

MIT Open Access Articles

Nearly Work-Efficient Parallel DFS in Undirected Graphs

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Ghaffari, Mohsen, Grunau, Christoph and Qu, Jiahao. 2023. "Nearly Work-Efficient Parallel DFS in Undirected Graphs."

As Published: <https://doi.org/10.1145/3558481.3591094>

Publisher: ACM|Proceedings of the 35th ACM Symposium on Parallelism in Algorithms and Architectures

Persistent URL: <https://hdl.handle.net/1721.1/150981>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution



Nearly Work-Efficient Parallel DFS in Undirected Graphs

Mohsen Ghaffari
MIT
Cambridge, MA, USA
ghaffari@mit.edu

Christoph Grunau
ETH Zurich
Zurich, Switzerland
cgrunau@ethz.ch

Jiahao Qu
ETH Zurich
Zurich, Switzerland
jiahao.qu@inf.ethz.ch

ABSTRACT

We present the first parallel depth-first search algorithm for undirected graphs that has near-linear work and sublinear depth. Concretely, in any n -node m -edge undirected graph, our algorithm computes a DFS in $\tilde{O}(\sqrt{n})$ depth and using $\tilde{O}(m+n)$ work. All prior work either required $\Omega(n)$ depth, and thus were essentially sequential, or needed a high poly(n) work and thus were far from being work-efficient.

CCS CONCEPTS

• Theory of computation → Parallel algorithms.

KEYWORDS

Parallel algorithms, depth-first search, work-efficient

ACM Reference Format:

Mohsen Ghaffari, Christoph Grunau, and Jiahao Qu. 2023. Nearly Work-Efficient Parallel DFS in Undirected Graphs. In *Proceedings of the 35th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '23)*, June 17–19, 2023, Orlando, FL, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3558481.3591094>

1 INTRODUCTION

Depth-first search (DFS) is one of the basic algorithmic techniques for graph problems, with a wide range of applications. In an n -node and m -edge undirected graph, a simple sequential algorithm computes a DFS in $O(m+n)$ time. This is often covered in introductory algorithmic courses. Unfortunately, the state-of-the-art parallel algorithms for computing a DFS require at least $\Omega(\sqrt{n})$ processors to run faster than this sequential algorithm. This significantly limits the applicability of these parallel DFS algorithms, because having $\Omega(\sqrt{n})$ or more processors is quite a high requirement for most plausible applications. In this paper, we describe a parallel DFS algorithm that runs faster than the sequential algorithm with just poly($\log n$) processors. Indeed, so long as the number of available processors is in the range $[1, \Theta(\sqrt{n})]$ —which arguably captures a wide range of practical settings of interest—this parallel DFS algorithm provides the best possible speedup over the sequential algorithm, up to logarithmic factors.

We next overview notions of work and depth in parallel algorithms and how they determine the time complexity given a number of processors. Then, we review prior work on parallel DFS computations. Afterward, we formally state our contributions.

1.1 Background: Work & depth

Work and depth in parallel algorithms. We follow the standard work-depth terminology [7].¹ For an algorithm \mathcal{A} , the *work* $W(\mathcal{A})$ is defined as the total number of operations. The *depth* $D(\mathcal{A})$ is defined as the longest chain of operations with sequential dependencies, in the sense that the $(i+1)^{th}$ operation depends on (and should wait for) the results of operations i in the chain. The work and depth bounds determine the time $T_p(\mathcal{A})$ for running the algorithm when we have p processor. We have $T_1(\mathcal{A}) = W(\mathcal{A})$ and $T_\infty(\mathcal{A}) = D(\mathcal{A})$. A simple observation, known as Brent’s principle [8], gives the following general bound:

$$W(\mathcal{A})/p \leq T_p(\mathcal{A}) \leq W(\mathcal{A})/p + D(\mathcal{A}).$$

Work-efficient parallel algorithms. Parallel algorithms with work $W(\mathcal{A})$ asymptotically equal to the complexity of their sequential counterpart are known as *work-efficient*. If this equality holds up to a poly($\log n$) factor, the algorithm is called *nearly work-efficient*. (Nearly) work-efficient parallel algorithms enjoy asymptotically optimal speed-up over sequential algorithms (up to logarithmic factors) for a small number of processors. Once the number of processors exceeds some threshold—asymptotically equal to $W(\mathcal{A})/D(\mathcal{A})$ —the time complexity bottoms out at $D(\mathcal{A})$. Thus, an ultimate goal in devising parallel algorithms is to obtain (nearly) work-efficient algorithms with depth as small as possible.

1.2 State of the art for parallel DFS algorithms

DFS is quite hard for parallel algorithms. Reif [14] showed that computing the lexicographically first DFS—where the DFS should visit the neighbors of a node according to their numbers—is P-complete.² That is, if there is a poly($\log n$)-depth poly(n) work algorithm for this problem, then there is such an algorithm for all problems in P. This is why his paper was titled “*Depth-first search is inherently sequential.*” Follow-up work on parallel DFS algorithms thus focused on computing an arbitrary DFS, where the order of visiting neighbors of a node gets chosen by the algorithm. This is also the version of the DFS problem that we consider. The most relevant prior work is by Aggarwal and Anderson [3] and a follow-up by Goldberg, Plotkin, and Vaidya [9]. These give parallel DFS algorithms, though focusing almost exclusively on depth and require work much higher than the sequential algorithm.

Aggarwal and Anderson [3], building partially on a prior work of Anderson[5], presented a randomized DFS algorithm for undirected



This work is licensed under a Creative Commons Attribution-NonDerivs International 4.0 License.

SPAA '23, June 17–19, 2023, Orlando, FL, USA
© 2023 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9545-8/23/06.
<https://doi.org/10.1145/3558481.3591094>

¹More concretely, we describe our work assuming the strongest PRAM variant, CRCW with arbitrary writes. This is done for simplicity. The results can be extended easily to the weaker variants, e.g., EREW, as the latter can simulate the stronger variants at the cost of a logarithmic factor loss in depth and work. We did not attempt to optimize the logarithmic factors in our results.

²Technically, the P-completeness is for a decision variant which asks whether a vertex v is visited before another vertex u or not.

graphs, with $\text{poly}(\log n)$ depth and $\text{poly}(n)$ work.³ The latter is a high and unspecified polynomial in n , which is at least $\Omega(n^3)$. This high work complexity is, in part, due to its use of exact maximum weight matching as a subroutine, for which known parallel algorithms need a high $\text{poly}(n)$ work [12]. Goldberg et al. [9] devised a deterministic variant of [3] with depth $\tilde{O}(\sqrt{n})$ and work $\text{poly}(n)$. A closer inspection indicates that this deterministic algorithm requires $\tilde{O}(m\sqrt{n})$ work.

To summarize, the state-of-the-art parallel DFS algorithms take work much higher than the sequential algorithm. Thus, they require a high number of processors to run faster than the sequential algorithm. Even in the work of Goldberg et al. [9], one would need $\Omega(\sqrt{n} \cdot \text{poly}(\log n))$ processors to see any time advantage over the sequential algorithm. This issue significantly limits the applicability of these algorithms [3, 9] in the current (or plausible future) settings of parallel computation.

1.3 Our contribution

We present the first nearly work-efficient parallel algorithm with sublinear depth. This algorithm will outperform the sequential counterpart as soon as we have $\text{poly}(\log n)$ processors. In fact, it exhibits an optimal speedup over the sequential algorithm, up to logarithmic factors, if the number of processors is at most $\Theta(\sqrt{n})$. Arguably, this range covers most of the practically relevant settings.

THEOREM 1.1. *There is a randomized parallel algorithm that, in any n -node m -edge undirected graph $G = (V, E)$, given a root node r , computes a depth-first search tree of G rooted at node r , using $\tilde{O}(m + n)$ work and $\tilde{O}(\sqrt{n})$ depth, with high probability.*

In the full version of this paper, we sketch how we can achieve the same statement as Theorem 1.1 using a deterministic algorithm. This is by replacing some randomized subroutines in algorithms that we use from prior work with deterministic counterparts, and it increases the depth and work bounds by only logarithmic factors.

Our algorithm for Theorem 1.1 follows the outer shell of the approach of Aggarwal and Anderson [3]. The novelty is in the internal ingredients. We adjust some parts of the algorithm to make it nearly work-efficient, and in particular, we use (and develop) certain batch-dynamic parallel data structures in several parts. The latter allows us to ensure that the total work remains $\tilde{O}(m + n)$. We provide an algorithm overview in Section 3.

2 PRELIMINARIES

2.1 Basic definitions

We use two basic definitions from prior work [3]:

Definition 2.1. (Initial DFS segment) *Consider a graph $G = (V, E)$ and a root node $r \in V$. An **initial DFS segment**, or simply an **initial segment**, is a tree $T' = (V', E')$ rooted in node r , where $V' \subseteq V$ and $E' \subseteq E$, such that T' can be extended to some depth-first search tree T'' rooted in node r . That is, there exists a depth-first search tree $T'' = (V, E'')$ such that $E'' \supseteq E'$.*

Observation 2.2. *A rooted tree T' is an initial segment iff there are no paths between different branches of T' using vertices in $V - T'$.*

³Aggarwal, Anderson, and Kao [4] later provided an extension to directed graphs, with similar bounds. In this paper, our focus is on undirected graphs.

More formally, there should be no path connecting two incomparable nodes of T' and made of internal nodes in $V - T'$. Here, incomparable means two nodes of T' , neither of which is an ancestor of the other.

Definition 2.3. (Separator) *For a graph H with n' vertices, a subset of vertices Q is called a **separator** if the largest connected component of $H - Q$ has size at most $\frac{n'}{2}$. We call a separator Q a **k -path separator** if Q is made of k vertex disjoint paths.*

2.2 Basic tools from prior work

Prefix sum on a linked list. The prefix sums in an n -item linked list can be computed in $O(\log n)$ depth and $O(n)$ work [6]:

Lemma 2.4. *Given a linked list (x_1, \dots, x_k) where each element x_i is associated with a number y_i , there is a deterministic algorithm that computes the prefix sum in $O(\log n)$ depth and $O(n)$ work. Moreover, the value $\sum_{j=1}^i y_j$ can be accessed directly from x_i .*

Maximal matching. Luby [13] provides an efficient deterministic maximal matching algorithm with $O(\log^5 n)$ depth and $\tilde{O}(m)$ work. We will use this in a black-box manner.

Lemma 2.5. *Given a graph $G = (V, E)$, there is a deterministic algorithm that computes a maximal matching in $O(\log^5 n)$ depth and $O(m \log^5 n)$ work.*

3 ALGORITHM OVERVIEW

The outer shell of our algorithm is based on the classic approach of Aggarwal and Anderson [3], which provides a $\text{poly}(\log n)$ -depth DFS algorithm but uses $\text{poly}(n)$ work for a high polynomial. The approach is recursive. Let G be the input graph, and r be the DFS root. We gradually grow an *initial DFS segment* of G until it becomes a complete DFS tree of G . See Section 2.1 for definitions. At the start, the segment is simply the root r . At each point, the initial segment is extended such that the problem is reduced to finding a new DFS tree in each connected component of the remaining graph, and such that each component has size at most half of the previous size. Thus, within $\log n$ recursions, the whole DFS is constructed.

The algorithm for extending an initial segment T' to a full DFS tree works as follows. For a connected component C in $G - T'$, by Observation 2.2, there is a unique vertex $x \in T'$ with the lowest depth that has a neighbor $y \in C$. We construct a DFS rooted in y for component C , then connect it to T' using the edge (x, y) . This is done in parallel for different connected components of $G - T'$.

The core of the algorithm is to construct an initial DFS segment that forms a separator for each component. This involves two parts: The first part is to find a separator that consists of a small number of vertex disjoint paths. The second part is constructing an initial segment from this set of paths. We next discuss each of these parts separately, in Section 3.1 and Section 3.2, and comment on how our algorithm differs from that of Aggarwal and Anderson and achieves $\tilde{O}(m)$ work.⁴

⁴From now on, we focus on connected graphs. This implies that $m = \Omega(n)$ and thus allows us to state the work bound simply as $\tilde{O}(m)$, instead of $\tilde{O}(m + n)$. Note that connected components can be identified in $\tilde{O}(m)$ work and $\text{poly}(\log n)$ depth via classical parallel algorithms [11].

3.1 Separator construction

The first part is to construct a separator with few paths.

Aggarwal-Anderson. Aggarwal and Anderson [3] present a $\text{poly}(n)$ -work and $\text{poly}(\log n)$ depth algorithm that computes an $O(1)$ -path separator, as follows: We start with the trivial n -path separator that consists of one path for each vertex. Then, in each iteration, we reduce the number of paths by a constant factor while ensuring that the paths form a separator. We continue this until at most a constant number of paths are left. The core of the process is reducing the number of paths by a constant factor. To do that, in a rough sense, the basic idea is to match up and merge pairs of paths iteratively. Aggarwal and Anderson [3] reduced this problem to a minimum weight perfect matching problem. The latter is known to be solvable using $\text{poly}(\log n)$ depth and a high $\text{poly}(n)$ amount of work [12], which is at least $\Omega(n^3)$ work. However, this approach does not yield a work-efficient parallel algorithm as we are unaware of any $\tilde{O}(m)$ -work algorithm for minimum-weight perfect matching, even with a sublinear depth.

Our separator algorithm. To keep the work bound $\tilde{O}(m)$, our separator will consist of $\tilde{O}(\sqrt{n})$ paths, instead of $O(1)$. In Section 4, we present a parallel algorithm that computes an $O(\sqrt{n})$ -path separator using near-linear work and $\tilde{O}(\sqrt{n})$ depth:

THEOREM 3.1 (SEPARATOR THEOREM). *There is an algorithm that finds an $O(\sqrt{n})$ -path separator Q in $O(\sqrt{n} \log^8 n)$ depth and $O(m \log^7 n)$ work. Each path is stored as one doubly-linked list.*

3.2 Construting an initial segment from the separator paths

The second part assumes that we have a separator consisting of several paths, and absorbs these paths into the current partial DFS tree, forming a new initial DFS segment that includes all vertices of these paths (and potentially some more).

Aggarwal-Anderson. Aggarwal and Anderson find a separator Q that consists of only $O(1)$ paths. Then, they add these paths to the partial DFS tree essentially one by one. They find a path p from the lowest node in the partial tree T' to a path l in the separator. Then, by adding the path p to the partial DFS tree T' , they can absorb at least half of the vertices of l into T' . These operations can be done using basic parallel spanning tree algorithms, in $\tilde{O}(m)$ depth and $\text{poly}(\log n)$ work. After repeating the above procedure at most $O(\log n)$ times, all vertices in the separator get absorbed into T' .

Our absorption algorithm. In contrast to the $O(1)$ -path separator of Aggarwal and Anderson, our work-efficient separator construction produces $O(\sqrt{n})$ paths. If we were to trivially absorb these paths, the work required with the basic solution would be $\tilde{O}(m\sqrt{n})$. To make this part work-efficient, we need that the total work over all absorptions is $\tilde{O}(m)$. We will perform each absorption using depth $\text{poly}(\log n)$ and work linearly proportional to the sum of the number of vertices on the paths and the number of edges adjacent to the vertices on the path. These vertices and edges get deleted from the remaining graph due to the absorption, and hence we can argue that the overall work is $\tilde{O}(m)$. The key algorithmic ingredient will be certain batch-dynamic parallel data structures, which can perform large batches of updates to the graph using depth

$\text{poly}(\log n)$ and work proportional to the total number of changes in the graph. In Section 5 and Section 6, we present the algorithms that provide this and prove the following theorem statement.

THEOREM 3.2 (ABSORPTION THEOREM). *Given an $O(\sqrt{n})$ -path separator, where each path is stored as a doubly-linked list, there is an algorithm that constructs an initial segment $T' \subseteq G$, where T' is also a separator for G . The algorithm uses $O(\sqrt{n} \log^3 n)$ depth w.h.p. and $O(m \log^3 n)$ work in expectation.*

4 SEPARATOR CONSTRUCTION

In this section, we prove the following theorem.

THEOREM 3.1 (SEPARATOR THEOREM). *There is an algorithm that finds an $O(\sqrt{n})$ -path separator Q in $O(\sqrt{n} \log^8 n)$ depth and $O(m \log^7 n)$ work. Each path is stored as one doubly-linked list.*

PROOF. We start with Q consisting of one path for each vertex. By Lemma 4.1, when Q consists of more than $48\sqrt{n}$ paths, we can repeatedly reduce the number of paths in Q by a constant factor in $O(\sqrt{n} \log^7 n)$ depth and $O(m \log^6 n)$ work. After $O(\log n)$ iterations, we get a separator Q consisting of at most $48\sqrt{n}$ paths. \square

4.1 Path Reduction

4.1.1 Review of Path Reductions in [3]. The crucial part of the separator construction is to reduce the number of paths by a constant fraction while preserving the separator property of Q .

Suppose that initially, Q has k paths. Aggarwal and Anderson [3] first divide the paths into a set L of long paths and a set S of short paths. Initially, $\frac{1}{4}k$ of the paths are placed in L and the rest $\frac{3}{4}k$ are placed in S . The idea is to find a set of vertex disjoint paths P between L and S . Each path p in P has one end on a long path and the other end on a short path. All internal vertices of p are not contained in Q , and each path in Q intersects at most one path in P . We refer to a set of paths P that satisfy all properties stated above as *valid*. Suppose the path p joins the path $l \in L$ and $s \in S$. Let $l = l'xl''$ and $s = s'ys''$ where x and y are the endpoints of p , and s' is equal or longer than s'' . Then l is replaced by $l'ps'$, and s by s'' , while we discard the path l'' .

Let \hat{L} and \hat{S} be the long and short paths that are joined, and L^* be the part of long paths that get discarded. Besides P being *valid*, Aggarwal and Anderson [3] want two more properties:

- (1) The set of paths P are maximal.
- (2) There is no path between L^* and $S - \hat{S}$.

Suppose that at least $\frac{1}{12}k$ of the paths are joined, and Q remains a separator after the replacement. Then the length of at least $\frac{1}{9}$ of the short paths are reduced by $\frac{1}{2}$. In at most $O(\log n)$ such steps, at least $\frac{1}{4}k$ of the short paths are removed. When this happens, one can stop because the number of paths in the separator has been reduced by a constant fraction.

If either of the above two conditions fails, then they [3] can directly find a different Q' such that it consists of less than a constant fraction of paths than Q .

4.1.2 Path Reduction Algorithm. The bottleneck for the work in [3] is the algorithm that finds a maximal P . They reduce this problem to the problem of minimum weight perfect matching, and solve

it in $\text{poly}(\log n)$ depth and using $\Omega(n^3)$ work via known perfect matching algorithms [12]. In our paper, we crucially want work-efficient algorithms, which use $\tilde{O}(m)$ work:

Lemma 4.1. *Given a separator Q that consists of k paths, with $k > 48\sqrt{n}$, there is an algorithm that finds Q' that consists of at most $\frac{47}{48}k$ paths in $O(\sqrt{n} \log^7 n)$ depth and $O(m \log^6 n)$ work.*

To prove this lemma, we find a set of vertex disjoint paths P that are valid (as defined above) with the exception that some of the paths in P have no endpoint in S . We divide P into P_1 and P_2 where P_1 are the “matched” paths that have the other end on S , while P_2 are the “unmatched” paths, meaning that their other ends are not on S .

Similarly, let \hat{L}_1 be the paths of L that are joined with S using P_1 . Suppose that the long path $l = l'xl'' \in \hat{L}_1$ and the short path $s = s'ys'' \in S$ are joined by $p \in P_1$. We would like to learn whether $|s'| \geq |s''|$ so that we can determine which part of the short path to join the long path. Since s is provided as a doubly-linked list, we make a copy of s and only keep one direction of the doubly-linked direction. Without loss of generality, let $s = s'ys''$ be the direction that is kept. Then we assign the value 1 to each element on the copied linked list, and we invoke Lemma 2.4 on the copied s to learn y 's rank on the list. If y 's rank is greater or equal to $\frac{1}{2}$ of the rank of the last vertex on the list, then $|s'| \geq |s''|$. Otherwise, $|s''| > |s'|$. This operation can be performed simultaneously for all paths with work proportional to the length of the paths and depth $O(\log n)$. Without loss of generality we assume $|s'| \geq |s''|$, we then replace l by $l'ps'$, s by s'' , and we add the discarded part l'' to L_1^* .

For the path $l = l'xl'' \in \hat{L}_2$ and the unmatched path $p \in P_2$. We replace l with $l'p$, discard l'' . We add the discarded part l'' to L_2^* . Denote the discarded parts as $L^* = L_1^* \cup L_2^*$, the long paths that are attached to P as $\hat{L} = \hat{L}_1 \cup \hat{L}_2$, and the short paths that are attached to P as \hat{S} . Let $D = V - P - L - S$ be the set of vertices not on any paths. A picture for illustration is shown in Fig. 1. The three properties we need are

- (1) The set of paths P are maximal in the sense that there is no path from $L - \hat{L}$ to $S - \hat{S}$ with all internal vertices being in D .
- (2) There is no path between the discarded part L^* and $S - \hat{S}$ with all internal vertices being in D .
- (3) The number of unmatched paths is small; concretely the number of paths in P_2 is less or equal to $\frac{1}{48}k$.

Suppose that these three conditions hold; we argue that either the new $Q = L \cup P \cup S - L^*$ remains a separator or we can directly find a different Q' consisting of less than a constant fraction of paths in Q . There are potentially two problems that can arise. In each case, we find a different separator with a constant reduction in the number of paths. We list the two potential problems below:

- (1) The discarded parts of L^* caused Q to no longer be a separator.
- (2) The number of matched paths is too small, meaning the number of paths in P_1 is less than $\frac{1}{12}k$.

Aggarwal and Anderson [3] show that in either case, we can immediately find another separator consisting of less than $\frac{47}{48}k$ paths. In our algorithm, we can follow the same solutions for these issues. For a review of their explanations, please see the full version of our paper. For the rest of this section, we can assume that neither

problem arises. Thus Q remains a separator and at least $\frac{1}{12}k$ short paths belong to \hat{S} . Recall that originally, there are $\frac{3}{4}k$ short paths. So at least $(\frac{k}{12})/(\frac{3k}{4}) = \frac{1}{9}$ fraction of the original short paths have their size cut in half. Thus, in $O(\log n)$ iterations, all short paths will be absorbed.

We will show in Lemma 4.2 that finding the desired set P can be done in $O(\sqrt{n} \log^6 n)$ depth and $O(m \log^5 n)$ work. For short paths of the form $s = s'ys''$, determining whether $|s'| \geq |s''|$ can be done in $O(\log n)$ depth and $O(n \log n)$ work across all short paths. Since we repeat the procedure at most $9 \log n$ times, we prove the work and depth bound of Theorem 3.1.

4.2 Path Merging

We first give a high-level description of our path-merging algorithm. In the beginning, each long path chooses one end as its head. Then each path simultaneously tries to extend itself from the head until it reaches a short path. If a path cannot extend, it kills the head vertex and backtracks one vertex. We continue the procedure until the number of “active” long paths is less than \sqrt{n} . This way, we ensure that roughly at least \sqrt{n} of the work is parallelized, which in turn bounds the total depth of the algorithm.

Lemma 4.2. *Assume $k \geq 48\sqrt{n}$. There is a parallel algorithm that finds the set of paths P satisfying the following three properties using depth $O(\sqrt{n} \log^6 n)$ and work $O(m \log^5 n)$.*

- (1) The set P of paths is maximal, i.e., there is no path from $L - \hat{L}$ to $S - \hat{S}$ such that the internal vertices of the path are in D .
- (2) There is no path between the discarded part L^* and $S - \hat{S}$ such that the internal vertices of the path are in D .
- (3) The number of unmatched paths is small; concretely the number of paths in P_2 is less or equal to $\frac{1}{48}k$.

We work on an auxiliary graph G' where we contract each short path $s \in S$ into a single vertex v_s . Throughout the process, each vertex in G' is in one of three states, namely *available*, *contained in a (long) path*, or *dead*. Initially, all vertices not contained in a long path are *available*. For each long path l , we pick an arbitrary end as the head of that path, denoted as u_l . Now, in each step, a head vertex u_l either gets matched to a neighbor v_l that is still *available* or it doesn't get matched. If u_l gets matched to v_l , then v_l joins the path l , and v_l becomes the new “head” of the path l . If u_l does not get matched, then only because of two possible reasons: Either, the node u_l corresponds to a contracted short vertex v_s . If that happens, we say that l succeeded. If l succeeded, then its head vertex does not get matched and l remains unchanged. Otherwise, if u_l does not correspond to a contracted short vertex v_s , then u_l has not been matched because all available neighboring vertices of u_l have been matched to different head vertices. In that case, u_l dies and is removed from the path l , and its predecessor w_l in the path l becomes the new head l . If u_l was the only node in l , then the path l does not participate in the matching process anymore. If the number of long paths attempting matching (or equivalently, the number of head vertices not corresponding to a contracted vertex v_s) is less than \sqrt{n} , then the process terminates, and we go to the postprocessing phase described below.

For a path l that successfully joins a short path, let l_{old} be its form before the above update, and $l'_{new} = \{v_1, \dots, v_k\}$ be its form




Figure 1: Merging long and short paths.

after the update. The last vertex $v_k \in G'$ on l_{new} corresponds to a contracted short path $s \in S$ in the original graph G . Thus, the edge $(v_{k-1}, v_k) \in G'$ corresponds to an edge $(v_{k-1}, v'_k) \in G$. We replace v_k with v'_k to get the corresponding path in l_{new} in the graph G . Then $p_1 = (l_{new} - l_{old}) \cup x$, where x is defined to be the only vertex in $l_{old} \cap l_{new}$ that was a head vertex during the path merging algorithm. So p_1 belongs to P_1 and path l_{old} belongs to \hat{L}_1 .

Similarly, for a path l that still participates in the matching process when the algorithm terminates, let l_{old} be its old form and l_{new} be its form after the update. Let $p_2 = (l_{new} - l_{old}) \cup x \in P_2$ where x is defined analogously. Path l_{old} belongs to \hat{L}_2 . For a path l that is not participating because all the vertices on the paths are dead, it belongs to $L - \hat{L}_1 - \hat{L}_2$.

We see that P_1 has one end on a long path and one end on a short path. And P_2 has one end on a long path and another on a vertex not in L or S . Now we show why the paths we find satisfy the above three properties. Property 3 holds because $k \geq 48\sqrt{n}$ and the number of paths in P_2 is less than \sqrt{n} . So we have that the number of paths in P_2 is less than $\frac{1}{48}k$. For properties 1 and 2, we observe that all vertices in $L - \hat{L}_1 - \hat{L}_2$ or L^* are dead. Notice that a vertex can only become dead if it was a head vertex of a long path but failed to get matched in the matching process. Hence, we get properties 1 and 2 from the following lemma.

Lemma 4.3. *Suppose that a vertex v becomes dead during the above algorithm, then there is no path from v to a vertex in $S - \hat{S}$ such that the internal vertices of the path are in D .*

PROOF. For a vertex v to die, it must be a head vertex of a long path during the matching process. It becomes dead if all its neighbors are dead or they joined some other long paths. For the sake of contradiction, suppose that there is a path $p = (v_1 = v, v_2, \dots, v_k = v_s)$ from v to a contracted vertex v_s formed from a short path with all internal vertices being in D . As v_1 is dead and v_k is available, there exists some index i with v_i being dead and v_{i+1} being available. Since v_{i+1} is available at the end of the matching process, it must have been in the available state throughout the whole matching process. This yields a contradiction as when v_i was attempting matching, it could have been matched to v_{i+1} which means v_i would not be dead. \square

PROOF OF LEMMA 4.2. We have shown that the output satisfies all the guarantees of Lemma 4.2. Lemma 4.4, proven in the next section, shows that there exists a parallel implementation of the procedure with work $O(m \log^5 n)$ and depth $O(\sqrt{n} \log^6 n)$, which finishes the proof of Lemma 4.2. \square

4.3 Parallel Implementation

It remains to discuss the parallel implementation of the procedure described in Section 4.2. In particular, the remaining part of this section is dedicated to proving the lemma below. We note that some parts of the matching procedure are nontrivial. The main reason the parallel implementation is nontrivial is that the number of steps for computing the paths can be up to $\Theta(\sqrt{n})$. However, we can only afford $O(m \log^5 n)$ work overall. Thus, the algorithm cannot afford to read the whole input in each iteration, and we should ensure that each edge is read only poly($\log n$) times, in an amortized sense.

Lemma 4.4. *There is a parallel algorithm that implements the procedure described in Section 4.2 with $O(m \log^5 n)$ work and $O(\sqrt{n} \log^6 n)$ depth.*

Recall that, at any point in time, a node is in one of three states and can change its state at most twice. Moreover, in each step, the number of vertices changing their state is equal to the number of head vertices attempting to get matched. As the process stops once less than \sqrt{n} head vertices attempt to get matched, this implies that at least \sqrt{n} vertices change their state in each step. Hence, the total number of steps is upper bounded by $2n/\sqrt{n} = O(\sqrt{n})$. Therefore, it suffices to show that each step can be implemented with $O(\log^6 n)$ depth. We say that an edge changes its state if one of its endpoints changes its state. In particular, each edge changes its state $O(1)$ times. Therefore, it suffices to show that each step can be implemented with work $O(N_{change} \log^5 n)$, where N_{change} is the total number of vertices and edges changing their state. Achieving this bound is nontrivial; just reading all available neighbors of a vertex u_i trying to get matched might already exceed it. Thus, our matching routine makes use of a data structure that allows to efficiently get access to a subset of u_i 's neighbors that are still available. We use the data structure from the lemma below.

Lemma 4.5. *There is a data structure with the following guarantees: The initial input is a graph $G = (\{v_1, v_2, \dots, v_n\}, \{e_1, e_2, \dots, e_m\})$*

with $n \geq 2$ where all vertices are active in the beginning. The data structure supports the following operations after initialization:

- **MakeInactive**($\{i_1, i_2, \dots, i_k\}$) takes an array consisting of $k \geq 1$ distinct indices between 1 and n and marks the corresponding vertices as inactive. This operation can be done in $O((k + \sum_{j=1}^k \deg_G(v_{i_j})) \log n)$ work and $O(\log n)$ depth.
- **Query**(i_1, i_2, \dots, i_k, t) takes an array consisting of k distinct indices between 1 and n and a number t . The output is an array a_j for every $j \in [k]$ containing t distinct active neighbors of vertex v_{i_j} . If v_{i_j} has fewer than t active neighbors, then a_j contains one entry for each active neighbor of v_{i_j} . The work is $O(k \cdot t \cdot \log n)$, and the depth is $O(\log n)$.

Initialization takes $O((m+n) \log n)$ work and $O(\log n)$ depth.

The data structure uses a common technique, and we defer the detailed description of it to the full version of this paper. On a high level, the adjacency list of each node v is augmented with a balanced binary tree. The leaves correspond to v 's neighbors, and each internal node keeps track of how many neighbors in the corresponding subtree are still active. In our concrete case, the input graph is the graph G' . We also maintain the invariant that a node is available if and only if it is marked as being active in the data structure. Note that we can initialize the data structure right at the beginning with $O(m \log n)$ work and $O(\log n)$ depth. Now, let's consider an arbitrary step of the algorithm. We denote by U the set consisting of all head vertices u_i trying to get matched. Our matching procedure builds the matching gradually in $O(\log n)$ phases. In each phase, it uses Luby's deterministic parallel maximal matching algorithm (Lemma 2.5) as a black box on a graph that is constructed with the help of the data structure. In more detail, in phase i , each node $u \in U$ that has not been matched in previous phases first selects 2^i arbitrary neighbors that are still available (in particular, which haven't been matched in previous phases). If u has fewer such neighbors, it selects all of them. By making use of the data structure, we can do the selection using $O(2^i \log n)$ work per node $u \in U$ that has not been matched and $O(\log n)$ depth. Let H_i be the bipartite graph where one side of the bipartition consists of all nodes in U that have not been previously matched, and the other side consists of all available nodes that have been selected by at least one node. Moreover, there is an edge between a head vertex u and an available vertex v if and only if u has selected v . We then compute a maximal matching of H_i in $O(|V(H_i)| + |E(H_i)| \log^5 n)$ work and $O(\log^5 n)$ depth using the algorithm of Lemma 2.5. In particular, one has $O(2^i \log^5 n)$ work per vertex in U that has not been matched before. Then, the data structure is updated by marking all previously available vertices that have been matched as inactive. Note that we can pay for this operation by charging $O(\log n)$ to each such vertex and $O(\log n)$ to each incident edge. We can do this charging as all previously available vertices that have been matched change their state (and thus also the incident edges). Also, after each phase, we remove vertices from U that don't have any available neighbors. Both the work and the depth of the algorithm are dominated by the $O(\log n)$ invocations of Luby's deterministic maximal matching algorithm. In particular, the overall depth is $O(\log n) \cdot O(\log^5 n) = O(\log^6 n)$. To upper bound the work, consider some vertex $u \in U$. First, consider that u has been matched in some phase i . Then, informally speaking, the algorithm has done $O(2^i \log^5 n)$ work

on behalf of u in all phases combined. If u has been matched in the first phase, then we charge $O(\log^5 n)$ work to the node that u matched with. If u has not been matched in the first phase, then at least 2^{i-1} neighbors of u got matched in phase $i-1$. Thus, at least 2^{i-1} edges incident to u change their state, and therefore u can charge each such edge $O(\log^5 n)$. We can use a similar charging argument if u has not been matched. Therefore, the work in each step is $O(N_{\text{change}} \log^5 n)$. This shows the work and depth bound of Lemma 4.4. It remains to argue correctness. In the last phase i , we have $2^i \geq n$. Therefore, each vertex $u \in U$ that has not been matched before selects all its neighbors that are still available. Thus, if the maximal matching in H_i does not match u , then all its available neighbors have been matched to other head vertices, which shows that the final matching indeed satisfies the guarantees stated in the previous section.

5 CONSTRUCTING AN INITIAL SEGMENT FROM THE SEPARATOR PATHS

In this section, we prove Theorem 3.2, which shows how we can add the $O(\sqrt{n})$ paths in the separator to the partial DFS one by one, using $\text{poly}(\log n)$ depth for each path, and $O(m)$ work overall. For the sake of readability, we first restate the lemma.

THEOREM 3.2 (ABSORPTION THEOREM). *Given an $O(\sqrt{n})$ -path separator, where each path is stored as a doubly-linked list, there is an algorithm that constructs an initial segment $T' \subseteq G$, where T' is also a separator for G . The algorithm uses $O(\sqrt{n} \log^3 n)$ depth w.h.p. and $O(m \log^3 n)$ work in expectation.*

We want to construct an initial segment T' that contains all the vertices in Q —the set of all vertices of the separator paths. We do this by absorbing paths of Q into T' one by one. To have a work-efficient algorithm, we want to ensure that in each iteration of absorbing a path, the total work is near-linear with respect to the number of edges adjacent to the path that got absorbed into T' .

Our algorithm uses a batch-dynamic parallel data structure. We next describe the interface of this data structure and use that to provide a proof for Theorem 3.2. The actual data structure that proves this lemma is presented in Section 6.

Lemma 5.1. *Given a graph $G = (V, E)$, a separator Q that consists of some paths, and a root $r \in G$ that forms the initial partial tree T' , there is a data structure that supports the following operations.*

- **FindCC()** returns a connected $C \subseteq G - T'$ such that $C \cap Q \neq \emptyset$, if no such C exists, the function returns *Success*.
- **LowestNode(C)** takes a connected component $C \subseteq G - T'$, and returns the vertex $v \in C$ that is adjacent to a vertex in $x \in T'$. The vertex x is the unique vertex with the lowest depth that is adjacent to C .
- **FindPathS2P(C, x)** takes a connected component $C \subseteq G - T'$ and a vertex $x \in T$ and returns a path p from x to a vertex $q \in Q$. All the vertices in p are not in Q except for q .
- **BatchDelete(p)** takes a path p consists of vertices (v_1, \dots, v_k) and deletes the vertices from $G - T'$.

Moreover, here are the work and depth of the above operations.

- **FindCC()** has work and depth $O(1)$.
- **LowestNode(C)** has work and depth $O(1)$.

- **FindPathS2P**(C, x) has depth $O(\log n)$ w.h.p. If the function returns a path p , the work is $O(|p| \log n)$ where $|p|$ is the number of vertices on the path p
- **BatchDelete**(p) has work $O(|E(p)| \log^3 n)$ in expectation and depth $O(\log^2 n)$ w.h.p. where $|E(p)|$ are the number of edges adjacent to the vertices in p .

Having this data structure, we can now prove Theorem 3.2.

PROOF OF THEOREM 3.2. To absorb the separator into the tree, we sequentially find a path p from a vertex $q \in Q \cap (G - T')$ to a vertex $x \in T'$, such that all the internal vertices of the path are in $G - T'$. Suppose that q belongs to the path $l = l'ql''$ in the separator, and $|l'| \geq |l''|$. Then we take the longer half l' and incorporate it into the tree T' , by adding the path pql' to the tree T' . For the new T' to remain an initial segment, we need x to be the lowest vertex in T' adjacent to the connected component containing q .

For each path $l = l'ql''$, we need to learn whether $|l'| \geq |l''|$, so that we can determine which part of the path to absorb in the initial segment. Since l is provided as a doubly-linked list, we make a copy of l and only keep one direction of the doubly-linked direction. Without loss of generality, let $l = l'yl''$ be the direction that is kept. Then we assign the value 1 to each element on the copied linked list, and we invoke Lemma 2.4 on the copied l to learn the rank of q in the list. If this rank is greater or equal to $\frac{1}{2}$ of the rank of the last vertex on the list, then $|l'| \geq |l''|$. Otherwise, $|l''| > |l'|$. This check can be performed using depth $O(\log n)$ and work proportional to the length of the path.

We also need to absorb the path pql' to T' , and vertices on pql' should learn their depth in T' . Suppose the absorption is through the edge (x, y) with $y \in T'$ where x is the first vertex on the path pql' . We next make each vertex in pql' learn its depth in T' , using $O(\log n)$ depth and work proportional to the length of pql' . For a prefix sum computation, we initiate vertex x with the value $depth_y$ —i.e., the depth of vertex y in the existing partial tree T' —and we initiate each remaining vertex on the pql' path with value 1. Then we invoke Lemma 2.4 on the path to compute the prefix sum on the list. As a result, all vertices on pql' learn their depth in T' . This operation can be done in $O(\log n)$ depth and $O(|pql'|)$ work where $|pql'|$ is the length of the path.

We repeatedly call FindCC() to find a connected component containing a vertex from $Q \cap (G - T')$, then find x using LowestNode(C) and call FindPathS2P(C, x) to find the desired path p . Then we determine whether $|l'| \geq |l''|$, and join the path pql' to T' . Finally, we call BatchDelete($\{pql'\}$) to delete the vertices from $G - T'$.

Lemma 5.1 shows that the first two operations have depth and work $O(1)$, the third operation has work $O(|p| \log n)$ and depth $O(\log n)$, and the last operation has $O(|E(p)| \log^3 n)$ amortized work in expectation and depth $O(\log^2 n)$, with high probability. Here $|E(p)|$ are the number of edges adjacent to the vertices in p . Determining whether $|l'| \geq |l''|$ and joining the path pql' to T' can be done in $O(\log n)$ depth for one joining and $O(n)$ work in total overall joining operations. Since we repeat the above sequence of operations $O(\sqrt{n} \log n)$ times, the algorithm uses $O(\sqrt{n} \log^3 n)$ depth w.h.p. and $O(m \log^3 n)$ work. \square

6 DATA STRUCTURES

The data structure we use is based on a combination of the ones developed by Acar et al. in [2] and [1], though we also need some further modifications.

We work on $G - T'$, and this graph undergoes vertex and edge deletions. During the construction of the initial segment in Section 5, we need to repeatedly find a path from the lowest vertex in the partial tree T' to a vertex in the separators, such that the path is made of internal vertices in $G - T'$. To do that efficiently, we maintain the connectivity structure of $G - T'$. We would like to keep one spanning tree for each connected component of $G - T'$. The are two main problems that we want to solve using this data structure. The first is that after joining a path to the partial DFS tree T' , thus deleting its vertices from $G - T'$, we need to update the connectivity structure of $G - T'$ and in particular, as some edges get deleted from the respective tree, we might have to find replacement edges for them. The second problem is to maintain the connectivity structure of $G - T'$ such that we can answer path queries of the following type efficiently: Given a set of vertices C and a vertex x , we need to report a path in $G - T'$ connecting x to a vertex in C , using $\text{poly}(\log n)$ depth and work proportional to the path length.

In the batch-dynamic setting, a batch of updates (or queries) are applied simultaneously; in our case, we will have batches of deletions to $G - T'$. For each such batch, we would like the work to be near-linearly proportional to the number of updates while keeping the depth $O(\text{poly} \log n)$. For this purpose, we use a modified and combined version of the data structures provided in [1] and [2]. We first provide a brief recap of these and then present the combined and adapted data structure that we need.

6.1 Connectivity and rake-and-compress data structures

Before recapping the data structures of [1, 2], let us remark on a small subtlety: their algorithms are written as supporting edge deletions. However, we can generally treat vertex deletions as deleting all edges adjacent to the vertex.

6.1.1 Parallelized connectivity data structure Algorithm. Consider a graph undergoing (batches of) edge deletions, and suppose we want to maintain a spanning tree for each connected component of it. Acar et al. [1] provide a solution for this, which is essentially a parallelized version of the sequential algorithm developed by Holm, de Lichtenberg, and Thorup (HDT) [10]. The HDT algorithm maintains a maximal spanning forest, certifying the graph's connectivity. We also note that although the work in the following lemma is stated in expectation, by increasing the work by a factor of $O(\log n)$, the work bound also holds w.h.p.

Lemma 6.1. *There is a parallel batch-dynamic connectivity data structure that maintains a maximal forest for a graph undergoing vertex and edge deletions. Given any batch of edge deletions, the data structure uses $O(\log^2 n)$ expected amortized work per edge deletion. The depth to process a batch of edge deletions is $O(\log^3 n)$ w.h.p.*

The main challenge is that, when an edge of the spanning forest is deleted, which essentially breaks the tree of the component into pieces, we need to see if there is a *replacement* edge that connects

the pieces. Furthermore, we might have to do several of these simultaneously, for a batch of edge deletions. To solve this efficiently, the HDT algorithm maintains a set of $\log n$ nested forests. The topmost level of the nest forest represents a spanning forest of the entire graph. Each level contains all tree edges stored in levels below it. A key invariant we keep is that the largest component size at the level i forest is at most 2^i . When an edge is searched and fails to become a replacement edge in the forest, we decrease the level of that edge by 1. This way, we ensure that an edge is searched at most $\log n$ times before it is deleted from our graph. Acar et al. [1] parallelize the task of finding replacement non-tree edges by examining multiple potential replacements at once, which gives the lemma we stated above.

6.1.2 Parallelized Rake and Compress Tree. The rake and compress operation for a static tree is a simple recursive procedure that allows one to “process” a tree in $O(\log n)$ simple iterations. The rake operation removes all leaves from the tree, except in the case of a pair of adjacent degree-one vertices where it removes the one with a smaller vertex id. The compress operation removes an independent set of vertices of degree two that are not adjacent to leaves. It is desired that this independent set has size within a constant fraction of the maximum independent set (in expectation). A simple analysis shows that after $O(\log n)$ iterations of rake and compress, the tree shrinks to a single vertex. This is because a constant fraction of the vertices in a forest are either leaves or degree 2 vertices due to the degree constraint.

Rake-and-compress as a low-depth hierarchical clustering. The rake and compress can be viewed as a recursive clustering process. A cluster is a connected subset of vertices and edges of the original forest. We note that a cluster may contain an edge without containing both of its endpoints. The boundary vertices of a cluster C are the vertices $v \in C$ that are adjacent to an edge $e \in C$. The vertices and edges of the original forest form the *base clusters*. Initially, each vertex and each edge form its own cluster. In the course of rake and compress, clusters are merged using the following rule: Whenever a vertex v is removed during rake or compress operations, all of the clusters with v as a boundary vertex are merged with the base cluster containing v ; we say v represents the new cluster. The children of a vertex are clusters that merged together to form it, which we store in an adjacency list. Thus we will have a collection of hierarchical forest $\{T_1, \dots, T_k\}$ where T_1 is the original forest, and T_{i+1} is generated from T_i after one iteration of the rake-and-compress. T_k is the final forest in which every connected component in the original forest is clustered into a single cluster. An example rake and compress tree are depicted in Fig. 2

We need a dynamic version of this, which will allow us to answer some type of queries in the tree, e.g., reporting a path in the tree between two nodes. Consider a forest that undergoes edge deletions and insertions (with the promise that all edges present after any of these updates form a forest). We would like to maintain the results of the rake and compresses operations during the $O(\log n)$ iterations, while the forest undergoes these updates.

Acar et al. [2] first observe that the decision of whether to remove a vertex v only depends on v 's neighbors $N(v)$ and the leaf status of $N(v)$. Then they argue that after one edge insertion/deletion,

only a constant number of vertices will have neighbors that change their leaf status. These vertices are called affected vertices. Meanwhile, for all other vertices, the decisions of whether to remove themselves are unaffected. In each subsequent iteration of the rake-and-compress process, the number of affected vertices due to that edge insertion/deletion grows by a constant additive factor. Thus the work induced by one edge insertion/deletion is only $O(\log n)$, yielding the following lemma. Like before, we note that although the work in the following lemma is stated in expectation, by increasing the work by a factor of $O(\log n)$, the work bound also holds w.h.p.

Lemma 6.2. *There is a dynamic data structure that maintains the hierarchical forests $\{T_1, \dots, T_k\}$ where T_1 is the original forest, and T_{i+1} is generated from T_i after one iteration of the rake-and-compress process. And T_k is the final forest in which every connected component in the original forest is clustered into a single cluster. For a n -node forest that undergoes edge insertions and deletions, processes batch insertions and deletions of k edges in $O(k \log n)$ work in expectation and $O(\log(n) \log^*(n))$ span w.h.p.*

6.2 Overview of the combined data structure

The combined and adapted data structure we next present proves Lemma 5.1. We first provide an overview of the data structure, and we then discuss each of the operations and their complexities.

To explain the data structure, let us briefly summarize what we will use it for. Our data structure is meant to be working on $G - T'$, the graph induced by vertices that are not in the current initial segment. We construct and maintain a parallelized HDT connectivity forest (as overviewed in Section 6.1.1). In addition, we keep a rake-and-compress representation (as overviewed in Section 6.1.2) for a copy of the HDT forest. The parallelized HDT forest is responsible for keeping a maximal forest of $G - T'$ as it undergoes vertex and edge detection. Whenever some edge deletions happen, the replacement edges found by this parallelized HDT are also fed into the rake-and-compress representations of the trees. The rake-and-compress representation allows one to efficiently answer queries about the component, e.g., finding a path between two vertices.

We augment the RC-tree with two different flags (from now on, we use the phrase RC-Tree to refer to the tree's rake-and-compress representation). The first flag is the separator flag which helps us to determine whether a connected component contains a vertex from the separator Q . In the RC-Tree, each vertex has a flag indicating whether it is in the separator Q . In the original tree T_1 , a vertex has a *separator flag* if it is in Q . When removing a vertex v from the RC tree, we merge all the clusters with v as the boundary vertex to form a new cluster. If any of the above clusters or v has a separator flag, the newly formed cluster will have the separator flag. Moreover, in the adjacency list, we sort the children cluster such that the children with separator flags appear in front of the children without separator flags. Finally, the clusters representing connected components are sorted into a linked list where the connected components with separator flags appear before the ones without the flags.

The second augmentation is the lowest neighbor augmentation which helps us find a connection vertex of a connected component C to the partial tree T' . Before the first stage of the rake-and-compress

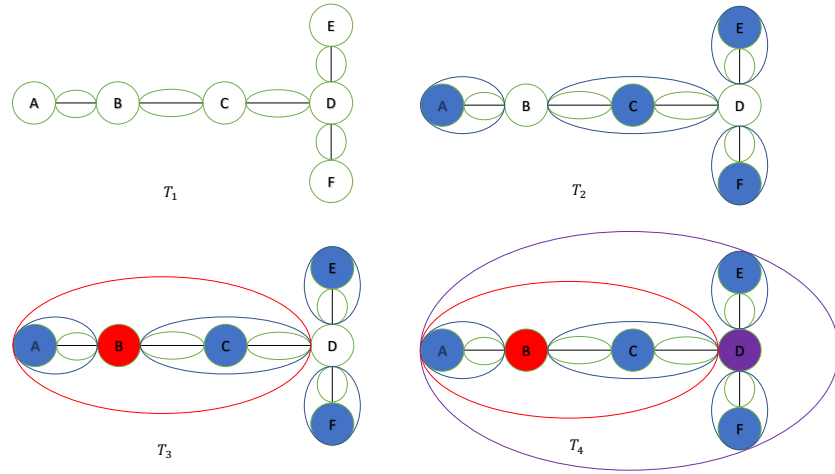


Figure 2: A recursive clustering of a tree. The tree consists of vertices $\{A, B, C, D, E, F\}$ and those vertices are connected by black edges. We use circles to represent different clusters. T_1 : The vertices and edges in the original tree form the green clusters. T_2 : The blue vertices are removed to form the blue clusters. Vertices $\{A, E, F\}$ are removed through the rake operation. Vertex $\{C\}$ is removed through the compress operation. T_3 : $\{B, D\}$ are both leaf vertices in the tree. Through tie-breaking, $\{B\}$ is removed to form the red cluster. T_4 : The only (leaf) vertex D in the tree is removed through the rake operation to form the purple cluster.

process, if a vertex v has neighbors in the partial tree T' , it is augmented with (v, d_v) where d_v is the depth of the lowest depth tree neighbor of v . If v has no tree neighbor, it is augmented with *None*. When merging clusters formed by removing v , suppose any of the clusters or v is augmented with anything other than *None*. Let $\{(v_1, d_{v_1}), \dots, (v_k, d_{v_k})\}$ be the augmentations of v 's children clusters, the cluster formed by removing v is augmented with (v_i, d_{v_i}) where d_{v_i} has (one of) the lowest depth among d_{v_1}, \dots, d_{v_k} . Otherwise, it is augmented with *None*. If v is augmented with (v_i, d_{v_i}) , then v_i has a neighbor $y \in T'$ such that y has the lowest depth among all tree neighbors of vertices in the cluster formed by v .

Next, we discuss the complexities of the operations provided by the data structure. We first discuss `FindCC()`, `LowestNode()`, and `BatchDelete()`, as they are simpler and shorter. Then, in a separate subsection, we discuss `FindPathS2P()`.

6.3 FindCC(), LowestNode(), and BatchDelete()

`FindCC()` has work and depth $O(1)$. This is because the connected components of $G - T'$ are sorted such that connected components with separator flags appear before those without. So we can just check the separator flag of the first connected component in the list. If it has the flag, then we return the connected component. Otherwise, the function returns *success*.

`LowestNode(C)` has work and depth $O(1)$ because it is augmented with (v, d_v) . As described above, v is the unique vertex with the lowest depth that is adjacent to a vertex contained in the cluster, and that cluster is the connected component C .

For `BatchDelete(p)`, vertices in $G - T'$ adjacent to p learn the depth of their new tree neighbors, so their lowest neighbor augmentations are updated. Suppose vertices on p have k total neighbors. Lemma 6.1 shows that deleting edges from the parallelized HDT and finding the replacement edges have amortized work $O(k \log^3 n)$. Since we can find a total of at most k replacement edges, by Lemma 6.2, inserting k edges to the RC-Tree can be done

in $O(k \log(1 + n/k))$ work in expectation and $O(\log(n) \log^*(n))$ depth w.h.p. This is because the vertices with augmentation updates are the vertices adjacent to the deleted edges.

6.4 FindPathS2P

We first describe how to report a path between two given vertices. Then we explain how to report a path between one set of vertices and another vertex, which provides exactly `FindPathS2P(C, x)`.

6.4.1 Point-to-point path queries. We first show that the Point-to-Point queries `FindPathP2P(x, y)` between two vertices x and y in the RC-Tree can be done using depth $O(\log n)$ and work proportional to $d(x, y)$, where $d(x, y)$ is the distance between x and y in the original tree.

Recall that the RC algorithm will produce a collection of hierarchical forests $\{T_1, \dots, T_k\}$ where T_1 is the original forest, and T_{i+1} is generated from T_i after one iteration of the rake-and-compress process. Also, T_k is the final forest in which every connected component in the original tree is clustered into a single cluster. Notice that in forest T_i , the edges are either edges in the original forest or clusters formed by the compression of the vertices. The vertices are either vertices in the original forest or clusters formed by the rake-and-compress of the vertices.

Lemma 6.3. *For two vertices x, y in the 1st level RC tree, there is an algorithm that finds the path between x and y in work $O(d(x, y) \log n)$ and depth $O(\log n)$ w.h.p.*

We first show that we can answer the path query between adjacent vertices in T_i .

Lemma 6.4. *Let x, y be two neighboring vertices in the i^{th} level of the RC-Tree T_i connected by the cluster edge E . There is an algorithm that finds the path in the original tree between x and y in work $O(d(x, y) \cdot i)$ and depth $O(i)$.*

PROOF. We prove the above statement by induction. Suppose that x, y are adjacent vertices in T_1 . Then the path between x, y is

just (x, y) , and the operation has work and depth $O(1)$. Now in T_i , if E is an edge from the original tree, then the path is still (x, y) . Suppose that E is a cluster formed by the compression of the vertex z in T_j for $j \leq i - 1$. Then z is still neighboring with both x and y in T_j . This is because once a vertex x becomes a boundary vertex of a cluster E in T_j , it will remain so in subsequent trees until E is merged with either x or another vertex. Moreover, the path between x, y in the original tree is the union of the paths formed by (x, z) and (z, y) . By induction, we can find the paths between (x, z) and (z, y) in parallel in depth $O(i - 1)$, and work $O(d(x, z) \cdot (i - 1))$ and $O(d(z, y) \cdot (i - 1))$ respectively. Since $d(x, y) = d(x, z) + d(z, y)$, the work of finding the path between x, y in T_i is

$$O(d(x, z) \cdot (i - 1)) + O(d(z, y) \cdot (i - 1)) + C = O(d(x, y) \cdot i)$$

where C is the constant overhead cost, thus proving Lemma 6.4. \square

Consider the general case where x, y are two vertices in the original tree. We recursively generate a list $(X_1, \dots, X_k), (Y_1, \dots, Y_k)$ where $X_1 = x, Y_1 = y$, and

$$X_{i+1} = \begin{cases} X_i & \text{if } X_i \text{ is not merged in the formation of } T_{i+1} \\ C & \text{if } X_i \text{ is merged to form the cluster } C \text{ in } T_{i+1} \end{cases}$$

We find the largest i such that $X_i \neq Y_i$. Since $X_{i+1} = Y_{i+1}$, we have that X_i and Y_i are children clusters of $X_{i+1} = Y_{i+1}$. Suppose X_{i+1} is formed from the removal of vertex x_{i+1} . Finding the path between x and y is equivalent to finding the paths between x and x_{i+1} , and between y and x_{i+1} . Moreover, X_i and Y_i both have x_{i+1} as one of their boundary vertices. In the following lemma, we show that this can be done in work $O(d(x, x_{i+1}) \cdot i)$ and $O(d(y, x_{i+1}) \cdot i)$ respectively, thus proving lemma Lemma 6.3.

Lemma 6.5. *Consider a vertex x in the RC-Tree with $(x = X_1, \dots, X_k)$, where*

$$X_{i+1} = \begin{cases} X_i & \text{if } X_i \text{ is not merged in the formation of } T_{i+1} \\ C & \text{if } X_i \text{ is merged to form the cluster } C \text{ in } T_{i+1} \end{cases}$$

We denote by x_i the vertex that is removed to form X_i . Suppose X_i has y as a one boundary vertex. Then there is an algorithm that finds a path between x and y in work $O(d(x, y) \cdot i)$ and depth $O(i)$.

PROOF. The proof is by induction. If $i = 1$, x_1 cannot have y as a boundary vertex. This is because only clusters formed by the edges can have boundary vertices. If $i = 2$, then the path between x and y is just the single edge path (x, y) .

Now we consider the case where $i > 2$. Suppose $X_i = X_{i-1}$. Then X_{i-1} still has y as its boundary vertex in T_{i-1} . This is because once a vertex v becomes a boundary vertex of a cluster C in T_j , it will remain so in subsequent trees until C is merged with either v or another vertex. Thus we have the result from induction.

Suppose $X_i \neq X_{i-1}$. We know that X_{i-1} is a child cluster of X_i formed by removing x_i in T_{i-1} . In T_{i-1} , if X_{i-1} has y as one of its boundary vertices, we have the result from induction. Suppose X_{i-1} has x_i as its boundary vertex and (y, x_i) are neighbors in T_{i-1} . In this case, the path between x and y is the union of path between (x_i, y) and (x, x_i) . The former path can be found in work $O(d(x_i, y) \cdot (i - 1))$ by Lemma 6.4, and the latter path can be found by induction in work $O(d(x, x_i) \cdot (i - 1))$, so the total work is still

$$O(d(x_i, y) \cdot (i - 1)) + O(d(x, x_i) \cdot (i - 1)) + C = d(x, y) \cdot i$$

where C is the constant overhead cost, thus proving Lemma 6.5. \square

6.4.2 Set-to-point path queries. Recall that in the original tree T_1 each vertex is augmented with a flag. In T_1 , a vertex has the separator flag if it is in Q . Moreover, a cluster has the separator if any of its child clusters has the separator flag.

If x has the separator flag, we just return x . Otherwise, we recursively generate a list $(x = X_1, \dots, X_k)$ where

$$X_{i+1} = \begin{cases} X_i & \text{if } X_i \text{ is not merged in the formation of } T_{i+1} \\ C & \text{if } X_i \text{ is merged to form the cluster } C \text{ in } T_{i+1} \end{cases}$$

We find the largest i such that X_i does not have the separator flag. There are two cases to consider.

- (1) Suppose that $x_{i+1} \notin Q$. Since X_{i+1} has a separator flag, it must have a child cluster Y that has the separator flag, and the cluster X_i does not have the separator flag. So we know that there is a path from x to some vertex $q \in Y \cap Q$ using x_{i+1} , so we call $\text{FindPathP2P}(x, x_{i+1})$ and $\text{FindPath}'_i(Y, x_{i+1})$, where we will describe below, to find the paths. The function then returns the union of the above two paths.
- (2) Suppose that $x_{i+1} \in Q$. Then we can call $\text{FindPathP2P}(x, x_{i+1})$ directly to find the path.

For the function $\text{FindPath}'_i(Y, z)$, the subscript i indicates that the function is called on T_i . We assume Y has z as its boundary vertex and Y has the separator flag. Suppose Y is formed from the removal of vertex y in T_j . Let E be the (cluster) edge that connects y and z in T_j , if E has the separator flag, we go to the lower level tree and call $\text{FindPath}'_{i-1}(E, z)$. Otherwise, we know E does not have any vertex that belongs to Q . Let E' be another child of Y that has the separator flag, then we can call $\text{FindPathP2P}(z, y)$ and $\text{FindPath}'_{i-1}(y, E')$ because the union of the above two paths is the desired path.

Thus the $\text{FindPath}'$ function follows the following recursive (informal) relationship.

$$\text{FindPath}'_i = \text{FindPath}'_{i-1} + \text{FindPathP2P}$$

The depth of $\text{FindPathS2P}(C, x)$ query is $O(\log n)$ w.h.p. This is because the main function FindPathS2P only calls the function $\text{FindPath}'_i$ at most once for some i . For the function $\text{FindPath}'_i$, as shown above, the depth of the recursion is at most $i = O(\log n)$ w.h.p. In the meanwhile, we can run the FindPathP2P queries in parallel, and those operations have depth $O(\log n)$ w.h.p.

For the function $\text{FindPathS2P}(C, x)$, suppose that the returned path is a path between x and a vertex $y \in C$. Then the work is $O(d \log n)$ w.h.p. where d is the distance between (x, y) in the original tree. This is because the work cost of FindPathS2P is the sum of all the work from FindPathP2P and $O(\log n)$ overhead cost. The work of FindPathP2P is proportional to its length, and the final path returned by the function is the union of the paths returned by all calls of FindPathP2P .

7 ACKNOWLEDGMENT

C. G. was supported by the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation program (grant agreement No. 853109).

REFERENCES

- [1] Umut A Acar, Daniel Anderson, Guy E Blelloch, and Laxman Dhulipala. 2019. Parallel batch-dynamic graph connectivity. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. 381–392.
- [2] Umut A Acar, Daniel Anderson, Guy E Blelloch, Laxman Dhulipala, and Sam Westrick. 2020. Parallel Batch-Dynamic Trees via Change Propagation. In *Annual European Symposium on Algorithms (ESA)*.
- [3] A. Aggarwal and R. Anderson. 1987. A Random NC Algorithm for Depth First Search. In *ACM Symposium on Theory of Computing (STOC)*. 325–334.
- [4] Alok Aggarwal, Richard J Anderson, and M-Y Kao. 1989. Parallel depth-first search in general directed graphs. In *ACM symposium on Theory of Computing (STOC)*. 297–308.
- [5] Richard Anderson. 1985. A parallel algorithm for the maximal path problem. In *ACM Symposium on Theory of Computing (STOC)*. 33–37.
- [6] Richard J. Anderson and Gary L. Miller. 1990. A simple randomized parallel algorithm for list-ranking. *Inform. Process. Lett.* 33, 5 (1990), 269–273. [https://doi.org/10.1016/0020-0190\(90\)90196-5](https://doi.org/10.1016/0020-0190(90)90196-5)
- [7] Guy E Blelloch. 1996. Programming parallel algorithms. *Commun. ACM* 39, 4 (1996), 85–97.
- [8] Richard P Brent. 1974. The parallel evaluation of general arithmetic expressions. *Journal of the ACM (JACM)* 21, 2 (1974), 201–206.
- [9] AV Goldberg, SA Plotkin, and PM Vaidya. 1988. Sublinear-time parallel algorithms for matching and related problems. In *IEEE Symposium on Foundations of Computer Science (FOCS)*. 174–185.
- [10] Jacob Holm, Kristian De Lichtenberg, and Mikkel Thorup. 2001. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *Journal of the ACM (JACM)* 48, 4 (2001), 723–760.
- [11] Joseph Jaja. 1992. *An introduction to parallel algorithms*. Vol. 17. Addison-Wesley Reading.
- [12] Richard M Karp, Eli Upfal, and Avi Wigderson. 1985. Constructing a perfect matching is in random NC. In *ACM symposium on Theory of Computing (STOC)*. 22–32.
- [13] Michael Luby. 1993. Removing randomness in parallel computation without a processor penalty. *J. Comput. System Sci.* 47, 2 (1993), 250–286. [https://doi.org/10.1016/0022-0000\(93\)90033-5](https://doi.org/10.1016/0022-0000(93)90033-5)
- [14] John H Reif. 1985. Depth-first search is inherently sequential. *Inform. Process. Lett.* 20, 5 (1985), 229–234.