreme, if $q_u = 0$ for all nodes, then advice is always correct. This case allows to find the treasure in d moves, by simply following each encountered advice. Alternatively, it also allows to find the treasure using $\mathcal{O}(\log n)$ queries, by performing a separator based search. On the other extreme, if $q_u = 1$ for all nodes, then advice is essentially meaningless, and the search cannot be expected to be efficient. An intriguing question is therefore to identify the largest value of q that allows for efficient search.

The reader wishing to increase its familiarity with the model and its algorithmic challenges may read Appendix A, where a natural greedy algorithm is analyzed, and shown to be inefficient.

1.2 Our Results

Consider the noisy advice model on trees with maximum degree Δ and depth D . Roughly speaking, we show that $1/\sqrt{\Delta}$ is the threshold for the noise parameter q , in order to obtain search algorithms with low expected complexities.

The proof that there is no algorithm with low expected complexities when the noise exceeds $1/\sqrt{\Delta}$ is rather simple, and in fact, holds even if the algorithm has access to the advice of all internal nodes. Intuitively, the argument is as follows (the formal proof appears in Section 4.1). Consider a complete Δ -ary tree of depth D and assume that the treasure τ is placed at a leaf. The first observation (Lemma 10) is that the expected number of leaves having more advice point to them than to τ is a lower bound on the query complexity. The next observation is that there are roughly Δ^D leaves whose distance from τ is $2D$. For each of those leaves u , the probability that more advice points towards it than towards τ can be approximated by the probability that all nodes on path connecting u and τ are faulty. As this latter probability is q^{2D} , the expected number of leaves that have more pointers leading to them is roughly $\Delta^D q^{2D}$, which explodes when $q \gg 1/\sqrt{\Delta}$. This essentially establishes the lower bound for the noise regime.

The main technical difficulties we had to face appeared when we aimed to show that the $1/\sqrt{\Delta}$ lower bound is, in fact, tight, and moreover, that there exist extremely efficient algorithms when the noise is above the threshold. In this regard, we note two technical contributions. The first appears in the construction of the moving algorithm A_{walk} . Even though the algorithm should be designed to quickly find an adversarially placed treasure, it is in fact based on a Bayesian approach. The challenging part is identifying the correct prior. Constructing algorithms that ensure worst-case guarantees through a Bayesian approach was done in [4] which studies a closely related, yet much simpler problem of search on the line. Apart from [4] we are unaware of other works that follow this approach. The second technical contribution corresponds to the query setting, where we mimic the resampling of advice at separator nodes, by locally applying the moving algorithm.

² The success probability after a fixed number of rounds is another quantity of interest. It is left for future work.

130 **1.2.1 Upper Bounds**

131 In Section 2, we present a walking algorithm that is optimal up to a constant factor for
 132 the regime of noise below the threshold. Furthermore, this algorithm does not require prior
 133 knowledge of either the tree's structure, or the values of Δ , q , d , or n .

134 Using this walking algorithm, we derive two query algorithms (in Section 3). The first
 135 is optimal up to a factor of $\mathcal{O}(\log^2(\Delta) \log n)$ and the second is restricted to regular trees,
 136 but is optimal up to a factor of $\mathcal{O}(\log(\Delta) \log \log n)$. Note that the query algorithms use the
 137 knowledge of the tree structure, as well as bounds on the regime of noise.

138 Before stating our theorems, we need the following definition.

► **Definition 1.** Condition (\star) holds with parameter $0 < \varepsilon < 1$ if for every node v , we have

$$q_v < \frac{1 - \varepsilon - \Delta_v^{-\frac{1}{4}}}{\sqrt{\Delta_v} + \Delta_v^{\frac{1}{4}}}.$$

139 Note that since $\Delta_v \geq 2$, the condition is always satisfiable when taking a small enough ε . In
 140 the following theorems the \mathcal{O} notation hides only a polynomial a polynomial term in $1/\varepsilon$.

141 Note, all our algorithms are deterministic, hence, expectation is taken with respect only
 142 to the sampling of the advice.

143 ► **Theorem 2.** *There exists a deterministic walking algorithm A_{walk} such that for any con-*
 144 *stant $\varepsilon > 0$, if Condition (\star) holds with parameter ε then $\mathcal{M}(A_{walk}) = \mathcal{O}(\sqrt{\Delta}d)$.*

145 ► **Theorem 3. 1.** *For any $\varepsilon > 0$, there exists a deterministic query algorithm A_{sep} such*
 146 *that if Condition (\star) holds with parameter ε then the query complexity is $\mathcal{Q}(A_{sep}) =$*
 147 *$\mathcal{O}(\sqrt{\Delta} \log \Delta \cdot \log^2 n)$.*

148 **2.** *Assume that $q < c/\sqrt{\Delta}$ for a small enough constant $c > 0$. Then there exists a determ-*
 149 *inistic query algorithm $A_{2-layers}$ such that, restricted to (not necessarily complete) Δ -ary*
 150 *trees, $\mathcal{Q}(A_{2-layers}) = \mathcal{O}(\sqrt{\Delta} \log n \cdot \log \log n)$.*

151 **1.2.2 Lower Bounds**

152 We establish (in Section 4 and Appendix G) the following lower bounds.

153 ► **Theorem 4.** *The following holds for any randomized algorithm A and any integer $\Delta \geq 3$.*

154 **1. Exponential complexity above the threshold.**

155 *Consider a complete Δ -ary tree. For every constant $\varepsilon > 0$, if $q \geq \frac{1+\varepsilon}{\sqrt{\Delta-1}} \cdot (1 + \frac{1}{\Delta-1})$, then*
 156 *both $\mathcal{Q}(A)$ and $\mathcal{M}(A)$ are exponential in D .*

157 **2. Lower bounds for any q .**

158 (a) *Consider a complete Δ -ary tree. Then $\mathcal{Q}(A) = \Omega(q\Delta \log_{\Delta} n)$.*

159 (b) *For any integer d , there is a tree with at most $d\Delta$ nodes, and a placement of the*
 160 *treasure at depth d , such that $\mathcal{M}(A) = \Omega(dq\Delta)$.*

161 Observe that taken together, Theorems 2,4,3 and Condition (\star) imply that for any $\varepsilon > 0$ and
 162 large enough Δ , efficient search can be achieved if $q < (1-\varepsilon)/\sqrt{\Delta}$ but not if $q > (1+\varepsilon)/\sqrt{\Delta}$.

163 **1.2.3 Memory-less Algorithms**

164 Query algorithms assume the knowledge of the tree and hence cannot avoid memory com-
 165 plexity which is linear in n . In contrast, our walking algorithm A_{walk} uses memory that is
 166 composed of advice accumulated during the walk, and hence remains low, in expectation.

167 Finally, we analyse the performance of simple memoryless algorithms called *probabilistic*
 168 *following*, suggested in [14]. At every step, the algorithm follows the advice at the current
 169 vertex with some fixed probability λ , and performs a random walk step otherwise. It turns
 170 out that such algorithms can perform well, but only in a very limited regime of noise.
 171 Specifically, we prove:

172 ► **Theorem 5.** *There exist positive constants c_1 , c_2 and c_3 such that the following holds. If*
 173 *for every vertex u , $q_u < c_1/\Delta_u$ then there exists a probabilistic following algorithm that finds*
 174 *the treasure in less than c_2d expected steps. On the other hand, if $q > c_3/\Delta$ then for any*
 175 *probabilistic following strategy the move complexity on a complete Δ -ary tree is exponential*
 176 *in the depth of the tree.*

177 Since this algorithm is randomized, expectation is taken over both the randomness involved
 178 in sampling advice and the possible probabilistic choices made by the algorithm.

179 Interestingly, when $q_u < c_1/\Delta_u$ for all vertices, this algorithm works even in a semi-
 180 adversarial model. In fact, it turns out that in the semi-adversarial model, probabilistic
 181 following algorithms are the best possible, as the threshold for efficient search, with respect
 182 to any algorithm, is roughly $1/\Delta$. Due to lack of space these results are discussed and proved
 183 in Appendix H.

184 1.3 Related Work

185 In computer science, search algorithms have been the focus of numerous works. Due to their
 186 importance, trees are particularly popular structures to investigate search, see e.g., [19, 3,
 187 21, 20]. Within the literature on search, many works considered noisy queries [11, 18, 10],
 188 however, it was typically assumed that noise can be *resampled* at every query. As mentioned,
 189 dealing with permanent faults incurs challenges that are fundamentally different from those
 190 that arise when allowing queries to be resampled. To illustrate this difference, consider the
 191 simple example of a star graph and a constant $q < 1$. Straightforward amplification can
 192 detect the target in $\mathcal{O}(1)$ expected number of queries. In contrast, in our model, it can be
 193 easily seen that $\Omega(n)$ is a lower bound for both the move and the query complexities, for
 194 any constant noise parameter.

195 A search problem on graphs in which the set of nodes with misleading advice is chosen
 196 by an adversary was studied in [15, 16, 17], as part of the more general framework of the *liar*
 197 *models* [1, 2, 5, 8, 22]. Data structures with adversarial memory faults have been investigated
 198 in the so called faulty-memory RAM model introduced in [13]. In particular, data structures
 199 supporting the same operations as search trees with adversarial memory faults were studied
 200 in [12, 7]. Interestingly, the data structures developed in [7] can cope with up to $O(\log n)$
 201 faults, happening at any time during the execution of the algorithm, while maintaining
 202 optimal space and time complexity. All these worst case models are, however, significantly
 203 different from the randomized one we consider, both in terms of techniques and results. The
 204 subject of queries with probabilistic memory faults, as the ones we study here, has been
 205 explicitly studied in the context of sorting [6].

206 The noisy advice model considered in this paper actually originated in the recent biolo-
 207 gically centered work [14], aiming to abstract navigation relying on guiding instructions in
 208 the context of collaborative transport by ants. There, a group of ants carry a large load of
 209 food aiming to transport it to their nest, while basing their navigation on unreliable advice
 210 given by pheromones that are laid on the terrain. In that work, the authors modelled ant
 211 navigation as a probabilistic following algorithm, and noticed that an execution of such an
 212 algorithm can be viewed as an instance of Random Walks in Random Environment (RWRE)

213 [23, 9]. Relying on results from this subfield of probability theory, the authors showed that
 214 when tuned properly, such algorithms enjoy linear move complexity on grids, provided that
 215 the bias towards the correct direction is sufficiently high.

216 1.4 Notations

217 Denote $p = 1 - q$, and for a node u , $p_u = 1 - q_u$. For two nodes u, v , let $\langle u, v \rangle$ denote the
 218 simple path connecting them, excluding the end nodes, and let $[u, v] = \langle u, v \rangle \cup \{u\}$ and
 219 $[u, v] = [u, v] \cup \{v\}$. For a node u , let $T(u)$ be the subtree rooted at u . We denote by
 220 $\overrightarrow{\text{adv}}(u)$ (resp. $\overleftarrow{\text{adv}}(u)$) the set of nodes whose advice points towards (resp. away from) u .
 221 By convention $u \notin \overrightarrow{\text{adv}}(u) \cup \overleftarrow{\text{adv}}(u)$. Unless stated otherwise, \log is the natural logarithm.

222 1.5 Organization

223 In Section 2 we present our optimal walking algorithm. Section 3 presents our query al-
 224 gorithms, while most of the details regarding the more elaborated algorithm on regular trees
 225 are deferred to Appendix E. In Section 4 we show the lower bounds for both the move
 226 and query complexities. In Section 5, we give a list of open problems. Theorem 5 and the
 227 threshold of $\Theta(1/\Delta)$ that applies to the semi-adversarial setting are proved in Appendix H.

228 2 Optimal Walking Algorithm

229 In this section we prove Theorem 2. At a very high level, at any given time, the walking
 230 algorithm processes the advice seen so far, identifies a promising node to continue from on
 231 the border of the already discovered connected component, moves to that node, and explores
 232 one of its neighbors.

233 2.1 Algorithm Design following a Greedy Bayesian Approach

234 In our setting the treasure is placed by an adversary. However, we can still study algorithms
 235 induced by assuming that it is placed in one of the vertices according to some known distri-
 236 bution and see how they measure up in our worst case setting. As mentioned, this approach
 237 is similar to [4], which studies the closely related, yet much simpler problem of search on
 238 the line. Of course, the success of this scheme highly depends on the choice of the prior
 239 distribution we take.

240 To make our life easier, let us first assume that the structure of the tree is known to the
 241 algorithm. Also, we assume the treasure is placed in one of the leaves of the tree according
 242 to some known distribution θ , and denote by adv the advice on the nodes we have already
 243 visited. Aiming to find the treasure as fast as possible, a possible greedy algorithm explores
 244 the vertex that, given the advice seen so far, has the highest probability of having the
 245 treasure in its subtree.

246 We extend the definition of θ to internal nodes by defining $\theta(u)$ to be the sum of $\theta(w)$
 247 over all leaves w of $T(u)$. Given some u that was not visited yet, and given the previously
 248 seen advice adv , the probability of the treasure being in u 's subtree $T(u)$, is:

$$\begin{aligned}
 249 \quad \mathbb{P}(\tau \in T(u) \mid \text{adv}) &= \frac{\mathbb{P}(\tau \in T(u))}{\mathbb{P}(\text{adv})} \mathbb{P}(\text{adv} \mid \tau \in T(u)) \\
 250 &= \frac{\theta(u)}{\mathbb{P}(\text{adv})} \prod_{w \in \overrightarrow{\text{adv}}(u)} \left(p_w + \frac{q_w}{\Delta_w} \right) \prod_{w \in \overleftarrow{\text{adv}}(u)} \frac{q_w}{\Delta_w}. \\
 251
 \end{aligned}$$

The last factor is q_w/Δ_w because it is the probability that the advice at w points exactly the way it does in adv , and not only away from τ . Note that the advice seen so far is never for vertices in $T(u)$ as we consider a walking algorithm, and u has not been visited yet. Therefore, if $\tau \in T(u)$ then correct advice in adv points to u . We ignore the term $p_w + q_w/\Delta_w$ as it is normally quite close to 1, and applying a log we can approximate the relative strength of a node by:

$$\log(\theta(u)) + \sum_{w \in \overleftarrow{\text{adv}}(u)} \log\left(\frac{q_w}{\Delta_w}\right).$$

We do not want to assume that our algorithm knows q_w , but we do assume that in the worst scenario $q_w \sim 1/\sqrt{\Delta_w}$. Assigning this value and rescaling we finally define:

$$\text{score}(u) = \frac{2}{3} \log(\theta(u)) - \sum_{w \in \overleftarrow{\text{adv}}(u)} \log(\Delta_w).$$

When comparing two specific vertices u and v , $\text{score}(u) > \text{score}(v)$ iff:

$$\sum_{w \in \langle u, v \rangle \cap \overrightarrow{\text{adv}}(u)} \log(\Delta_w) - \sum_{w \in \langle u, v \rangle \cap \overrightarrow{\text{adv}}(v)} \log(\Delta_w) > \frac{2}{3} \log\left(\frac{\theta(v)}{\theta(u)}\right).$$

252 This is because any advice that is not on the path between u and v contributes the same to
 253 both sides, as well as advice on vertices on the path that point sideways, and not towards u
 254 or v ³. Since we use this score to compare two vertices that are neighbors of already explored
 255 vertices, and our algorithm is a walking algorithm, then we will always have all the advice
 256 on this path. In particular, the answer to whether $\text{score}(u) > \text{score}(v)$, does not depend
 257 on the specific choices of the algorithm, and does not change throughout the execution of
 258 the algorithm, even though the scores themselves do change. The comparison depends only
 259 on the advice given by the environment.

Let us try and justify the score criterion at an intuitive level. Consider the case of a complete Δ -ary tree, with θ being the uniform distribution on the leaves⁴. Here $\text{score}(u) > \text{score}(v)$ if (cheating a little by thinking of $\log(\Delta)$ and $\log(\Delta - 1)$ as equal):

$$|\overrightarrow{\text{adv}}(u) \cap \langle u, v \rangle| - |\overrightarrow{\text{adv}}(v) \cap \langle u, v \rangle| > \frac{2}{3} (d(u) - d(v)).$$

260 If, for example, we consider two vertices $u, v \in T$ at the same depth, then $\text{score}(u) > \text{score}(v)$
 261 if there is more advice pointing towards u than towards v . If the vertices have different
 262 depths, then the one closer to the root has some advantage, but it can still be beaten.

263 For general trees, perhaps the most natural θ to take is the uniform distribution on all
 264 nodes (or just on all leaves - this choice is actually similar). It is also a generalization of the
 265 example above. Unfortunately, however, while this works well on the complete Δ -ary tree,
 266 we show in Appendix C that this approach fails on other (non-complete) Δ -ary trees.

³ It is tempting to define $\text{score}(u)$ as the sum of weighted advice from the root to u . However, when comparing two vertices, the advice of their least common ancestor would be counted twice, which we prefer to avoid.

⁴ Actually, a similar formula could be derived choosing θ to be the uniform distribution over all nodes, but for technical reasons it is easier to restrict it to leaves only.

267 **2.2 Algorithm A_{walk}**

In our context, there is no distribution over treasure location and we are free to choose θ as we like. We take θ to be the distribution defined by a simple random process. Starting at the root, at each step, walk down to a child uniformly at random. until reaching a leaf. For a leaf v , define $\theta(v)$ as the probability that this process eventually reaches v . Our extension of θ can be interpreted as $\theta(v)$ being the probability that this process passes through v . Formally, $\theta(\sigma) = 1$, and $\theta(u) = (\Delta_\sigma \prod_{w \in \langle \sigma, u \rangle} (\Delta_w - 1))^{-1}$. It turns out that this choice, slightly changed, works remarkably well, and gives an optimal algorithm in noise conditions that practically match those of our lower bound. For a vertex $u \neq \sigma$, define:

$$\beta(u) = \prod_{w \in [\sigma, u)} \Delta_w.$$

268 It is a sort of approximation of $1/\theta(u)$, which we prefer for technical convenience. Indeed,
 269 for all u , $1/\beta(u) \leq \theta(u)$. A wonderful property of this β (besides the fact that it gives rise
 270 to an optimal algorithm) is that to calculate $\beta(v)$ (just like θ), one only needs to know the
 271 degrees of the vertices from v up to the root. It is hard to imagine distributions on leaves
 272 that allow us to calculate the probability of being in a subtree without knowing anything
 273 about it!

274 In the walking algorithm, if v is a candidate for exploration, these nodes must have been
 275 visited already and so the algorithm does not need any a priori knowledge of the structure
 276 of the tree. The following claim will be soon useful:

► **Claim 6.** The following two inequalities hold for every $c < 1$:

$$\sum_{v \in T} \frac{c^{d(v)}}{\beta(v)} \leq \frac{1}{1-c}, \quad \sum_{v \in T} \frac{d(v)c^{d(v)}}{\beta(v)} \leq \frac{c}{(1-c)^2}.$$

277 **Proof.** To prove the first inequality, follow the same random walk defining θ . Starting at the
 278 root, at each step choose uniformly at random one of the children of the current vertex. Now,
 279 while passing through a vertex v collect $c^{d(v)}$ points. No matter what choices are made, the
 280 number of points is at most $1 + c + c^2 + \dots = 1/(1-c)$. On the other hand, $\sum_{v \in T} \theta(v)c^{d(v)}$ is
 281 the expected number of points gained. The result follows since $1/\beta(v) \leq \theta(v)$. The second
 282 inequality is derived similarly, using the fact that $c + 2c^2 + 3c^3 + \dots = c/(1-c)^2$. ◀

For a vertex $u \in T$ and previously seen advice \mathbf{adv} define:

$$\mathbf{score}(u) = \frac{2}{3} \log \left(\frac{1}{\beta(u)} \right) - \sum_{w \in \overleftarrow{\mathbf{adv}}(u)} \log(\Delta_w).$$

283 Algorithm A_{walk} keeps track of all vertices that are children of the vertices it explored so
 284 far, and repeatedly walks to and then explores the one with highest score according to the
 285 current advice, breaking ties arbitrarily. Note that the algorithm does not require prior
 286 knowledge of either the tree's structure, or the values of Δ , q , d or n .

287 **2.3 Analysis**

288 Recall the definition of Condition (\star) from Definition 1. The next lemma provides a large
 289 deviation bound tailored to our setting. The proof appears in Appendix B.

► **Lemma 7.** Consider independent random variables X_1, \dots, X_ℓ , where X_i takes the values $(-\log \Delta_i, 0, \log \Delta_i)$ with respective probabilities $(p_i + \frac{q_i}{\Delta_i}, q_i(1 - \frac{2}{\Delta_i}), \frac{q_i}{\Delta_i})$, for parameters $p_i, q_i = 1 - p_i$ and $\Delta_i > 0$. Assume that Condition (\star) holds for some $\varepsilon > 0$. Then for every integer (positive or negative) m ,

$$\mathbb{P}\left(\sum_{i=1}^{\ell} X_i \geq m\right) \leq \frac{(1-\varepsilon)^\ell}{e^{\frac{3m}{4}}} \prod_{i=1}^{\ell} \frac{1}{\sqrt{\Delta_i}}.$$

290 The next theorem states that A_{walk} is optimal up to a constant factor for the regime of noise
291 below the threshold. It establishes Theorem 2.

292 ► **Theorem 8.** Assume that Condition (\star) holds for some fixed $\varepsilon > 0$. Then $\mathcal{M}(A_{walk}) =$
293 $\mathcal{O}(d\sqrt{\Delta})$, where the constant hidden in the \mathcal{O} notation only depends polynomially on $1/\varepsilon$.

294 **Proof.** Denote the vertices on the path from σ to τ by $\sigma = u_0, u_1, \dots, u_d = \tau$ in order.
295 Denote by E_k the expected time to reach u_k once u_{k-1} is reached. We will show that for all
296 k , $E_k = \mathcal{O}(\sqrt{\Delta})$, and by linearity of expectation this concludes the proof.

Once u_{k-1} is visited, A_{walk} only goes to some of the nodes that have score at least as high as u_k . We can therefore bound E_k from above by assuming we go through all of them, and this expression does not depend on the previous choices of the algorithm and the nodes it saw before seeing u_k . The length of this tour is bounded by twice the sum of distances of these nodes from u_k . Hence,

$$E_k \leq 2 \sum_{i=1}^k \sum_{u \in C(u_i)} \mathbb{P}(\text{score}(u) \geq \text{score}(u_k)) \cdot d(u_k, u).$$

297 Where $C(u_k) = T(u_{k-1}) \setminus T(u_k)$, and so $\cup_{i=1}^k C(u_i) = T \setminus T(u_k)$. Recall that scores are
298 defined so that u has a larger score than u_k , if the sum of weighted arrows on the path $\langle u_k, u \rangle$
299 is at least $\frac{2}{3} \log(\beta(u)/\beta(u_k))$. Setting m to be this value, Lemma 7 allows to calculate this
300 probability exactly. Indeed, a vertex x on the path should point towards u_k : this happens
301 with probability $p_x + q_x/\Delta_x$. Otherwise, it points towards u with probability q_x/Δ_x , and
302 elsewhere with probability $q_x(1 - 2/\Delta_x)$. Denoting $c = 1 - \varepsilon$,

$$\begin{aligned} \frac{E_k}{2} &\leq \sum_{i=1}^k \sum_{u \in C(u_i)} \frac{c^{d(u_k, u)-1}}{e^{\frac{3}{4} \cdot \frac{2}{3} \log(\frac{\beta(u)}{\beta(u_k)})}} \sqrt{\prod_{v \in \langle u, u_k \rangle} \frac{1}{\Delta_v}} \cdot d(u_k, u) \\ &= \frac{1}{c} \sum_{i=1}^k \sum_{u \in C(u_i)} \frac{c^{d(u_k, u)}}{\sqrt{\frac{\beta(u)}{\beta(u_k)}}} \sqrt{\frac{\Delta_{u_i}}{\frac{\beta(u_k)}{\beta(u_i)} \cdot \frac{\beta(u)}{\beta(u_i)}}} \cdot d(u_k, u) \\ &\leq \frac{\sqrt{\Delta}}{c} \sum_{i=1}^k c^{d(u_k, u_i)} \sum_{u \in C(u_i)} c^{d(u_i, u)} \frac{\beta(u_i)}{\beta(u)} \cdot (d(u_k, u_i) + d(u_i, u)). \end{aligned}$$

307 By Claim 6, applied to the tree rooted at u_i , we get:

$$\sum_{u \in C(u_i)} \frac{c^{d(u_i, u)} \beta(u_i)}{\beta(u)} < \frac{1}{1-c}, \quad \text{and} \quad \sum_{u \in C(u_i)} \frac{c^{d(u_i, u)} \beta(u_i)}{\beta(u)} d(u_i, u) < \frac{c}{(1-c)^2}.$$

309 And so:

$$\begin{aligned} \frac{E_k}{2} &\leq \frac{\sqrt{\Delta}}{c(1-c)} \sum_{i=1}^k c^{d(u_k, u_i)} d(u_k, u_i) + \frac{\sqrt{\Delta}}{(1-c)^2} \sum_{i=1}^k c^{d(u_k, u_i)} \\ &\leq \frac{(1+c)\sqrt{\Delta}}{(1-c)^3} \leq \frac{2\sqrt{\Delta}}{\varepsilon^3} = \mathcal{O}(\sqrt{\Delta}), \end{aligned}$$

311
312

313 where we again used the equality $c + 2c^2 + 3c^3 + \dots = c/(1 - c)^2$. ◀

314 **3 Query Algorithms**

315 **3.1 An $\mathcal{O}(\sqrt{\Delta} \log \Delta \log^2 n)$ Queries Algorithm**

316 Our next goal is to prove the first item in Theorem 3. As is common in search on trees, our
 317 technique in this section is based on separators. We say a node u is a *separator* of T if all
 318 the connected components of $T \setminus \{u\}$ are of size at most $|T|/2$. It is well known that such
 319 a node exists. Assume there is some local procedure, that given a vertex u decides with
 320 probability $1 - \delta$ in which one of the connected components of $T \setminus \{u\}$, the treasure resides.
 321 Applying this procedure on a separator of the tree, and then focusing the search recursively
 322 only on the component it pointed out, results in a type of algorithm we call a *separator based*
 323 algorithm. It uses the local procedure at most $\lceil \log_2 n \rceil$ times, and by a union bound, finds
 324 the treasure with probability at least $1 - \lceil \log_2 n \rceil \delta$. Broadly speaking, we will be interested
 325 in the expected running time of this sort of algorithm conditioned on it being successful.
 326 This sort of conditioning complicates matters slightly. In what follows, we assume that the
 327 set of separators for the tree is fixed.

328 **Proof.** (of the first item in Theorem 3) The algorithm we build is denoted \mathbf{A}_{sep} . It runs a
 329 separator based algorithm in parallel to some arbitrary exhaustive search algorithm. The
 330 meaning of *in parallel* here simply means that the two algorithms are run in an alternating
 331 fashion. Fix some small h . The local exploration procedure, denoted \mathbf{local}_h , for a vertex u
 332 proceeds as follows.

333 **Procedure $\mathbf{local}_h(u)$.** Consider the tree $T_h(u)$ rooted at u consisting of all vertices satisfy-
 334 ing $\log_{\Delta} \beta(v) < h$ together with their children. So a leaf of $v \in T_h(u)$ is either a leaf of T , or
 335 satisfies $\Delta^h \leq \beta(v) < \Delta^{h+1}$. Denote the second kind a *nominee*. Call a nominee *promising* if
 336 the number of weighted arrows pointing to v is large, specifically, if $\sum_{w \in [u, v]} X_w \geq \frac{2}{3} h \log \Delta$,
 337 where $X_w = \log \Delta_w$ if the advice at w is pointing to v , $X_w = -\log \Delta_w$ if it is pointing to u ,
 338 and $X_w = 0$ otherwise. Viewing it as a query algorithm, we now run the walking algorithm
 339 \mathbf{A}_{walk} on $T_h(u)$ (starting at its root u) until it either finds the treasure or finds a promising
 340 nominee. In the latter case, $\mathbf{local}_h(u)$ declares that the treasure is on the connected com-
 341 ponent of $T \setminus \{u\}$ containing this nominee. If $\tau \in T_h(u)$ then set $\tau_u = \tau$. Otherwise let τ_u
 342 be the leaf of $T_h(u)$ closest to the treasure, and so in this case τ_u is a nominee. Say that u is
 343 *h -misleading* if either (1) $\tau \notin T_h(u)$ and τ_u is not promising, or (2) there is some promising
 344 nominee $v \in T_h(u)$ that is not in the same connected component of $T \setminus \{u\}$ as τ_u . Note
 345 that if u is not h -misleading then $\mathbf{local}_h(u)$ necessarily outputs the correct component of
 346 $T \setminus \{u\}$, namely, the one containing the treasure. The proof of the following lemma appears
 347 in Appendix D. The part regarding regular trees will be needed later.

349 **► Lemma 9.** For any u , $\mathbb{P}(u \text{ is } h\text{-misleading}) \leq (\Delta + 1)(1 - \varepsilon)^h$. Also, for any event X such
 350 that X occurring always implies that u is not misleading, we have $\mathbb{P}(X) \mathcal{Q}(\mathbf{local}_h(u) \mid X) =$
 351 $\mathcal{O}(\sqrt{\Delta} \log \Delta \cdot h)$. In the case the tree is regular, these bounds become $2(1 - \varepsilon)^h$ and $\mathcal{O}(\sqrt{\Delta} \cdot h)$
 352 respectively. The constant hidden in the \mathcal{O} notation only depends polynomially on $1/\varepsilon$.

353 Taking $h = -3 \log(2n) / \log(1 - \varepsilon)$, gives $\mathbb{P}(u \text{ is misleading}) \leq 1/n^2$. Denote by **Good**
 354 the event that none of the separators encountered are misleading. By a union bound,
 355 $\mathbb{P}(\mathbf{Good}^c) \leq 1/n$.

$$356 \quad \mathcal{Q}(\mathbf{A}_{sep}) = \mathbb{P}(\mathbf{Good}) \mathcal{Q}(\mathbf{A}_{sep} \mid \mathbf{Good}) + \mathbb{P}(\mathbf{Good}^c) \mathcal{Q}(\mathbf{A}_{sep} \mid \mathbf{Good}^c). \quad (1)$$

357 As A_{sep} runs an exhaustive search algorithm in parallel, the second term is $\mathcal{O}(1)$. For the first
 358 term, note that conditioning on \mathbf{Good} , all local procedures either find the treasure or give
 359 the correct answer, and so there are $\mathcal{O}(\log n)$ of them and they eventually find the treasure.
 360 Denote by u_i the i -th vertex that \mathbf{local}_h is executed on. By linearity of expectation, and
 361 applying Lemma 9, the first term of (1) is $\mathbb{P}(\mathbf{Good}) \sum_i \mathcal{Q}(\mathbf{local}_h(u_i) \mid \mathbf{Good}) = \mathcal{O}(\log n \cdot$
 362 $\sqrt{\Delta} \log \Delta \cdot h) = \mathcal{O}(\sqrt{\Delta} \log \Delta \log^2 n)$. As $\log(1+x) > x$ always, then $-1/\log(1-\varepsilon) \leq 1/\varepsilon$,
 363 and the hidden factor in the \mathcal{O} is as stated. \blacktriangleleft

364 3.2 An Almost Tight Result for Regular Trees

365 We now turn our attention to the second item in Theorem 3. Due to space constraints
 366 its full proof is deferred to Appendix E. At a high level, we run two algorithms in parallel
 367 (i.e., in an alternating fashion): A_{fast} , and A_{mid} . Algorithm A_{fast} is actually A_{sep} applied
 368 with parameter $h = \Theta(\log \log n)$ instead of $\Theta(\log n)$. Using Lemma 9, with probability
 369 $1 - 1/\log^{\mathcal{O}(1)}(n)$, the local procedure of A_{fast} always detects the correct component for each
 370 separator, and A_{fast} needs an expected number of $\mathcal{O}(\sqrt{\Delta} \cdot \log n \cdot \log \log n)$ queries to find
 371 the treasure. This is the running time we are aiming for.

372 Algorithm A_{mid} is similar to A_{sep} except it uses a different subroutine for local explora-
 373 tion. It then remains to show that it finds the treasure using a relatively low expected
 374 number of queries even conditioning on the event that caused A_{fast} to fail, namely, the
 375 event that there is a misleading separator at the scale $h = \Theta(\log \log n)$. The query complex-
 376 ity of A_{mid} does blow up under this event but we show that the blowup is not that bad, and
 377 can be compensated by the fact that the bad event has small probability. This is the core of
 378 the proof, and what requires most work. In fact, the complexity of the arguments led us to
 379 restrict the discussion to regular trees and also modify the subroutine for local exploration
 380 to ease the analysis.

381 4 Lower Bounds

382 We next prove Items (1) and (2a) of Theorem 4. Item (2b) is proved in Appendix G.2.

383 4.1 Exponential Complexity Above the Threshold

384 We wish to prove Item (1) in Theorem 4. Namely, that for every fixed $\varepsilon > 0$, and for every
 385 complete Δ -ary tree, if $q \geq \frac{1+\varepsilon}{\sqrt{\Delta-1}} \cdot (1 + \frac{1}{\Delta-1})$, then every randomized search algorithm has
 386 query (and move) complexity which is both exponential in the depth d of the treasure and
 387 polynomial in n . In fact, this lower bound holds even if the algorithm has access to the
 388 advice of all internal nodes. The following lemma is proved in Appendix G.1:

389 \blacktriangleright **Lemma 10.** *Assume the treasure is placed in a leaf τ of the complete Δ -ary tree. Denote by*
 390 \mathbf{adv} *the random advice on all internal nodes, then the expected number of leaves u satisfying*
 391 $|\mathbf{adv}(u)| > |\mathbf{adv}(\tau)|$, *is a lower bound on the query complexity of any algorithm.*

392 Using Lemma 10, all we need to do is approximate the number of leaves u satisfying
 393 $|\mathbf{adv}(u)| > |\mathbf{adv}(\tau)|$. When comparing the number of pointers that point towards each of
 394 two different nodes, only the pointers of the internal nodes on the path between them may
 395 influence on the result. The probability that a leaf u “beats” the treasure τ in the sense of
 396 Lemma 10, is at least the probability that exactly one node on the path points to u and

397 none of the rest point towards the treasure. This probability is at least

$$398 \quad \frac{q}{\Delta} \cdot \left(q \cdot \left(1 - \frac{1}{\Delta} \right) \right)^{d(u,\tau)-2}.$$

399 There are precisely $(\Delta - 1)^D$ leaves whose distance from the treasure is $2D$. Therefore, the
400 expected number of leaves that beat the treasure is at least:

$$401 \quad \frac{q}{\Delta} (\Delta - 1)^D q^{2D-2} \cdot \left(1 - \frac{1}{\Delta} \right)^{2D-2} = \frac{\Delta}{q(\Delta - 1)^2} \cdot \left(\frac{q^2(\Delta - 1)^3}{\Delta^2} \right)^D \geq \frac{\Delta}{q(\Delta - 1)^2} \cdot (1 + \varepsilon)^2 D.$$

403 Item (1) in Theorem 4 follows. ◀

404 4.2 A Lower Bound of $\Omega(\sqrt{\Delta} \cdot \log_{\Delta} n)$ when $q \sim 1/\sqrt{\Delta}$

405 We now prove Item (2a) in Theorem 4. We wish to prove that for $\Delta \geq 3$, on the complete
406 Δ -ary tree of depth D , any algorithm needs $\Omega(q\Delta D)$ queries in expectation. Note that,
407 in particular, when q is roughly $1/\sqrt{\Delta}$, and n is the tree size, the query complexity is
408 $\Omega(\sqrt{\Delta} \cdot \log_{\Delta} n)$. Before proving this lower bound, we need the following observation (proved
409 in Appendix G.3)

410 ► **Observation 11.** Any randomized algorithm that finds a uniformly chosen treasure between
411 k identical objects needs an at least $(k + 1)/2$ queries in expectation.

412 To prove the lower bound of $\Omega(q\Delta D)$, consider the complete Δ -ary tree of depth D . We
413 prove by induction on D , that if the treasure is placed uniformly at random. in one of the
414 leaves, then the expected query complexity of any algorithm is at least $q(\Delta/2 - 1)D$. If
415 $D = 0$, then there is nothing to show. Assume this is true for D , and we shall prove it for
416 $D + 1$. Let $T_1, \dots, T_{\Delta-1}$ be the subtrees hanging down from the root (in the induction, the
417 “root” is actually an internal node, and so has $\Delta - 1$ children), each having depth D . Let i
418 be the index such that $\tau \in T_i$, and denote by Q the number of queries before the algorithm
419 makes its first query in T_i . We will assume that the algorithm gets the advice in the root
420 for free. Denote by Y the event that the root is faulty. In this case, Observation 11 applies,
421 and we need at least $\Delta/2 - 1$ queries to hit the correct tree. We subtracted one query from
422 the count because we want to count the number of queries strictly before querying inside T_i .
423 We therefore get $\mathbb{E}[Q] \geq \mathbb{P}(Y) \cdot \mathbb{E}[Q | Y] \geq q(\Delta/2 - 1)$. By linearity of expectation, using
424 the induction hypothesis, we get the result for a uniformly placed treasure over the leaves,
425 and so it holds also in the adversarial case. ◀

426 5 Open Problems

427 Closing the small gap between the upper and lower bounds for the query setting remains
428 open. The noisy advice model may well be interesting to study in other search settings.
429 In particular, obtaining efficient search algorithms for general graphs is highly intriguing.
430 Even though the likelihood of a node being the treasure under a uniform prior can still be
431 computed in principle, it is not so easy to compare two nodes as in Theorem 8 because there
432 may be more than a single path between them.

433 In a limited regime of noise, we believe that memoryless strategies might very well be
434 efficient also on general graphs, and we pose the following conjecture. Proving it may require
435 the use of tools from the theory of RWRE, which seem to be lacking in the context of general
436 graph topologies.

437 ► **Conjecture 12.** There exists a probabilistic following algorithm that finds the treasure in
438 expected linear time on any undirected graph assuming $q < c/\Delta$ for a small enough $c > 0$.

439 — References —

- 440 1 Andrei Asinowski, Jean Cardinal, Nathann Cohen, Sébastien Collette, Thomas Hackl, Michael Hoffmann, Kolja B. Knauer, Stefan Langerman, Michal Lason, Piotr Micek, Günter Rote, and Torsten Ueckerdt. Coloring hypergraphs induced by dynamic point sets and bottomless rectangles. *CoRR*, abs/1302.2426, 2013.
- 441 2 Javed A. Aslam and Aditi Dhagat. Searching in the presence of linearly bounded errors. In *STOC*, pages 486–493. ACM, 1991.
- 442 3 Yosi Ben-Asher, Eitan Farchi, and Ilan Newman. Optimal search in trees. *SIAM J. Comput.*, 28(6):2090–2102, 1999.
- 443 4 Michael Ben-Or and Avinatan Hassidim. The bayesian learner is optimal for noisy binary search (and pretty good for quantum as well). In *FOCS*, pages 221–230, 2008.
- 444 5 Ryan S. Borgstrom and S. Rao Kosaraju. Comparison-based search in the presence of errors. In *STOC*, pages 130–136. ACM, 1993.
- 445 6 Mark Braverman and Elchanan Mossel. Noisy sorting without resampling. In *SODA*, pages 268–276, 2008.
- 446 7 Gerth Stølting Brodal, Rolf Fagerberg, Irene Finocchi, Fabrizio Grandoni, Giuseppe F. Italiano, Allan Grønlund Jørgensen, Gabriel Moruz, and Thomas Mølhave. Optimal resilient dynamic dictionaries. In *Algorithms - ESA 2007, 15th Annual European Symposium, Eilat, Israel, October 8-10, 2007, Proceedings*, pages 347–358, 2007.
- 447 8 Ferdinando Cicalese and Ugo Vaccaro. Optimal strategies against a liar. *Theor. Comput. Sci.*, 230(1-2):167–193, 2000.
- 448 9 Alexander Drewitz and Alejandro F. Ramírez. Selected topics in random walk in random environment. *Topics in Percolative and Disordered Systems, Springer Proceedings in Mathematics and Statistics*, 69:23–83, 2014.
- 449 10 Ehsan Emamjomeh-Zadeh, David Kempe, and Vikrant Singhal. Deterministic and probabilistic binary search in graphs. In *STOC*, pages 519–532, 2016.
- 450 11 Uriel Feige, Prabhakar Raghavan, David Peleg, and Eli Upfal. Computing with noisy information. *SIAM J. Comput.*, 23(5):1001–1018, October 1994.
- 451 12 Irene Finocchi, Fabrizio Grandoni, and Giuseppe F. Italiano. Resilient search trees. In *SODA*, pages 547–553, 2007.
- 452 13 Irene Finocchi and Giuseppe F. Italiano. Sorting and searching in the presence of memory faults (without redundancy). In *STOC*, pages 101–110, 2004.
- 453 14 Ehud Fonio, Yael Heyman, Lucas Boczkowski, Aviram Gelblum, Adrian Kosowski, Amos Korman, and Ofer Feinerman. A locally-blazed ant trail achieves efficient collective navigation despite limited information, *eLife* 2016;5:e20185. 2016.
- 454 15 Nicolas Hanusse, David Ilcinkas, Adrian Kosowski, and Nicolas Nisse. Locating a target with an agent guided by unreliable local advice: How to beat the random walk when you have a clock? In *PODC*, pages 355–364, 2010.
- 455 16 Nicolas Hanusse, Dimitris Kavvadias, Evangelos Kranakis, and Danny Krizanc. Memoryless search algorithms in a network with faulty advice. *Theoretical Computer Science*, 402(2–3):190 – 198, 2008.
- 456 17 Nicolas Hanusse, Evangelos Kranakis, and Danny Krizanc. Searching with mobile agents in networks with liars. *Discrete Applied Mathematics*, 137(1):69–85, 2004.
- 457 18 Richard M. Karp and Robert Kleinberg. Noisy binary search and its applications. In *SODA*, pages 881–890, 2007.
- 458 19 Eduardo Sany Laber and Loana Tito Nogueira. Fast searching in trees. In *Electronic Notes on Discrete Mathematics*, 2001.
- 459 20 Shay Mozes, Krzysztof Onak, and Oren Weimann. Finding an optimal tree searching strategy in linear time. In *SODA*, pages 1096–1105, 2008.

- 488 **21** Krzysztof Onak and Pawel Parys. Generalization of binary search: Searching in trees and
489 forest-like partial orders. In *FOCS*, pages 379–388, 2006.
- 490 **22** Andrzej Pelc. Searching games with errors - fifty years of coping with liars. *Theor. Comput.*
491 *Sci.*, 270(1-2):71–109, 2002.
- 492 **23** Alain-Sol Snitzman. Topics in random walks in random environment. *ICTP Lecture Notes*
493 *Series*, 2004.

494 Appendix

495 **A** A Natural Attempt that Fails

496 The following example provides useful intuition. Consider the permissive scenario where
497 the tree is known and all the advice at the internal nodes of the tree is available to the
498 algorithm. Since advice is sampled independently at each node, it seems natural, at least
499 at first glance, to associate each node with a “likelihood rank”, being the total number of
500 advice pointers pointing to it in the whole tree. A natural query algorithm A_{natural} would
501 then be to query the nodes one by one, according to the resulting ranking. In other terms
502 A_{natural} goes at each step to the unvisited node having most arrows pointing to it among
503 the neighbors of previously seen nodes. Unfortunately, this naive strategy turns out highly
504 inefficient. To see why, consider a complete Δ -ary tree of depth D , except for one child
505 of the root, which is turned into a leaf, trimming its $(\Delta - 1)^{D-1}$ descendants. Assume
506 further that this particular child is in fact the treasure location τ . For any leaf $v \neq \tau$, with
507 probability $\frac{q}{\Delta} \cdot q^{D-1} (1 - \frac{1}{\Delta})^{D-1}$ the root points towards v and all the rest of the nodes on
508 the path $\langle \sigma, \ell \rangle$ do not point upward (towards the treasure). This makes all nodes of the path
509 better ranked than τ , and so A_{natural} would query them all before querying τ . There are
510 $(\Delta - 1)^D$ such leaves, and hence the expected number of nodes queried before the treasure
511 is at least $q^D (\Delta - 1)^{D-1} (1 - \frac{1}{\Delta})^D$ in expectation. Even for q as small as c/Δ this number
512 is exponential in the depth of the tree.

513 The main reason why A_{natural} fails is that, although the suggested ranking reflects the
514 correct likelihood under a uniform prior of the treasure, this uniform prior is not tailored to
515 our setting of adversarial treasure location. Instead, the algorithm we present in Section 2
516 relies on a ranking which incorporates in a clever way the tree structure.

517 **B** Proof of the Chernoff Estimate

► **Lemma 7 (restated).** Consider independent random variables X_1, \dots, X_ℓ , where X_i takes the values $(-\log \Delta_i, 0, \log \Delta_i)$ with respective probabilities $(p_i + \frac{q_i}{\Delta_i}, q_i(1 - \frac{2}{\Delta_i}), \frac{q_i}{\Delta_i})$, for parameters $p_i, q_i = 1 - p_i$ and $\Delta_i > 0$. Assume that Condition (\star) holds for some $\varepsilon > 0$. Then for every integer (positive or negative) m ,

$$\mathbb{P} \left(\sum_{i=1}^{\ell} X_i \geq m \right) \leq \frac{(1 - \varepsilon)^\ell}{e^{\frac{3m}{4}}} \prod_{i=1}^{\ell} \frac{1}{\sqrt{\Delta_i}}.$$

518 **Proof.** For any $s \in \mathbb{R}$,

$$\begin{aligned}
 519 \quad \mathbb{P}\left(\sum_{i=1}^{\ell} X_i \geq m\right) &= \mathbb{P}\left(e^s \sum_{i=1}^{\ell} X_i \geq e^{sm}\right) \leq \frac{\mathbb{E}\left[e^s \sum_{i=1}^{\ell} X_i\right]}{e^{sm}} = \frac{\prod_i \mathbb{E}\left[e^{sX_i}\right]}{e^{sm}} \\
 520 \quad &= \frac{1}{e^{sm}} \prod_{i=1}^{\ell} \left(\frac{p_i + \frac{q_i}{\Delta_i}}{e^{\log(\Delta_i)s}} + q_i \left(1 - \frac{2}{\Delta_i}\right) + \frac{q_i}{\Delta_i} e^{\log(\Delta_i)s} \right) \\
 521 \quad &\leq \frac{1}{e^{sm}} \prod_{i=1}^{\ell} \left(\frac{1}{\Delta_i^s} + q_i + q_i \Delta_i^{s-1} \right). \\
 522
 \end{aligned}$$

523 We take $s = \frac{3}{4}$, and get:

$$524 \quad \mathbb{P}\left(\sum_{i=1}^{\ell} X_i \geq m\right) \leq \frac{1}{e^{\frac{3m}{4}}} \prod_{i=1}^{\ell} \left(\Delta_i^{-\frac{3}{4}} + q_i + q_i \Delta_i^{-\frac{1}{4}} \right) \leq \frac{1}{e^{\frac{3m}{4}}} \prod_{i=1}^{\ell} \frac{1-\varepsilon}{\sqrt{\Delta_i}}$$

525 Where for the last step we used Condition (\star) which says:

$$526 \quad q_i < \frac{1-\varepsilon - \Delta_i^{-\frac{1}{4}}}{\sqrt{\Delta_i} + \Delta_i^{\frac{1}{4}}} \implies q_i \Delta_i^{\frac{1}{2}} + q_i \Delta_i^{\frac{1}{4}} + \Delta_i^{-\frac{1}{4}} < 1-\varepsilon \implies \Delta_i^{-\frac{3}{4}} + q_i + q_i \Delta_i^{-\frac{1}{4}} < \frac{1-\varepsilon}{\sqrt{\Delta_i}}$$

527 ◀

528 **C** Taking θ to be the Uniform Distribution is not a Good Idea

529 As mentioned at the end of Section 2.1, when our tree is a complete Δ -ary tree, choosing θ
 530 to be the uniform distribution over the leaves results in an efficient algorithm with respect
 531 to the worst case placement of the treasure. Trying to tackle more general trees, perhaps
 532 the most natural a priori distribution is the uniform one over the nodes of the tree. As our
 533 technical presentation accommodates only distributions on leaves, we take θ to be uniform
 534 over the leaves only, and remark that the same result we get here applies to the former case.

535 Unfortunately, we show that this variant may take exponentially many queries before
 536 finding the treasure no matter what q is. This in fact follows from a similar argument to
 537 the one mentioned in Section A. The instance we consider is a complete Δ -ary tree of depth
 538 D , except for one child of the root, which is turned into a leaf, trimming its $(\Delta - 1)^{D-1}$
 539 descendants. We consider the case that this particular child is in fact the treasure location

540 τ .

541 Recall from Section 2.1 that $\text{score}(u) > \text{score}(\tau)$ iff:

$$542 \quad \sum_{w \in \langle u, \tau \rangle \cap \vec{\text{adv}}(u)} \log(\Delta_w) - \sum_{w \in \langle u, \tau \rangle \cap \vec{\text{adv}}(\tau)} \log(\Delta_w) > \frac{2}{3} \log \left(\frac{\theta(\tau)}{\theta(u)} \right), \quad (2)$$

544 where $\theta(u)$ is now understood as the ratio between the number of leaves in $T(u)$ divided by
 545 the total number of leaves in T , as opposed to the total number of leaves if the tree was a
 546 complete tree.

547 In particular, consider any node u at distance $a \cdot D$ from the root for some $a < 2/5$.
 548 This node u owns a tree $T(u)$ of size $(\Delta - 1)^{D(1-a)}$, hence $\theta(u) \sim (\Delta - 1)^{-a}$. In contrast
 549 $\theta(\tau) = (\Delta - 1)^{-D}$. Therefore,

$$550 \quad \frac{2}{3} \log \left(\frac{\theta(\tau)}{\theta(u)} \right) = -\frac{2}{3} D(1-a) \log(\Delta),$$

551

552 where we write $\log(\Delta)$ in place of $\log(\Delta - 1)$ as it does not change the nature of the result,
 553 only the choice of the constant $2/5$. On the other hand there are only $a \cdot \Delta$ nodes on the path
 554 $\langle u, \tau \rangle$ so the left side of (2) is always greater than $-aD \log \Delta$. In other words if $a < 2/5$
 555 then $-\frac{2}{3}D(1-a) \log \Delta < -aD \log \Delta$, and any node u at depth $a \cdot D$ has a better score
 556 than τ , regardless of the advice on the path $\langle \tau, u \rangle$ which means that our algorithm needs
 557 $(\Delta - 1)^{2/5D}$ steps at least.

558 D Proof of Lemma 9

559 ► **Lemma 9 (restated).** For any u , $\mathbb{P}(u \text{ is } h\text{-misleading}) \leq (\Delta + 1)(1 - \varepsilon)^h$. Also, for any
 560 event X such that X occurring always implies that u is not misleading, we have $\mathbb{P}(X) \mathcal{Q}(\text{local}_h(u) \mid X) =$
 561 $\mathcal{O}(\sqrt{\Delta} \log \Delta \cdot h)$. In the case the tree is regular, these bounds become $2(1 - \varepsilon)^h$ and $\mathcal{O}(\sqrt{\Delta} \cdot h)$
 562 respectively. The constant hidden in the \mathcal{O} notation only depends polynomially on $1/\varepsilon$.

563 **Proof.** To check the probability that u is misleading, consider two cases:

564 1. $\tau \notin T_h(u)$, and τ_u is not promising. By Lemma 7, and recalling that $\Delta^h \leq \beta(\tau_u)$, the
 565 probability τ_u is not promising is:

$$566 \quad \mathbb{P} \left(\sum_{w \in \langle u, \tau_u \rangle} -X_w \leq \frac{2}{3} h \log(\Delta) \right) = \mathbb{P} \left(\sum_{w \in \langle u, \tau_u \rangle} X_w \geq -\frac{2}{3} \cdot h \log(\Delta) \right)$$

$$567 \quad \leq \prod_{w \in \langle u, \tau_u \rangle} \frac{1 - \varepsilon}{\sqrt{\Delta_w}} \cdot e^{\frac{3}{4} \cdot \frac{2}{3} h \log(\Delta)} = \frac{(1 - \varepsilon)^{d(u, \tau_u)}}{\sqrt{\beta(\tau_u)}} \Delta^{\frac{h}{2}} \leq (1 - \varepsilon)^{d(u, \tau_u)}.$$

569 As $d(u, \tau_u) \geq \log_{\Delta} \beta(\tau_u) \geq h$, this is at most $(1 - \varepsilon)^h$.

570 2. If v is a nominee that is not in the same connected component of $T \setminus \{u\}$ as τ_u , then by
 571 Lemma 7, the probability that v is promising is

$$572 \quad \mathbb{P} \left(\sum_{w \in \langle u, v \rangle} X_w \geq \frac{2}{3} \log \Delta \cdot h \right) \leq \prod_{w \in \langle u, v \rangle} \frac{1 - \varepsilon}{\sqrt{\Delta_w}} \cdot e^{-\frac{3}{4} \cdot \frac{2}{3} h \log \Delta}$$

$$573 \quad = \frac{(1 - \varepsilon)^{d(u, v)}}{\sqrt{\beta(v)}} \Delta^{-\frac{h}{2}} \leq \frac{(1 - \varepsilon)^{d(u, v)}}{\Delta^h}.$$

575 However, denote by L the set of nominees in T_u . As they are a subset of the leaves of
 576 T_u , by the way θ is defined:

$$577 \quad 1 \geq \sum_{x \in L} \theta(v) \geq \sum_{x \in L} \frac{1}{\beta(v)} \geq \sum_{x \in L} \frac{1}{\Delta^{h+1}} = \frac{|L|}{\Delta^{h+1}} \quad (3)$$

578 So, $|L| \leq \Delta^{h+1}$. Therefore, by a union bound, the probability that there exists a nominee
 579 v that renders u misleading is at most $\Delta(1 - \varepsilon)^h$.

580 The probability that u is misleading is then at most $(1 + \Delta)(1 - \varepsilon)^h$ as stated. In the case
 581 where the tree is regular, the analysis is the same, except that in (3), $\beta(v) = \Delta^h$, and so
 582 following the same logic, $|L| \leq \Delta^h$, and this part contributes only $(1 - \varepsilon)^h$.

583 For the second part of the lemma, consider some event X where u is not misleading. As
 584 τ_u is either the actual treasure or promising, and acts as the treasure in the eyes of \mathbf{A}_{walk} ,
 585 then the local procedure stops when it encounters τ_u . It might actually stop before (because
 586 it found another promising node), so, $\mathcal{Q}(\text{local}_h(u) \mid X) \leq \mathcal{Q}(\mathbf{A}_{walk}(T_h(u)) \mid X)$. Therefore,

$$587 \quad \mathbb{P}(X) \mathcal{Q}(\text{local}(u) \mid X) \leq \mathbb{P}(X) \mathcal{Q}(\mathbf{A}_{walk}(T_h(u)) \mid X) \leq \mathcal{Q}(\mathbf{A}_{walk}(T_h(u))) = \mathcal{O}(\sqrt{\Delta} \cdot \text{depth}(T_h(u)))$$

588 But the depth of T_u is at most $\mathcal{O}(h \log \Delta)$, since its leaves satisfy $\beta(v) < \Delta^{h+1}$, and $\beta(v) \geq$
 589 $2^{\text{depth}(v)}$. For the case of a regular tree, $\beta(v) = \Delta^{\text{depth}(v)}$ and so the depth of T_u is at most
 590 h , giving the result. \blacktriangleleft

591 **E** A More Involved $\mathcal{O}(\sqrt{\Delta} \log n \cdot \log \log n)$ Algorithm for Regular Trees

592 In this section we give a formal proof for the second item in Theorem 3. That is, we
 593 present algorithm $\mathbf{A}_{2\text{-layers}}$ which is designed for (not necessarily complete) Δ -ary trees,
 594 and performs extremely well in the regime where $q < c/\sqrt{\Delta}$ for some small enough positive
 595 constant c . Specifically, in that regime, it finds the treasure in $\mathcal{O}(\sqrt{\Delta} \log n \cdot \log \log n)$ queries
 596 in expectation. Before we continue, let us note that taking a small enough c , the condition
 597 $q < c/\sqrt{\Delta}$ we are using here actually implies⁵ Condition (\star) with $\varepsilon = (1 - 2^{-1/4})/2$.

598 Algorithm $\mathbf{A}_{2\text{-layers}}$ runs two algorithms in parallel, namely, \mathbf{A}_{fast} , and \mathbf{A}_{mid} . Algorithm
 599 \mathbf{A}_{fast} is actually \mathbf{A}_{sep} , except that it applies the local procedure with parameter h being
 600 $h_2 = \lceil \kappa_2 \log \log n \rceil$ rather than $\Theta(\log n)$. Algorithm \mathbf{A}_{mid} is similar to \mathbf{A}_{sep} , as it also
 601 uses h being $h_1 = \lceil \kappa_1 \log n \rceil$. However it uses a different local exploration procedure, see
 602 more details in Section E.1. κ_1 and κ_2 are constant independent of n whose value will be
 603 determined later. We will henceforth omit the ceiling $\lceil \cdot \rceil$ in the interest of readability.

604 Let us first recall some of the definitions that were introduced in Section 3.1. Since only
 605 regular trees are considered in this section, some of the definitions are simplified. Here $T_h(u)$
 606 denotes the tree of nodes at distance at most h from u . Call a leaf $v \in T_h(u)$ a *nominee* if
 607 its distance to u is exactly h . Denote by $\mathcal{U}(u)$ the set of nominees that are not in the same
 608 component as τ_u in $T \setminus \{u\}$. Call a nominee *promising* if $\sum_{w \in [u,v]} X_w \geq \frac{2}{3}h$, where $X_w = 1$ if
 609 the advice at w is pointing to v , $X_w = -1$ if it is pointing to u , and $X_w = 0$ otherwise. Note
 610 that X_u can never be -1 . Let τ_u be the leaf on $T_h(u)$ closest to τ if $\tau \notin T_h(u)$ and $\tau_u = \tau$
 611 otherwise. Recall also that u is called *h -misleading*, if one of the two following happens (1)
 612 $\tau_u \neq \tau$ and τ_u is not promising, or (2) There is some promising nominee in $\mathcal{U}(u)$.

613 Let **Excellent** be the event that no separator on the way to the treasure is h_2 -misleading.
 614 The following claim is a direct consequence of Lemma 9 (regular tree case) and linearity of
 615 expectation, summing the query complexity of the $\lceil \log n \rceil$ separators on the way to the
 616 treasure.

► Claim 13.

$$617 \quad \mathbb{P}(\text{Excellent}) \cdot \mathcal{Q}(\mathbf{A}_{fast} \mid \text{Excellent}) = \mathcal{O}\left(\sqrt{\Delta} \log n \cdot \log \log n\right).$$

618 To bound the total expected number of queries, we run in parallel algorithm \mathbf{A}_{mid} . All that
 619 remains is then to prove that $\mathbb{P}(\text{Excellent}^c) \cdot \mathcal{Q}(\mathbf{A}_{mid} \mid \text{Excellent}^c) = \mathcal{O}(\sqrt{\Delta} \log n)$.

620 E.1 Algorithm \mathbf{A}_{mid}

621 As mentioned, \mathbf{A}_{mid} is similar to \mathbf{A}_{sep} except that it uses a different local procedure. More
 622 precisely, recall that \mathbf{A}_{sep} executes Procedure $\text{local}_h(u)$ by running \mathbf{A}_{walk} on $T_h(u)$ until it
 623 either finds the treasure or finds a promising nominee, and in the latter case, it declares that
 624 the treasure is on the connected component of $T \setminus \{u\}$ containing this nominee. In the context

⁵ Indeed, recall that for regular trees, Condition (\star) reads $q < \frac{1-\varepsilon-\Delta^{-1/4}}{\sqrt{\Delta}+\Delta^{1/4}}$. Now, $\Delta \geq 2$ implies that
 $1 - \Delta^{-1/4} \geq 1 - 2^{-1/4}$ and $\Delta^{1/4} \leq \sqrt{\Delta}$. Hence $\frac{1-\varepsilon-\Delta^{-1/4}}{\sqrt{\Delta}+\Delta^{1/4}} \geq \frac{1-2^{-1/4}-\varepsilon}{2} \frac{1}{\sqrt{\Delta}}$. We may set $\varepsilon = \frac{1-2^{-1/4}}{2}$
 so that, as soon as $c < \frac{1-2^{-1/4}-\varepsilon}{2} = \frac{1-2^{-1/4}}{4}$, $q < c\Delta^{-1/2}$ implies Condition (\star) with that choice of ε .

625 of Algorithm A_{mid} , for technical commodity, we choose to run Procedure $\text{local}_h(u)$ with a
 626 simpler exploration routine which we call A_{loop} . It is less efficient than A_{walk} but its simplicity
 627 will be useful for analyzing its behaviour in various, “less clean”, circumstances. Indeed,
 628 we will need to analyse the performances of A_{loop} , conditioning on the event Excellent^c ,
 629 implying that some parts of the tree have to be pointing in the wrong direction.

630 The fact that A_{loop} is less efficient than A_{walk} will not affect the final bound, as its running
 631 time will dominate the total running time with very low probability.

632 **Algorithm A_{loop} .** Recall in this section we only deal with Δ -regular trees. Define *level i* as
 633 the set of all nodes at distance i from the root. At each round, A_{loop} only compares nodes
 634 within a given level i . Specifically, it goes to the node in level i with most arrows pointing at
 635 it among the non-visited nodes in level i . Note that it considers only vertices whose parent
 636 has been explored already. The index i is incremented modulo the depth of the tree D , on
 637 every round. Below is a description in pseudocode. The loop over i explains the name A_{loop} .

Algorithm 1: Algorithm A_{loop}

- 1 Continuously loop over levels $1, 2, \dots, D$
- 2 When considering level i , go to the yet unexplored reachable node at the current level (if one exists) that has most arrows pointing to it.

638 In what follows we will analyse Algorithm A_{loop} conditioning on some parts of the tree
 639 being misleading. For readability considerations, the interested reader might wish to first see
 640 how it behaves on a simpler scenario, without any conditioning. The proof of the following
 641 appears in Appendix F.2.

642 ► **Lemma 14.** *Consider a (not necessarily complete) Δ -ary tree. Then $\mathcal{Q}(A_{loop}) = \mathcal{O}(D^3\sqrt{\Delta})$.*

643 In fact, a slightly more refined analysis shows that $\mathcal{Q}(A_{loop}) = \mathcal{O}(D^2\sqrt{\Delta})$, but this not
 644 needed for our current purposes, and so we omit it.

645 E.2 Analysis of A_{mid} Conditioning on Excellent^c

646 To complete the proof of the second item in Theorem 3 we will show that if c small enough,
 647 then

$$648 \quad \mathbb{P}(\text{Excellent}^c) \cdot \mathcal{Q}(A_{mid} \mid \text{Excellent}^c) = \mathcal{O}(\sqrt{\Delta} \log n). \quad 649$$

650 E.2.1 Decomposing Excellent^c

651 At a high level, we seek to break Excellent^c into many elementary bad events. Denote
 652 u_1, \dots, u_ℓ the sequence of separators on the way to the treasure τ . Note that $\ell \leq \lceil \log n \rceil$.
 653 First,

$$654 \quad \text{Excellent}^c = \bigcup_{i \leq \ell} \{u_i \text{ is } h_2\text{-misleading}\}. \quad 655$$

656 Using the union bound argument in Section F (Claim 21),

$$657 \quad \mathcal{Q}(A_{mid} \cap \text{Excellent}^c) \leq \sum_{i \leq \ell} \mathcal{Q}(A_{mid} \cap u_i \text{ is } h_2\text{-misleading}), \quad (4) \quad 658$$

659 where, to keep the equation light we write $\mathcal{Q}(A \cap E)$ in place of $\mathcal{Q}(A \mid E) \cdot \mathbb{P}(E)$ where A is
 660 an algorithm and E is an event.

702 ► **Definition 16.** Let $a, b \in T$ be two nodes such that a is the closest one to τ out of the
 703 nodes in $[a, b]$. Noting that a vertex can never point to itself:

704 ■ For $S \subseteq \langle a, b \rangle$, denote by $M_{\text{sides}}^S(a, b)$ the event that the nodes of S neither point towards
 705 a nor towards b .

706 ■ For $S \subseteq [a, b)$, denote by $M_{\text{up}}^S(a, b)$ the event that the nodes of S all point towards b .

707 ► **Claim 17.** For any a, b and S as in Definition 16,

708 ■ $\mathbb{P}(M_{\text{sides}}^S(a, b)) \leq q^{|S|}$,

709 ■ $\mathbb{P}(M_{\text{up}}^S(a, b)) \leq \left(\frac{q}{\Delta}\right)^{|S|}$.

710 Let us now see in more detail what it means for a node u to be h_2 -misleading. First recall
 711 from the definition that $|[u, \tau_u]| = h_2$, as otherwise $\tau \in T_{h_2}(u)$ and u cannot be h_2 -misleading
 712 because of the path $[u, \tau_u]$. Several cases need to be considered.

713 1. τ_u is not promising, and so the sum of advice on $[u, \tau_u)$ is strictly less than $\frac{2}{3}h_2$. In this
 714 case, at least one of the following two must be true:

715 a. There are $\frac{1}{6}h_2$ locations on the path $[u, \tau_u)$ where the advice points outside of the
 716 path (the value of the corresponding X_i 's is 0). This corresponds to $M_{\text{sides}}^S(\tau_u, u)$ for
 717 some set $S \subseteq [u, \tau_u)$ of size⁶ $|S| = \frac{1}{6}h_2$.

718 b. There are $\frac{1}{12}h_2$ locations on $\langle u, \tau_u \rangle$ that point towards u (the value of the corres-
 719 ponding X_i 's is 1). This corresponds to $M_{\text{up}}^S(\tau_u, u)$ for some set $S \subseteq [u, \tau_u]$ of size
 720 $|S| = \frac{1}{12}h_2$.

721 2. Some $v \in \mathcal{U}(u)$ is promising. In this case there must be some $\frac{2}{3}h_2$ locations on $[u, v]$
 722 that point towards v . This corresponds to $M_{\text{up}}^S(u, v)$ for some $S \subseteq M_{\text{up}}^S([v, u])$ of size
 723 $|S| = \frac{2}{3}h_2$.

724 Define $\mathcal{C}(u) = \{S \subseteq [u, \tau_u] \mid |S| = \frac{1}{6}h_2\}$ and $\mathcal{D}(u) = \{S \subseteq [u, \tau_u] \mid |S| = \frac{1}{12}h_2\}$. Similarly
 725 define $\mathcal{E}(u) = \{(v, S) \mid v \in \mathcal{U}(u), S \subseteq [u, v], \text{ and } |S| = \frac{2}{3}h_2\}$. Combining Definition 16 with
 726 the previous paragraph, yields


$$\begin{aligned}
 727 \quad \{u \text{ is } h_2\text{-misleading}\} &\subseteq \{\tau_u \text{ is not promising}\} \cup \bigcup_{v \in \mathcal{U}(u)} \{v \text{ is promising}\} \\
 728 \quad &\subseteq \bigcup_{S \in \mathcal{C}(u)} M_{\text{sides}}^S(\tau_u, u) \cup \bigcup_{S \in \mathcal{D}(u)} M_{\text{up}}^S(\tau_u, u) \cup \bigcup_{(v, S) \in \mathcal{E}(u)} M_{\text{up}}^S(u, v). \\
 729 \quad &
 \end{aligned}$$

730 In fact, $\mathcal{E}(u)$ needs to be further decomposed. For each $v \in \mathcal{E}(u)$, let $k(v) = |[u, v] \cap$
 731 $[\sigma', \tau']|$. For each non-negative integer $k \geq 0$, let

$$732 \quad \mathcal{E}_k(u) = \{(v, S) \in \mathcal{E}(u) \mid k(v) = k\}.$$

733 Clearly, $\mathcal{E}(u) = \bigcup_{k=0}^{h_2} \mathcal{E}_k(u)$.

⁶ Here again we omit the $[\cdot]$.



■ **Figure 1** Different relative positions of u, τ_u and σ' . The path $[u, \tau_u]$ and different mistake patterns. In the left one mistakes (depicted as red stars) point outside of $[u, \tau_u]$, in the second they point towards u and in the third towards a nominee of $T_{h_2}(u)$, $v \in \mathcal{U}(u)$.

734 Using the union bound (Claim 21) as in Equation 4, the aforementioned decomposition
735 implies:

$$\begin{aligned}
 736 \quad \mathcal{Q}\left(\mathbf{A}_{loop}(T_{h_1}(\sigma')) \cap u \text{ is } h_2\text{-misleading}\right) &\leq \sum_{S \in \mathcal{C}(u)} \mathcal{Q}\left(\mathbf{A}_{loop}(T_{h_1}(\sigma')) \cap M_{sides}^S(\tau_u, u)\right) \\
 737 &+ \sum_{S \in \mathcal{D}(u)} \mathcal{Q}\left(\mathbf{A}_{loop}(T_{h_1}(\sigma')) \cap M_{up}^S(\tau_u, u)\right) \\
 738 &+ \sum_{k=0}^{h_2} \sum_{(v,S) \in \mathcal{E}_k(u)} \mathcal{Q}\left(\mathbf{A}_{loop}(T_{h_1}(\sigma')) \cap M_{up}^S(u, v)\right) \\
 739 &\hspace{15em} (6)
 \end{aligned}$$

740 To prove Lemma 15, our goal will be to show that each sum in the above equation is at
741 most $\mathcal{O}(\sqrt{\Delta}/\log n)$.

742 E.3 Analysing Atomic Expressions

743 To prove that each sum is indeed $\mathcal{O}(\sqrt{\Delta}/\log n)$ we use the following two lemmas (proved
744 in Appendix E.4), which encapsulate the core of this proof, namely, the resilience of \mathbf{A}_{loop}
745 to certain kinds of error patterns.

746 ► **Lemma 18.** Consider a tree T rooted at σ with treasure located at τ . Let $a, b \in T$ be two
747 nodes such that a is the closest one to τ out of the nodes in $[a, b]$. Then,

$$748 \quad \mathcal{Q}\left(\mathbf{A}_{loop} \mid M_{sides}^S(a, b)\right) = \mathcal{O}\left(D^4 \Delta^{\frac{|S|+1}{2}}\right).$$

750 ► **Lemma 19.** Consider a tree T rooted at σ with treasure located at τ . Let $a, b \in T$ be two
751 nodes such that a is the closest one to τ out of the nodes in $[a, b]$. Then,

$$752 \quad \mathcal{Q}\left(\mathbf{A}_{loop} \mid M_{up}^S(a, b)\right) = \mathcal{O}\left(D^4 \Delta^{K+\frac{1}{2}} 4^{|S|}\right),$$

754 where $K = |S \cap [\sigma, \tau]|$.

755 As a first step to bounding the three sums of Equation (6), note that:

$$756 \quad |\mathcal{C}(u)| \leq 2^{h_2} \tag{7}$$

$$757 \quad |\mathcal{D}(u)| \leq 2^{h_2}, \tag{8}$$

$$758 \quad |\mathcal{E}_k(u)| \leq 2^{h_2} \Delta^{h_2-k}. \tag{9}$$

54:22 Searching a Tree with Permanently Noisy Advice

760 Indeed, $\mathcal{C}(u), \mathcal{D}(u)$ are sets of subsets of a path of length h_2 . For the last term, the number
 761 of $v \in \mathcal{U}(u)$ at distance h_2 from u for which $k(v) = k$ is bounded by Δ^{h_2-k} . Now the three
 762 sums:

763 1. $S \in \mathcal{C}(u)$, so $S \subseteq [u, \tau_u]$ and $|S| = \frac{1}{6}h_2$, and τ_u is the closest to τ of all the nodes on the
 764 path. By Lemma 18,

$$765 \quad \mathcal{Q}(A_{loop}(T_{h_1}(\sigma')) \mid M_{sides}^S(\tau_u, u)) = \mathcal{O}\left(h_1^4 \Delta^{\frac{|S|+1}{2}}\right).$$

766 According to Claim 17,

$$767 \quad \mathbb{P}(M_{sides}^S(\tau_u, u)) \leq q^{|S|}.$$

768 Combining these bounds and (7) yields

$$769 \quad \sum_{S \in \mathcal{C}(u)} \mathcal{Q}(A_{loop}(T_{h_1}(\sigma')) \cap M_{sides}^S(\tau_u, u)) = \mathcal{O}\left(2^{h_2} \cdot q^{|S|} \cdot h_1^4 \Delta^{\frac{|S|+1}{2}}\right) \\ 770 \quad \quad \quad = \mathcal{O}\left(\sqrt{\Delta} \cdot 2^{h_2} \cdot c^{|S|} \cdot h_1^4\right),$$

771 because $q < c/\sqrt{\Delta}$. Recall that $h_1 = \kappa_1 \log n$, $h_2 = \kappa_2 \log \log n$, and $|S| = \frac{1}{6}h_2$. κ_1 was
 772 already set to be some constant. Taking a large enough κ_2 and a small enough c , both
 773 independent of n , the previous expression is $\mathcal{O}(\sqrt{\Delta}/\log n)$ as needed.

774 2. $S \in \mathcal{D}(u)$, so $S \subseteq [u, \tau_u]$ and $|S| = \frac{1}{12}h_2$. Therefore, by Lemma 19,

$$775 \quad \mathcal{Q}(A_{loop}(T_{h_1}(\sigma')) \mid M_{up}^S(\tau_u, u)) = \mathcal{O}\left(h_1^4 \Delta^{|S|+\frac{1}{2}} 2^{h_2}\right).$$

776 Because $K \leq |S|$ and $4^{|S|} \leq 2^{h_2}$. Combined with Claim 17 and (8):

$$777 \quad \sum_{S \in \mathcal{D}(u)} \mathcal{Q}(A_{loop}(T_{h_1}(\sigma')) \cap M_{up}^S(\tau_u, u)) = \mathcal{O}\left(2^{h_2} \cdot \left(\frac{q}{\Delta}\right)^{|S|} \cdot h_1^4 \Delta^{|S|+\frac{1}{2}} 2^{h_2}\right) \\ 778 \quad \quad \quad = \mathcal{O}\left(\sqrt{\Delta} \cdot 4^{h_2} \cdot q^{|S|} h_1^4\right)$$

779 Again, since $|S| = \frac{1}{12}h_2$, then c and κ_2 can be chosen so that this is $\mathcal{O}(\sqrt{\Delta}/\log n)$.

780 3. $(v, S) \in \mathcal{E}_k(u)$, where $v \in \mathcal{U}(u)$, $S \subseteq [u, v]$, and $|S| = \frac{2}{3}h_2$. Also, $|[u, v] \cap [\sigma', \tau']| = k$,
 781 and so $|S \cap [\sigma', \tau']| \leq k$. As $v \in \mathcal{U}(u)$, then u is the closest to treasure of the vertices on
 782 $[u, v]$. By Lemma 19,

$$783 \quad \mathcal{Q}(A_{loop}(T_{h_1}(\sigma')) \mid M_{up}^S(u, v)) = \mathcal{O}\left(h_1^4 \Delta^{k+\frac{1}{2}} 4^{h_2}\right)$$

784 Combined with (9) and Claim 17:


$$785 \quad \sum_{k=0}^{h_2} \sum_{(v, S) \in \mathcal{E}_k(u)} \mathcal{Q}(A_{loop}(T_{h_1}(\sigma')) \cap M_{up}^S(u, v)) = \mathcal{O}\left(\sum_{k \leq h_2} 2^{h_2} \Delta^{h_2-k} \cdot \left(\frac{q}{\Delta}\right)^{\frac{2}{3}h_2} h_1^4 \cdot \Delta^{k+\frac{1}{2}} 4^{h_2}\right) \\ 786 \quad \quad \quad = \mathcal{O}\left(\sqrt{\Delta} \cdot h_2 8^{h_2} h_1^4 (q^2 \Delta)^{\frac{1}{3}h_2}\right). \\ 787 \quad \quad \quad = \mathcal{O}\left(\sqrt{\Delta} \cdot h_2 8^{h_2} h_1^4 \cdot c^{\frac{1}{3}h_2}\right).$$

788 Similarly to the two previous sums, this whole expression can be made as small as
 789 $\mathcal{O}(\sqrt{\Delta}/\log n)$.

790 Note that we assumed for simplicity that u, τ_u and v are all inside $T_{h_1}(\sigma')$. If they are
 791 not, we take nodes that are the closest to them on this subtree, which can only improve the
 792 bounds.

793 This concludes the proof of Lemma 15 and hence completes the proof of the second item
 794 in Theorem 3.

798 **E.4 The Lemmas About the Resilience of A_{loop}**



799 **Figure 2** Notations introduced in the proof of Lemma 18 and 19. Points of S are depicted
 800 in red. On the figure $n(z_1) = 0, m(z_1) = 0, n(z_2) = 1, m(z_2) = 0, n(z_3) = 3, m(z_3) = 2$ and
 $n(z_4) = 3, m(z_4) = 0$.

799 **► Lemma 18 (restated).** Consider a tree T rooted at σ with treasure located at τ . Let
 800 $a, b \in T$ be two nodes such that a is the closest one to τ out of the nodes in $[a, b]$. Then,

801
$$\mathcal{Q}(A_{loop} \mid M_{sides}^S(a, b)) = \mathcal{O}\left(D^4 \Delta^{\frac{|S|+1}{2}}\right).$$

 802

803 **Proof.** As in the proof of Lemma 14, we break the number of queries made by A_{loop} condi-
 804 tioning on $M_{sides}^S(a, b)$ into a sum of random variables Q_j which correspond to the number
 805 of queries needed to discover the j -th node on the path $[\sigma, \tau]$ once the $(j - 1)$ -th was dis-
 806 covered. Each Q_j is bounded above by D times the expected number of competitors who
 807 beat this node. This is because each phase takes D steps, and only a subset of these
 808 nodes will actually be checked by A_{loop} on layer j before trying the correct node. Hence, a
 809 bound on the number of competitors who beat a given $u \in [\sigma, \tau]$ translates to a bound on
 810 $\mathcal{Q}(A_{loop}(T) \mid M_{sides}^S(\tau, \sigma))$ by multiplying it by D^2 .

811 Let u be such a node, and z be a competitor of u (i.e., it is at the same level as u).
 812 Define $k(z)$ as half the distance between z and u , and denote $n(z) := |S \cap [\sigma, \tau] \cap [u, z]|$ and
 813 $m(z) := |S \cap [\sigma, \tau]^c \cap [u, z]|$. See Figure 2 for illustration.

814 First note, that since all advice of $S \subseteq [a, b]$ points sideways w.r.t. to this path, then any
 815 of it which is on the path $[u, z]$ also points sideways w.r.t. it, except possibly at one point,
 816 which may actually point towards z . The different cases are seen in Figure 2:

- 817 ■ For z_1 , the paths do not intersect at all.
- 818 ■ In the case of z_2 , if the least common ancestor of u and z_2 was a member of S , then it
 819 could point towards z_2 , and that would be sideways w.r.t. $[a, b]$.
- 820 ■ For z_3 , the least common ancestor of b and z_3 could point towards z_3 .
- 821 ■ For z_4 , the least common ancestor of a and b could point towards z_4 .
- 822 ■ There is also the case where $a \notin [\sigma, \tau]$, which is not depicted on Figure 2. The analysis
 823 remains valid, and in fact $n(z) = 0$ for all competitors z .

824 This one special vertex, if it exists, conditioned on that it points sideways w.r.t. $[a, b]$, points
825 towards z with probability $1/(\Delta - 2)$, and otherwise points sideways w.r.t. $[u, z]$.

826 Fix k, n and m , and consider a competitor z such that $k(z) = k$, $n(z) = n$, and $m(z) = m$.
827 On the path $[u, z]$ the number of advice remaining to be sampled is $2k - n - m - 1$. By
828 Lemma 20:

$$\begin{aligned}
829 \quad \mathbb{P}(z \text{ beats } u) &\leq \left(1 - \frac{1}{\Delta - 2}\right) \mathbb{P}\left(\sum_{s=1}^{2k-1-n-m} X_s \geq 0\right) + \frac{1}{\Delta - 2} \mathbb{P}\left(\sum_{s=1}^{2k-1-n-m} X_s \geq -1\right) \\
830 \quad &= \left(\frac{1}{\sqrt{\Delta}}\right)^{2k-1-n-m} + \frac{4}{\Delta - 2} \left(\frac{1}{\sqrt{\Delta}}\right)^{2k-2-n-m} \\
831 \quad &= \left(1 + \frac{4\sqrt{\Delta}}{\Delta - 2}\right) \left(\frac{1}{\sqrt{\Delta}}\right)^{2k-1-n-m} \leq 7 \cdot \left(\frac{1}{\sqrt{\Delta}}\right)^{2k-1-n-m}, \\
832 \quad &
\end{aligned}$$

833 as $\Delta \geq 3$. For fixed k, n, m there are at most Δ^{k-m} nodes z with $k(z) = k$ and $m(z) = m$.
834 Also, for each such node, $n + m \leq 2k$. Hence, the total expected number of competitors
835 that beat u is at most:

$$836 \quad \sum_{k \leq D, n+m \leq 2k} \Delta^{k-m} \cdot 7 \left(\frac{1}{\sqrt{\Delta}}\right)^{2k-1-n-m}$$

838 For each choice of k there is exactly one corresponding value of n . This n satisfies $n \leq |S|$.
839 There are also at most D choices for m . Thus, the above is at most

$$840 \quad 7 \cdot \sum_{k \leq D, n+m \leq 2k} \Delta^{(n+1-m)/2} = \mathcal{O}\left(D^2 \Delta^{(|S|+1)/2}\right).$$

842 ◀

843 ► **Lemma 19 (restated).** Consider a tree T rooted at σ with treasure located at τ . Let
844 $a, b \in T$ be two nodes such that a is the closest one to τ out of the nodes in $[a, b]$. Then,

$$845 \quad \mathcal{Q}(A_{loop} \mid M_{up}^S(a, b)) = \mathcal{O}\left(D^4 \Delta^{K+\frac{1}{2}} 4^{|S|}\right),$$

846 where $K = |S \cap [\sigma, \tau]|$.

848 **Proof.** Let u be a node on the path $[\sigma, \tau]$. Our aim is to show that the expected number of
849 competitors of u that beat it is $\mathcal{O}(D^2 \Delta^{K+\frac{1}{2}} 4^{|S|})$.

850 As in the proof of Lemma 18, let z be a competitor of u . Define $k(z)$ as half the
851 distance between z and u , namely $k(z) := d(z, u)/2$. Denote $n(z) := |S \cap [\sigma, \tau] \cap [u, z]|$, and
852 $m(z) := |S \cap [\sigma, \tau]^c \cap [u, z]|$.

853 Fixing k, n and m , take a competitor z such that $k(z) = k$, $n(z) = n$, and $k(z) = k$. The
854 probability that such a z beats u is

$$855 \quad \mathbb{P}\left(\sum_{s=1}^{2k-1-n-m} X_s \geq -n - m\right) \leq 4^{n+m} \Delta^{n+m-k+\frac{1}{2}},$$

857 by Lemma 20. There are at most Δ^{k-m} such nodes z . We bound the probability that each
858 of these nodes z beats the treasure using the trivial bound 1 or the one above, depending
859 on whether $n + m \leq k$ or $n + m > k$. Hence the total expected number of competitors of u
860 who beat it is at most

$$861 \quad \sum_{k \leq D, n+m \leq k} \Delta^{k-m} \cdot 4^{n+m} \Delta^{n+m-k+\frac{1}{2}} + \sum_{k \leq D, n+m > k} \Delta^{k-m}.$$

862

863 Since $n + m \leq |S|$, and $n \leq K$, the first term is at most:

$$864 \quad 4^{|S|} \sum_{k \leq D, n+m \leq k} \Delta^{K+\frac{1}{2}} \leq 4^{|S|} \cdot D^2 \cdot \Delta^{K+\frac{1}{2}},$$

865 where we used the fact that there are most D distinct values for k and D distinct values for
866 m , while there is only one choice of n for each k . As for the second term, since $n + m > k$,
867 then it is at most:

$$868 \quad \sum_{k \leq D, n+m > k} \Delta^n \leq \sum_{k \leq D, n+m > k} \Delta^K \leq \sum_{k, m \leq D} \Delta^K \leq D^2 \cdot \Delta^K,$$

869 concluding the proof. ◀

870 **F** Complementary Proofs

871 **F.1** Another Large Deviation Estimate

872 Here, we introduce another large deviation estimate used for the analysis of the query
873 algorithm for regular trees. It gives better results for large h , yet works only for identical
874 random variables, and so suits regular trees, unlike Lemma 7.

875 ► **Lemma 20.** *Consider random variables X_i taking values $\{-1, 0, 1\}$ with respective prob-*
876 *abilities $(1 - q, q(1 - \frac{2}{\Delta}), \frac{q}{\Delta})$. If $q < \frac{c}{\sqrt{\Delta}}$ where $c < 1/64$, then for all $0 \leq h \leq l$,*

$$877 \quad \mathbb{P}\left(\sum_{i=1}^{\ell} X_i \geq -h\right) \leq (4\sqrt{\Delta})^h \Delta^{-\ell/2}$$

878 **Proof.** Assume $\sum_{i=1}^{\ell} X_i \geq -h$. Denote by $j := \{i \mid X_i = 1\}$. As the number of -1 's is at
879 least h , then $j \leq (\ell - h)/2$. There must also be at least $\ell - h - 2j$ zeros amongst what
880 remains, otherwise the sum is less than $-h$. Using a union bound over the value of j and
881 the locations of the ones and zeros we get:

$$882 \quad \mathbb{P}\left(\sum_{i=1}^{\ell} X_i \geq -h\right) \leq \sum_{j=0}^{\frac{\ell-h}{2}} \binom{\ell}{j} \binom{\ell-j}{\ell-h-2j} \left(q(1 - \frac{2}{\Delta})\right)^{\ell-h-2j} \left(\frac{q}{\Delta}\right)^j$$

$$883 \quad \leq \sum_{j=0}^{\frac{\ell-h}{2}} \binom{\ell}{j} \binom{\ell-j}{\ell-h-2j} q^{\ell-h-2j} \left(\frac{q}{\Delta}\right)^j$$

$$884 \quad \leq 3^{\ell} \sum_{j=0}^{\frac{\ell-h}{2}} q^{\ell-h-2j} \left(\frac{q}{\Delta}\right)^j \leq 3^{\ell} \cdot \frac{\ell}{2} \left(q^{\ell-h} + \left(\frac{q}{\Delta}\right)^{\frac{\ell-h}{2}}\right).$$

$$885$$

886 The last step uses the bound $\sum_{j=1}^N \rho^k \leq N \cdot (\rho + \rho^N)$. Note that $x/2 < (4/3)^x$ always, and
887 assigning $q < c/\sqrt{\Delta}$, this is at most:

$$888 \quad 4^{\ell} \frac{1}{\sqrt{\Delta}^{\ell-h}} \left(c^{\ell-h} + \left(\frac{\sqrt{c}}{\Delta^{3/2}}\right)^{\ell-h}\right) \leq 4^{\ell} \frac{1}{\sqrt{\Delta}^{\ell-h}} \left(c^{\ell-h} + \sqrt{c}^{\ell-h}\right) \leq 4^{\ell} \left(\frac{2\sqrt{c}}{\sqrt{\Delta}}\right)^{\ell-h}.$$

$$889$$

890 Since $c < 1/64$, then $2\sqrt{c} \leq 1/4$ which means that this is at most

$$891 \quad 4^h \left(\frac{1}{\sqrt{\Delta}}\right)^{\ell-h},$$

892 giving the desired bound. ◀

893 **F.2 Algorithm A_{loop} without Conditioning**

894 ► **Lemma 14 (restated).** Consider a (not necessarily complete) Δ -ary tree. Then $\mathcal{Q}(A_{loop}) =$
895 $\mathcal{O}(D^3\sqrt{\Delta})$.

896 **Proof.** Denote by $N_{layer}(u)$ the number of nodes on the same depth as u which have more
897 discovered arrows than u pointing to them. This definition is central because of the following
898 observation. The number of moves needed before finding u_{i+1} once u_i has been found is less
899 than $\mathcal{O}(DN_{layer}(u_i))$. Indeed, once u_i is discovered, only a subset of the nodes which have
900 more arrows pointing to them than u_{i+1} on layer $i+1$ are tried before u_{i+1} (at step (2) in
901 the pseudocode description). The loop over the levels (at step (1)) induces a multiplicative
902 factor of $\mathcal{O}(D)$.

903 Using linearity of expectation, it only remains to estimate $\mathbb{E}(N_{layer}(u_i))$ where u_i is the
904 ancestor of the treasure at depth $i \leq d$. There are at most Δ^ℓ nodes on layer i at distance
905 $2\ell - 1$ from u_i , for any $1 \leq \ell \leq i$. Moreover the probability that each of these nodes has
906 more arrows pointing towards it than u_i exactly corresponds to $\mathbb{P}\left(\sum_{j=1}^{2\ell-1} X_j \geq 0\right)$, with the
907 notations of Lemma 20.

908 Indeed, when comparing the amount of advice pointing to two different nodes u and v ,
909 only the nodes of $\langle u, v \rangle$ matter.

910 When estimating the probability that v beats u , each random variable X_j has to be
911 interpreted as taking value $+1$ if the advice points towards v , -1 if it points towards u , and
912 0 if it points neither to u nor v . In the case that $u = u_j$ and v is another node on layer j ,
913 these events happen respectively with probability q/Δ , $1 - q + q/\Delta$ and $q(1 - 2\frac{1}{\Delta})$.

914 This means that for each i ,

$$915 \quad \mathbb{E}(N_{layer}(u_i)) \leq \sum_{\ell=1}^i \mathbb{P}\left(\sum_{j=1}^{2\ell-1} X_j \geq 0\right) \Delta^\ell \leq \sum_{\ell=1}^d \mathbb{P}\left(\sum_{j=1}^{2\ell-1} X_j \geq 0\right) \Delta^\ell.$$

916
917 By Lemma 20 this is at most

$$918 \quad \mathcal{O}\left(\sum_{\ell=1}^d \Delta^{-\ell+\frac{1}{2}} \cdot \Delta^\ell\right) = \mathcal{O}(D\sqrt{\Delta}) = \mathcal{O}(D\sqrt{\Delta}).$$

919 ◀

920 **F.3 Special Form of Union Bound**

921 ► **Claim 21.** Let A be an event that can be decomposed as the union of events $(A_i)_{i \in I}$,
922 $A \subseteq \bigcup_{i \in I} A_i$. Let X be a random variable.

$$923 \quad \mathbb{E}(X | A)\mathbb{P}(A) \leq \sum_i \mathbb{E}(X | A_i)\mathbb{P}(A_i)$$

924
925 **Proof.** We denote by $\chi(B)$ the indicator function of event B . Then

$$926 \quad \mathbb{E}(X | A)\mathbb{P}(A) = \mathbb{E}(X \cdot \chi(A)) \leq \mathbb{E}\left(X \cdot \chi\left(\bigcup_i A_i\right)\right) \leq \mathbb{E}\left(X \cdot \sum_i \chi(A_i)\right)$$

$$927 \quad = \sum_i \mathbb{E}(X \cdot \chi(A_i)) = \sum_i \mathbb{E}(X | A_i)\mathbb{P}(A_i).$$

928
929 Where we used the union bound in the form $\chi(\bigcup_i A_i) \leq \sum_i \chi A_i$ and then linearity of
930 expectation. ◀

G

 Lower bounds

G.1 An Exponential Lower Bound Above the Threshold: Proof of Lemma 10

For the lower bound, assume the algorithm is given the advice \mathbf{adv} for all the internal nodes for free. By Yao's principle, instead of taking the worst case placement of the treasure for a randomized algorithm, we obtain a lower bound by considering only deterministic algorithms when the treasure is placed uniformly at random at one of the leaves.

In this simplified setting, an optimal algorithm can be described explicitly: It sorts the leaves according to $\mathbb{P}(\cdot | \mathbf{adv})$ (Claim 22) and tries them in this order. This order in fact corresponds to ranking nodes by how many arrows point to them (Claim 23). The expected number of nodes which are higher than the treasure in this ordering is therefore a lower bound for this algorithm, and thus for all algorithms.

Let \mathcal{L} be the set of leaves. For a given leaf $u \in \mathcal{L}$ and an advice configuration \mathbf{adv} , let $C(\mathbf{A}, \mathbf{adv}, u)$ be the cost (number of queries) of \mathbf{A} when the advice is equal to \mathbf{adv} and the treasure is located at u . We also define the cost $C(\mathbf{A}, u)$ of an algorithm \mathbf{A} when the treasure τ is located at u to be the expected cost of \mathbf{A} before finding τ where the expectation is over advice setting. That is:

$$C(\mathbf{A}, u) = \sum_{\mathbf{adv}} C(\mathbf{A}, \mathbf{adv}, u) \mathbb{P}(\mathbf{adv} | u).$$

In our setting, the expected number of queries of \mathbf{A} is:

$$C(\mathbf{A}) = \sum_{u \in \mathcal{L}} \mathbb{P}(u) \sum_{\mathbf{adv}} C(\mathbf{A}, \mathbf{adv}, u) \mathbb{P}(\mathbf{adv} | u).$$

► **Claim 22.** The algorithm \mathbf{A} that tries the locations u in the order given by $\mathbb{P}(u | \mathbf{adv})$, i.e., the most likely u is tried first and the least likely tried last, minimizes $C(\mathbf{A})$.

Proof. We can write

$$C(\mathbf{A}) = \sum_{\mathbf{adv}} \mathbb{P}(\mathbf{adv}) \sum_{u \in \mathcal{L}} C(\mathbf{A}, \mathbf{adv}, u) \mathbb{P}(u | \mathbf{adv}),$$

where it is understood that $\mathbb{P}(\mathbf{adv})$ is the marginal of $\mathbb{P}(\mathbf{adv}, u)$ with respect to the advice. Note that the term $\mathbb{P}(u | \mathbf{adv})$, standing for the probability of u holding the treasure given that the advice configuration is \mathbf{adv} , is only defined because we assume the treasure is placed according to a known distribution (uniform in our case). For a fixed advice setting \mathbf{adv} , it follows from the *rearrangement inequality* that $\sum_{u \in \mathcal{L}} C(\mathbf{A}, \mathbf{adv}, u) \mathbb{P}(u | \mathbf{adv})$ is minimized when $C(\mathbf{A}, \mathbf{adv}, u)$ and $\mathbb{P}(u | \mathbf{adv})$ are sorted in the same order with respect to u . This corresponds to algorithm \mathbf{A} trying the locations u in the order given by $\mathbb{P}(u | \mathbf{adv})$, which is exactly the statement of the claim. Hence, since we assume that all advice is known, the algorithm we have just described is feasible, and, in fact, optimal. Moreover, its query complexity is at least 1 plus the expected number of nodes which are strictly more likely than the treasure, where the expectation is taken over the randomness of the advice. ◀

It only remains to check that a node u is more likely than τ given an advice setting \mathbf{adv} iff more arrows point to u than τ . This will conclude the proof of Lemma 10 and hence of the exponential lower bound in Theorem 4.

► **Claim 23.** For two leaves $u, v \in \mathcal{L}$, and advice configuration \mathbf{adv} , $\mathbb{P}(u | \mathbf{adv}) > \mathbb{P}(v | \mathbf{adv})$ if and only if there is more advice pointing towards u than advice pointing towards v .

964 **Proof.** Recall that, by definition of the model

$$965 \quad \mathbb{P}(\text{adv} \mid \tau = u) = \left(p + \frac{q}{\Delta}\right)^{|\overrightarrow{\text{adv}}(u)|} \left(q\left(1 - \frac{1}{\Delta}\right)\right)^{|\overleftarrow{\text{adv}}(u)|},$$

966 In our regime it will always be the case that $p + \frac{q}{\Delta} > q\left(1 - \frac{1}{\Delta}\right)$, simply because we assume
967 $q < p$. Hence $\mathbb{P}(\text{adv} \mid \tau = u)$ is an increasing function of $|\overrightarrow{\text{adv}}(u)|$.

968 Since τ is placed uniformly at random, it follows from Bayes rule that $\mathbb{P}(\text{adv} \mid \tau = u) \propto$
969 $\mathbb{P}(\tau = u \mid \text{adv})$. The symbol \propto indicates that we omit the renormalizing factor. Hence, we
970 obtain that $\mathbb{P}(\tau = u \mid \text{adv}) > \mathbb{P}(\tau = v \mid \text{adv})$ if and only if $|\overrightarrow{\text{adv}}(u)| > |\overrightarrow{\text{adv}}(v)|$. ◀

971 G.2 A Lower Bound for the Move Complexity

972 ▶ **Observation 24.** For any Δ and d , there exists a tree of depth d and maximal degree
973 at most Δ for which any search algorithm A has move complexity $\mathcal{M}(A) = \Omega(dq\Delta)$. In
974 particular, when $q \sim 1/\sqrt{\Delta}$, we have $\mathcal{M}(A) = \Omega(d\sqrt{\Delta})$.

975 **Proof.** To see why the observation holds consider the caterpillar tree, composed of a path
976 of length n/Δ with each of its nodes being the center of a star graph of degree Δ . Assume
977 that the agent starts at one of the end sides of the path and the treasure at distance d
978 on the caterpillar spine. Recall that we assume that the algorithm does not know the tree
979 structure. In expectation, $\Omega(dq)$ nodes will point at an incorrect neighbor, and to pass from
980 any of those to the next node on the path, will require the agent to perform $\Omega(\Delta)$ trials in
981 expectation. ◀

982 G.3 Proof of Observation 11

983 A randomized strategy may be viewed as a convex combination of deterministic strategies.
984 The performance of a randomized strategy is thus a linear combination of the performance
985 of deterministic ones. Hence, it suffices to focus on deterministic strategies.

986 Since the treasure location is uniform over the k objects, a deterministic strategy finds
987 it after exactly i attempts with probability $1/k$ for any $i \leq k$. In other words, the number of
988 objects that are tried is distributed as a uniform random variable in $[1, k]$. Such a random
989 variable has mean $1/k \sum_{i=1}^k i = (k+1)/2$.

990 H Memoryless Algorithms and the Semi-Adversarial Model

991 In this section we present our results on the memoryless algorithms described in the in-
992 troduction. As mentioned, such algorithms can perform well also in a more difficult semi-
993 adversarial setting. Before we present these algorithms let us first describe formally the
994 semi-adversarial variant.

995 ▶ **Definition 25** (The Semi-Adversarial Model). As in the purely-probabilistic Noisy Advice
996 Model, each node is chosen to be *faulty* with probability q , and otherwise it is *sound*.
997 Also, similarly to the original model, a sound vertex always points at its correct neighbors.
998 However, in the semi-adversarial model, a faulty node u no longer points at a neighbor chosen
999 uniformly at random, and instead, the neighbor w which such a node points at is chosen by
1000 an adversary. Importantly, for each node u , the adversary must specify its potentially faulty
1001 advice w , before it is known which nodes will be faulty. In other words, first, the adversary
1002 specifies the faulty advice w for each node u , and then the environment samples which node
1003 is faulty and which is sound.

1004 H.1 Lower Bound in the Semi-Adversarial Variant

1005 The following result implies that if $q > 1/\Delta$ then any algorithm must have exponential query
1006 and move complexity in the depth D (or polynomial in n).

1007 ► **Theorem 26.** *Consider an algorithm in the semi-adversarial model. On the complete
1008 Δ -ary tree of depth D , the expected number of queries to find the treasure is $\Omega((q\Delta)^D)$.
1009 The lower bound holds even if the algorithm has access to the advice of all internal nodes in
1010 the tree.*

1011 **Proof.** Consider the complete Δ -ary tree and assume that the treasure is located at a leaf.
1012 The adversary behaves as follows. For any advice it gets a chance to manipulate, it would
1013 always make it point towards the root. With probability q^D the adversary gets to choose
1014 all the advice on the path between the root and the treasure. Any other advice points
1015 towards the root as well (either because it was correct to begin with or because it was set
1016 by the adversary). Hence with probability q^D the tree that the algorithm sees is the same
1017 regardless of the position of the treasure. It follows from Observation 11 that the time to
1018 find the treasure can only be linear in the number of leaves which is $\Omega(\Delta^D)$. ◀

1019 H.2 Probabilistic Following Algorithms

1020 Recall that a *Probabilistic Following (PF)* algorithm is specified by a *listening* parameter
1021 $\lambda \in (0, 1)$. At each step, the algorithm “listens” to the advice with probability λ and takes
1022 a uniform random step otherwise. The first item in the next theorem states that if the noise
1023 parameter is smaller than c/Δ for some small enough constant $0 < c < 1$, then there exists
1024 a listening parameter λ for which Algorithm PF achieves $\mathcal{O}(d)$ move complexity. Moreover,
1025 this result holds also in the semi-adversarial model. Hence, together with Theorem 26,
1026 it implies that in order to achieve efficient search, the noise parameter threshold for the
1027 semi-adversarial model is $\Theta(1/\Delta)$.

1028 ► **Theorem 27. 1.** *Assume that for every u , $q_u < 1/(10\Delta_u)$. Then PF with parameter $\lambda \in$
1029 $[0.7, 0.8]$ finds the treasure in less than $100d$ expected steps, even in the semi-adversarial
1030 setting.*

1031 **2.** *Consider the complete Δ -ary tree and assume that $q > 10/\Delta$. Then for any choice of
1032 λ the hitting time of the treasure by PF is exponential in the depth of the tree, even
1033 assuming the faulty advice is drawn at random.*

1034 **Proof.** Our plan is to show that the expected time to make one step in the correct direction
1035 is $\mathcal{O}(1)$, from any starting node. Conditioning on the advice setting, we make use of the
1036 Markov property to relate these elementary steps to the total travel time. The main delicate
1037 point in the proof stems from dealing with two different sources of randomness. Namely the
1038 randomness of the advice and that of the walk itself.

1039 It will be convenient to picture the tree as rooted at the target node τ . For any node u in
1040 the tree, we denote by u' the parent of u with respect to the treasure. With this convention,
1041 correct advice at a node u points at u' , while incorrect advice points at one of its children.
1042 The fact the walk moves on a tree means that for a given advice setting, the expected (over
1043 the walk) time it takes to reach u' from u can be written conveniently as a product of a
1044 variable involving the advice at u only and the advice on the set of u 's descendants (the two
1045 being independent).

1046 We denote by $t(u)$ the time it takes to reach node u . Manipulating average symbols such
1047 as \mathbb{E} requires extra care. Indeed, there are two sources of randomness, the first being the

1048 randomness used in drawing the advice and the second being the randomness used in the
 1049 walk itself. We write \mathbb{E} for averaging over the advice, while we use E_u to denote expectation
 1050 over the walk, conditioning on u being the starting node. As a remark, observe that $E_u(t(v))$
 1051 depends on the advice configuration, it is a random variable with respect to the advice, while
 1052 $\mathbb{E}E_u(t(v))$ really is just a number.

1053 The following is the central lemma of this section.

1054 ► **Lemma 28.** *Assume that for every vertex u , $q_u < 1/(10\Delta_u)$, and $\lambda \in [0.7, 0.8]$. Then for
 1055 all nodes u , $\mathbb{E}E_u t(u') \leq 100$. The result holds also in the semi-adversarial model.*

Let us now see how we can conclude the proof of the first item in Theorem 27, given the lemma. Consider a designated source σ . Let us denote by $\sigma = u_d, u_{d-1}, \dots, u_0 = \tau$ the nodes on the path from σ to τ . Let δ_i be the random variable indicating the time it takes to reach u_{i-1} after u_i has been visited for the first time. With these notations, the time to reach τ from σ is precisely $\sum_{i=1}^{d(\sigma, \tau)} \delta_i$. Hence, the expected time to reach τ from σ is, by linearity of expectation:

$$\sum_{i=1}^{d(\sigma, \tau)} \mathbb{E}[E_\sigma \delta_i].$$

1056 Conditioning on the advice setting, the process is a Markov chain and we may write

$$1057 \quad E_\sigma \delta_i = E_{u_i} t(u_{i-1}).$$

1059 Taking expectations over the advice (\mathbb{E}), under the assumptions of Lemma 28, it follows that
 1060 $\mathbb{E}(E_\sigma \delta_i) \leq 100$, for every $i \in [d(\sigma, \tau)]$. And this immediately implies a bound of $100 \cdot d(\sigma, \tau)$.

Proof of Lemma 28. We start with partitioning the nodes of the tree according to their distance from the root τ . More precisely, for $i = 1, 2, \dots, D$, where D is the depth of the tree, let

$$\mathcal{L}_i := \{u \in T : d(u, \tau) = i\}.$$

1061 The nodes in \mathcal{L}_i are referred to as *level- i* nodes. We treat the statement of the lemma
 1062 for nodes $u \in \mathcal{L}_i$ as an induction hypothesis, with i being the induction parameter. The
 1063 induction goes backwards, meaning we assume the assumption holds at level $i+1$ and show
 1064 it holds at level i . The case of the maximal level (base case for the induction) is easy since,
 1065 at a leaf the walk can only go up and so if u is a leaf $\mathbb{E}E_u(t(u')) = 1 < 100$.

1066 Assume now that $u \in \mathcal{L}_i$. We first condition on the advice setting. A priori, $E_u \tau(u')$
 1067 depends on the advice over the full tree, but in fact it is easy to see that only advice at
 1068 layers $\geq i$ matter. Recall from Markov Chain theory that an *excursion* to/from a point is
 1069 simply the part of the walk between two visits to the given point. We denote L_u the average
 1070 (over the walk only) length of an excursion from u to itself that does not go straight to u'
 1071 and we write N_u to denote the expected (over the walk only) number of excursions before
 1072 going to u' . We also refer to this number as a number of *attempts*. Note that N_u can be 0
 1073 if the walk goes directly to u' without any excursion. We decompose $t(u')$ in the following
 1074 standard way, using the Markov property

$$1075 \quad E_u t(u') = 1 + L_u \cdot N_u. \quad (10)$$

1077 Indeed the expectation $E_u t(u')$ can be seen as the expectation (over the walk) of $1 + \sum_{i=1}^T Y_i$
 1078 where the Y_i 's are the lengths of each excursion from u and T is the (random) number
 1079 of such excursions before hitting u' . The term $1 +$ accounts for the step from u to u' .
 1080 Note that $\{T \geq t\}$ is independent of Y_1, \dots, Y_t and so using Wald's identity we have that

1081 $E_u t(u') = 1 + E_u T \cdot E_u Y_1$. The term $E_u T$ is equal to N_u (by definition) while $E_u Y_1$ is equal
 1082 to L_u (by definition).

1083 We now want to average equality (10), which is only an average over the walk, by taking
 1084 the expectation over all advice in layers $\geq i$. To this aim, note that L_u can be written as

$$1085 \quad L_u = 1 + \sum_{v \neq u', v \sim u} p_{u,v} E_v t(u),$$

1086 where we write $u \sim v$ when u and v are neighbors in the tree and $p_{u,v}$ is the probability to
 1087 go straight from u to v given the advice setting. Note that, by assumption on the model,
 1088 $E_v t(u)$ depends on the advice at layers $\geq i + 1$ only, if we start at a node $v \in \mathcal{L}_{i+1}$, while
 1089 both $p_{u,v}$ and N_u depend only on the advice at layer $= i$ of the tree. This is true also in
 1090 the semi-adversarial model. Hence when we average, we can first average over layers $> i$ to
 1091 obtain, denoting $\mathbb{E}^{>i}$, the expectation over the layers $> i$,

$$1093 \quad \begin{aligned} \mathbb{E}^{>i} E_u t(u') &= 1 + \left(1 + \sum_{v \neq u', v \sim u} p_{u,v} \mathbb{E}^{>i} E_v t(u) \right) N_u, \\ 1094 \quad &= 1 + \left(1 + \sum_{v \neq u', v \sim u} p_{u,v} \mathbb{E} E_v t(u) \right) N_u. \end{aligned} \quad (11)$$

1095 and using the fact that,

$$1097 \quad \sum_{v \neq u'} p_{u,v} \leq 1, \quad (12)$$

1098 together with the induction assumption at rank $i + 1$, we obtain

$$1100 \quad \mathbb{E}^{>i} E_u t(u') \leq 1 + (1 + 100) N_u.$$

1101 From now on we replace 100 by a parameter $\kappa > 0$, for mere aesthetic reasons. Averaging
 1102 over the layer i of advice we obtain

$$1104 \quad \mathbb{E} E_u t(u') \leq 1 + (1 + \kappa) \mathbb{E} N_u.$$

1106 It only remains to analyse the term $\mathbb{E} N_u$. If the advice at u is correct, which happens with
 1107 probability $p_u = 1 - q_u$, then the number of attempts follows a (shifted by 1) geometric
 1108 law with parameter $\lambda + \frac{(1-\lambda)}{\Delta_u}$. In words, when the advice points to u' which happens with
 1109 probability at most 1, the walker can go to the correct node either because she listens to
 1110 the advice, which happens with probability λ , or because she did not listen, but still took
 1111 the right edge, which happens with probability $\frac{(1-\lambda)}{\Delta_u}$. Similarly, when the advice points to
 1112 a node $\neq u'$, which happens with probability at most q_u , then N_u follows a geometric law
 1113 (shifted by 1) with parameter $\frac{(1-\lambda)}{\Delta_u}$. The conclusion is that

$$1114 \quad \begin{aligned} \mathbb{E} N_u &\leq \left(\frac{1}{\lambda + \frac{(1-\lambda)}{\Delta_u}} - 1 \right) + q_u \left(\frac{\Delta_u}{1-\lambda} - 1 \right) \\ 1115 \quad &\leq \frac{1}{\lambda} - 1 + \frac{q_u \Delta_u}{1-\lambda} \end{aligned} \quad (13)$$

1116 And so it follows that

$$1118 \quad \mathbb{E} E_u t(u') \leq 1 + (1 + \kappa) \cdot \left(\frac{1}{\lambda} - 1 + \frac{q_u \Delta_u}{1-\lambda} \right)$$

54:32 Searching a Tree with Permanently Noisy Advice

1120 Hence if $q_u \Delta_u < 0.1$ and we choose $\lambda \in [0.7, 0.8]$ (for instance, we made no attempt in
 1121 optimizing these constants), we see that $\mathbb{E}N_u < 0.8$. This is because

$$1122 \quad \frac{1}{\lambda} - 1 + \frac{0.1}{1 - \lambda} \leq \frac{10}{7} - 1 + \frac{0.1}{1 - 0.8} < 0.93$$

1123 Hence it follows that

$$1124 \quad \mathbb{E}E_u t(u') \leq 1 + 0.93(1 + \kappa) < \kappa.$$

1126 The last inequality holds by choice of $\kappa = 100$. By our (backwards) induction, we have just
 1127 shown that, if $q < \frac{1}{10\Delta}$ and we set $\lambda \in [0.7, 0.8]$ then for all nodes u in the tree

$$1128 \quad \mathbb{E}E_u t(u') < 100.$$

1130 This concludes the proof of Lemma 28 and hence also of the first part of Theorem 27. ◀

1131 Let us explain how the lower bound in the second part of Theorem 27 is derived in the
 1132 case that $q\Delta > 10$. We assume we are in a complete Δ -ary tree under our usual uniform
 1133 noise model. With probability q there is fault at u and with probability $1 - \frac{1}{\Delta}$ the advice
 1134 does not point to u' . In this case, N_u follows a geometric law with parameter $\frac{1-\lambda}{\Delta}$. Hence

$$1135 \quad \mathbb{E}(N_u) \geq q\Delta \left(1 - \frac{1}{\Delta}\right) \frac{1}{1 - \lambda} - 1 \geq \frac{10(1 - \frac{1}{\Delta})}{1 - \lambda} - 1 \geq 10 \left(1 - \frac{1}{\Delta}\right) - 1 \geq 3,$$

1137 for any choice of λ , since $\Delta \geq 2$. We proceed very similarly, by induction, and use Equality
 1138 (11) together with the previous bound on $\mathbb{E}(N_u)$ to obtain that for any node on layer i , u
 1139 with parent u' , $\mathbb{E}E_u t(u') \geq 1 + 3 \min_{v \in \mathcal{L}_{i+1}} \mathbb{E}E_v t(v')$, so in particular $\min_{u \in \mathcal{L}_i} \mathbb{E}E_u t(u') \geq$
 1140 $1 + 3 \min_{v \in \mathcal{L}_{i+1}} \mathbb{E}E_v t(v')$. The expected hitting time of the target τ , even starting at one of
 1141 its children is therefore of order $\Omega(3^D)$. ◀

1142 ▶ **Remark.** Note that the proof uses crucially the tree structure and does not extend to
 1143 general graphs straightforwardly. Specifically, on a tree there is a single path from σ to τ
 1144 and so the points u_i are uniquely defined, they are not random. Moreover an excursion from
 1145 a node u at Layer i that does not visit its parent can only remain in layers $\geq i$. This was
 1146 used through the fact that $E_v t(u)$ depends only on the advice at layers $\geq i$, if we start at a
 1147 node $v \in \mathcal{L}_i$.