Brief Announcement: Energy Constrained Depth First Search

Shantanu Das[1]
LIS, Aix-Marseille University, University of Toulon, CNRS, Marseille, France
shantanu.das@lif.univ-mrs.fr

Dariusz Dereniowski[2]
Faculty of Electronics, Telecommunications and Informatics, Gdańsk University of Technology, Gdańsk, Poland
deren@eti.pg.edu.pl
https://orcid.org/0000-0003-4000-4818

Przemysław Uznański
Department of Computer Science, ETH Zürich, Zürich, Switzerland
przemyslaw.uznanski@inf.ethz.ch

Abstract

Depth first search is a natural algorithmic technique for constructing a closed route that visits all vertices of a graph. The length of such route equals, in an edge-weighted tree, twice the total weight of all edges of the tree and this is asymptotically optimal over all exploration strategies. This paper considers a variant of such search strategies where the length of each route is bounded by a positive integer B (e.g. due to limited energy resources of the searcher). The objective is to cover all the edges of a tree T using the minimum number of routes, each starting and ending at the root and each being of length at most B. To this end, we analyze the following natural greedy tree traversal process that is based on decomposing a depth first search traversal into a sequence of limited length routes. Given any arbitrary depth first search traversal R of the tree T, we cover R with routes $R_1, \ldots, R_l$, each of length at most B such that: R_i starts at the root, reaches directly the farthest point of R visited by R_{i-1}, then R_i continues along the path R as far as possible, and finally R_i returns to the root. We call the above algorithm *piecemeal-DFS* and we prove that it achieves the asymptotically minimal number of routes l, regardless of the choice of R. Our analysis also shows that the total length of the traversal (and thus the traversal time) of piecemeal-DFS is asymptotically minimum over all energy-constrained exploration strategies. The fact that R can be chosen arbitrarily means that the exploration strategy can be constructed in an online fashion when the input tree T is not known in advance. Each route R_i can be constructed without any knowledge of the yet unvisited part of T. Surprisingly, our results show that depth first search is efficient for energy constrained exploration of trees, even though it is known that the same does not hold for energy constrained exploration of arbitrary graphs.

2012 ACM Subject Classification Theory of computation → Graph algorithms analysis, Theory of computation → Online algorithms

Keywords and phrases DFS traversal, distributed algorithm, graph exploration, piecemeal exploration, online exploration

Digital Object Identifier 10.4230/LIPIcs.ICALP.2018.165

Related Version `https://arxiv.org/abs/1709.10146`

[1] Author partially supported by ANR project ANCOR (anr-14-CE36-0002-01)
[2] Author partially supported by National Science Centre (Poland) grant number 2015/17/B/ST6/01887.

1 Introduction

Consider a mobile robot (also called an *agent*) that needs to explore an initially unknown edge-weighted tree, where the weight of each edge is its length. Starting from a single vertex (the root), the robot must traverse all edges of and return to its initial location. Upon visiting a vertex v for the first time, the robot discovers the edges incident to v and can choose one of them to continue the exploration. Provided that the robot can remember the visited vertices and edges, a simple depth first search (DFS) is an efficient algorithm for exploring the tree, achieving the optimal cost of twice the sum of the lengths of edges in the tree. In a more interesting scenario, the robot has a limited source of energy (e.g. a battery) which allows it to traverse a path of length at most B (we say such a robot is *energy constrained*). Naturally, we assume that each vertex of the tree is at distance at most $B/2$ from the root, otherwise the tree cannot be fully explored. In this case, exploration is possible if the robot can recharge its battery whenever it returns back to the initial location. Thus, the exploration is a collection of routes of the robot, each of which starts and ends at the root, and has length at most B. We are interested in minimizing the number of such routes (i.e. the number of times the robot has to recharge) to completely explore the tree.

Related work. There exists extensive literature on graph traversal and exploration. We refer interested reader to works on several models that do not consider any energy limitation for the agents, including results on general graphs [22], trees [9, 17, 18, 19], lower bounds on exploration time [9, 13, 14, 20, 21], or exploration with little memory [1, 12].

The energy constrained exploration problem was first studied under the name of *Piecemeal Graph Exploration* [8], with the assumption that the route length $B \geq 2(1+\beta)r$, where r is the furthest distance from the starting vertex to other vertices, and $0 < \beta < 1$. That paper provided exploration algorithms for a special class of grid graphs with 'rectangular obstacles'. Awerbuch et al. [4] showed that, for general graphs, there exists an energy constrained exploration algorithm with a total cost of $O(m + n^{1+o(1)})$. This has been further improved (by an algorithm that is a combination of DFS and BFS) to $O(m + n \log^2 n)$ in [5]. Finally [15] provided an exploration algorithm for general unknown weighted graphs with total cost asymptotic to the sum of edge weights of the graph. Note that, as mentioned, all the above strategies require the length of each route to be strictly larger than the shortest path from the starting vertex to the farthest vertex. In other words, these algorithms fail in the extreme cases when the height of the explored tree (or the diameter of the graph) is equal to half of the energy budget, which seem to be the most challenging ones. The off-line version of the problem is NP-hard, see e.g. [18].

See [6, 11, 16] for works on the tree exploration model we consider, with one difference: each route, also of length at most B, starts at the root but may end at any vertex of the tree. Distributed algorithms for energy constrained agents have been a subject of recent investigation, see e.g. [2, 3, 7, 10].

2 Problem statement and DFS exploration

Let $T = (V(T), E(T), \omega\colon E(T) \to \mathbb{R}_+)$ be an edge-weighted tree with root r. We define a *route* $R = (v_0, v_1, \ldots, v_l)$ as a sequence of vertices that satisfies: (i) $\{v_i, v_{i+1}\} \in E(T)$ for each $i \in \{0, \ldots, l-1\}$, and (ii) $v_0 = v_l$ is the root r of T. Informally speaking, a route is a sequence of vertices forming a walk in T that starts and ends at the root. We define the *length* of R to be $\ell(R) = \sum_{i=1}^{l} \omega(\{v_{i-1}, v_i\})$. We say that a vertex v is *visited* by the route if $v = v_i$ for some $i \in \{1, \ldots, l\}$.

Given a tree T and a real number B, we say that $\mathcal{S} = (R_1, \ldots, R_k)$ is a *B-exploration strategy* for T if for each $i \in \{1, \ldots, k\}$, R_i is a route in T of length at most B, and each vertex of T is visited by some route in $\mathcal{S}$. We write $|\mathcal{S}|$ to refer to the number of routes in $\mathcal{S}$, $k = |\mathcal{S}|$. We formulate the combinatorial problem we study in this work as follows.

Energy Constrained Tree Exploration problem (ECTE)

> Given a real number $B > 1$ and an edge-weighted rooted tree T of height at most $B/2$ what is the minimum integer k such that there exists a B-exploration strategy that consists of k routes?

Our goal is to analyze a particular type of solution to this problem, namely, an exploration strategy that behaves like a depth first search traversal but adopted to the fact that route lengths are bounded by B. Let $R_{\rm DFS} = (v_0, v_1, \ldots, v_l)$ be a route in T that covers the tree T and performs a depth first search traversal of T. (Note that $R_{\rm DFS}$ is a route and thus we consider a depth first search traversal to have vertex repetitions.) For two vertices u and v of T, $d(u, v)$ denotes the distance between u and v understood as the sum of weights of the edges of the path connecting these vertices. We refer by $\mathrm{PDFS}(T) = (R_1, \ldots, R_k)$ (*Piecemeal Depth First Search*) to the following B-exploration strategy constructed iteratively for $i := 1, \ldots, k$:

(i) let $j_0 = 0$ i.e. $v_{j_0} = v_0 = r$,

(ii) R_i continues DFS exploration from where R_{i-1} *stopped making progress* (from the vertex $v_{j_{i-1}}$) as long as for currently visited v_p: $d(r, v_{j_{i-1}}) + \ell((v_{j_{i-1}}, v_{j_{i-1}+1}, \ldots, v_p)) + d(v_p, r) \leq B$,

(iii) furthest v_p (for $p \leq l$) that satisfies condition from ii is denoted as v_{j_i}, the vertex where R_i stopped making progress,

(iv) let $R_i = P_{i-1} \circ (v_{j_{i-1}}, v_{j_{i-1}+1}, \ldots, v_{j_i-1}, v_{j_i}) \circ P_i^R$, where P_{i-1} is the path from r to $v_{j_{i-1}}$, and P_i^R is the path from v_{j_i} to r.

Such a strategy $\mathrm{PDFS}(T)$ is called a *DFS B-exploration*.

We remark that different depth first search traversals $R_{\rm DFS}$ may result in different values of k (different number of routes) in the resulting DFS B-exploration, although for a particular choice of $R_{\rm DFS}$ the corresponding $\mathrm{PDFS}(T)$ is unique. We point out that our results stated below hold for an *arbitrary* choice of $R_{\rm DFS}$.

3 Our results

The following theorem provides the first main result of this work.

▶ **Theorem 1.** *Let T be a tree and let the longest path from the root to a leaf in T be at most $B/2$. It holds $|\mathrm{PDFS}(T)| \leq 10\,|\mathcal{R}|$, where $\mathcal{R}$ is a B-exploration strategy that consists of the minimum number of routes.*

The theorem refers to the number of routes in an exploration strategy. However, in order to analyze the behavior of $\mathrm{PDFS}(T)$, we introduce another parameter which turns out to be simpler to analyze. For any B-exploration strategy $\mathcal{S} = (R_1, \ldots, R_k)$ of T we will denote by $\xi(\mathcal{S})$ the *cost* of $\mathcal{S}$ defined as $\xi(\mathcal{S}) = \sum_{i=1}^{k} \ell(R_i)$. Then, $\mathrm{COPT}(T)$ is an optimal solution with respect to the cost, that is, a B-exploration strategy whose cost is minimum over all B-exploration strategies. Thus, in order to prove Theorem 1, we obtain, on route, the following second main result of our work.

▶ **Theorem 2.** *Let T be a tree and let $B/2$ be greater than or equal to the longest path from the root to a leaf in T. It holds $\xi(\mathrm{PDFS}(T)) \leq 10 \cdot \xi(\mathrm{COPT}(T))$.*

References

1 Christoph Ambühl, Leszek Gąsieniec, Andrzej Pelc, Tomasz Radzik, and Xiaohui Zhang. Tree exploration with logarithmic memory. *ACM Trans. Algorithms*, 7(2):17:1–17:21, 2011.

2 Julian Anaya, Jérémie Chalopin, Jurek Czyzowicz, Arnaud Labourel, Andrzej Pelc, and Yann Vaxès. Collecting information by power-aware mobile agents. In *Proceedings of Distributed Computing - 26th International Symposium, DISC 2012*, pages 46–60. Springer, 2012. doi:10.1007/978-3-642-33651-5_4.

3 Julian Anaya, Jérémie Chalopin, Jurek Czyzowicz, Arnaud Labourel, Andrzej Pelc, and Yann Vaxès. Convergecast and broadcast by power-aware mobile agents. *Algorithmica*, 74(1):117–155, 2016.

4 Baruch Awerbuch, Margrit Betke, Ronald L. Rivest, and Mona Singh. Piecemeal graph exploration by a mobile robot. *Inf. Comput.*, 152(2):155–172, 1999.

5 Baruch Awerbuch and Stephen G. Kobourov. Polylogarithmic-overhead piecemeal graph exploration. In *Proceedings of the Eleventh Annual Conference on Computational Learning Theory, COLT 1998*, pages 280–286. ACM, 1998. doi:10.1145/279943.279998.

6 Evangelos Bampas, Jérémie Chalopin, Shantanu Das, Jan Hackfeld, and Christina Karousatou. Maximal exploration of trees with energy-constrained agents. *CoRR*, abs/1802.06636, 2018. arXiv:1802.06636.

7 Andreas Bärtschi, Jérémie Chalopin, Shantanu Das, Yann Disser, Barbara Geissmann, Daniel Graf, Arnaud Labourel, and Matús Mihalák. Collaborative delivery with energy-constrained mobile robots. In *Proceedings of Structural Information and Communication Complexity - 23rd International Colloquium, SIROCCO 2016*, volume 9988 of *Lecture Notes in Computer Science*, pages 258–274, 2016. doi:10.1007/978-3-319-48314-6_17.

8 Margrit Betke, Ronald L. Rivest, and Mona Singh. Piecemeal learning of an unknown environment. *Machine Learning*, 18(2-3):231–254, 1995.

9 Peter Brass, Flavio Cabrera-Mora, Andrea Gasparri, and Jizhong Xiao. Multirobot tree and graph exploration. *IEEE Trans. Robotics*, 27(4):707–717, 2011.

10 Jurek Czyzowicz, Krzysztof Diks, Jean Moussi, and Wojciech Rytter. Communication problems for mobile agents exchanging energy. In *Proceedings of Structural Information and Communication Complexity - 23rd International Colloquium, SIROCCO 2016*, volume 9988 of *Lecture Notes in Computer Science*, pages 275–288, 2016. doi:10.1007/978-3-319-48314-6_18.

11 Shantanu Das, Dariusz Dereniowski, and Christina Karousatou. Collaborative exploration by energy-constrained mobile robots. In *Proceedings of Structural Information and Communication Complexity - 22nd International Colloquium, SIROCCO 2015*, volume 9439 of *Lecture Notes in Computer Science*, pages 357–369. Springer, 2015. doi:10.1007/978-3-319-25258-2_25.

12 Yann Disser, Jan Hackfeld, and Max Klimm. Undirected graph exploration with $\Theta(\log\log N)$ pebbles. In *Proceedings of the Twenty-seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016*, pages 25–39. SIAM, 2016. doi:10.1137/1.9781611974331.ch3.

13 Yann Disser, Frank Mousset, Andreas Noever, Nemanja Skoric, and Angelika Steger. A general lower bound for collaborative tree exploration. *CoRR*, abs/1610.01753, 2016.

14 Stefan Dobrev, Rastislav Královic, and Euripides Markou. Online graph exploration with advice. In *Proceedings of Structural Information and Communication Complexity - 19th International Colloquium, SIROCCO 2012*, volume 7355 of *Lecture Notes in Computer Science*, pages 267–278. Springer, 2012. doi:10.1007/978-3-642-31104-8_23.

15 Christian A. Duncan, Stephen G. Kobourov, and V. S. Anil Kumar. Optimal constrained graph exploration. *ACM Trans. Algorithms*, 2(3):380–402, 2006.

16 Mirosław Dynia, Mirosław Korzeniowski, and Christian Schindelhauer. Power-aware collective tree exploration. In *Proceedings of Architecture of Computing Systems - ARCS 2006, 19th International Conference*, volume 3894 of *Lecture Notes in Computer Science*, pages 341–351. Springer, 2006. doi:10.1007/11682127_24.

17 Mirosław Dynia, Jarosław Kutyłowski, Friedhelm Meyer auf der Heide, and Christian Schindelhauer. Smart robot teams exploring sparse trees. In *Proceedings of Mathematical Foundations of Computer Science, 31st International Symposium, MFCS 2006*, volume 4162 of *Lecture Notes in Computer Science*, pages 327–338. Springer, 2006. doi:10.1007/11821069_29.

18 Pierre Fraigniaud, Leszek Gąsieniec, Dariusz R. Kowalski, and Andrzej Pelc. Collective tree exploration. *Networks*, 48(3):166–177, 2006.

19 Yuya Higashikawa, Naoki Katoh, Stefan Langerman, and Shin-ichi Tanigawa. Online graph exploration algorithms for cycles and trees by multiple searchers. *J. Comb. Optim.*, 28(2):480–495, 2014.

20 Nicole Megow, Kurt Mehlhorn, and Pascal Schweitzer. Online graph exploration: New results on old and new algorithms. *Theor. Comput. Sci.*, 463:62–72, 2012.

21 Christian Ortolf and Christian Schindelhauer. A recursive approach to multi-robot exploration of trees. In *Proceedings of Structural Information and Communication Complexity - 21st International Colloquium, SIROCCO 2014*, volume 8576 of *Lecture Notes in Computer Science*, pages 343–354. Springer, 2014. doi:10.1007/978-3-319-09620-9_26.

22 Petrisor Panaite and Andrzej Pelc. Exploring unknown undirected graphs. *J. Algorithms*, 33(2):281–295, 1999.