

Adaptive Massively Parallel Algorithms for Cut Problems

ABSTRACT

We study the *Weighted Min Cut* problem in the *Adaptive Massively Parallel Computation* (AMPC) model. In 2019, Behnezhad et al. [3] introduced the AMPC model as an extension of the *Massively Parallel Computation* (MPC) model. In the past decade, research on highly scalable algorithms has had significant impact on many massive systems. The MPC model, introduced in 2010 by Karloff et al. [16], which is an abstraction of famous practical frameworks such as MapReduce, Hadoop, Flume, and Spark, has been at the forefront of this research. While great strides have been taken to create highly efficient MPC algorithms for a range of problems, recent progress has been limited by the *1-vs-2 Cycle Conjecture* [20], which postulates that the simple problem of distinguishing between one and two cycles requires $\Omega(\log n)$ MPC rounds. In the AMPC model, each machine has adaptive read access to a distributed hash table even when communication is restricted (i.e., in the middle of a *round*). While remaining practical [4], this gives algorithms the power to bypass limitations like the 1-vs-2 Cycle Conjecture.

We give the first sublogarithmic AMPC algorithm, requiring $O(\log \log n)$ rounds, for $(2 + \epsilon)$ -approximate weighted Min Cut. Our algorithm is inspired by the divide and conquer approach of Ghaffari and Nowicki [11], which solves the $(2 + \epsilon)$ -approximate weighted Min Cut problem in $O(\log n \log \log n)$ rounds of MPC using the classic result of Karger and Stein [15]. Our work is *fully-scalable* in the sense that the local memory of each machine is $O(n^\epsilon)$ for any constant $0 < \epsilon < 1$. There are no $o(\log n)$ -round MPC algorithms for Min Cut in this memory regime assuming the 1-vs-2 Cycle Conjecture holds. The *exponential speedup* in the AMPC runtime is the result of *decoupling* the different layers of the divide and conquer algorithm and solving all layers in $O(1)$ rounds in parallel. Finally, we extend our approach: we present an $O(k \log \log n)$ -round AMPC algorithms for approximating the Min k -Cut problem with a $4 + \epsilon$ approximation factor.

CCS CONCEPTS

• **Theory of computation** → **Massively parallel algorithms;**
Graph algorithms analysis.

KEYWORDS

Massively Parallel Computation, minimum cut, adaptivity

ACM Reference Format:

. 2018. Adaptive Massively Parallel Algorithms for Cut Problems. In *Woodstock '18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

Massively Parallel Computation (MPC) – introduced by Karloff et al. [16] in 2010 – is an abstract model that captures the capabilities of the modern parallel/distributed frameworks widely used in practice such as MapReduce [7], Hadoop [10], Flume [6], and Spark [21]. MPC has been at the forefront of the research on parallel algorithms in the past decade, and it is now known as the de facto standard computation model for the analysis of parallel algorithms.

In this paper, we focus on sublogarithmic-round algorithms for the Min Cut problem in the *Adaptive Massively Parallel Computation* (AMPC) model, which is a recent extension of MPC. In both MPC and AMPC, the input data is far larger than the memory of a single machine, and thus an input of size N is initially distributed across a collection of P machines. In the MPC model, the algorithm executes in several synchronous rounds, in which each machine executes local computations isolated from other machines, and the machines can only communicate at the end of a round. The total size of incoming/outgoing messages for each machine is limited by local memory constraints. We are interested in *fully-scalable* algorithms in which every machine is allocated a local memory of size $O(N^\epsilon)$ for any constant $0 < \epsilon < 1$. Moreover, we can often improve the round complexity¹ of the massively parallel algorithms by allowing a super-linear total memory $O(N^{1+\epsilon})$, for example, the filtering technique of Lattenzi et al. [17] in MPC or the maximal matching algorithm of Behnezhad et al. [4] in AMPC. So we are primarily interested in algorithms with $\tilde{O}(N)$ total memory, and therefore we assume there are $P = \tilde{O}(N^{1-\epsilon})$ machines.²

Recent developments in the hardware infrastructure and new technologies such as RDMA [8], eRPC, and Farm [9] allow for high-throughput, low-latency communication among machines in data centers, such that remote volatile memory accesses are becoming faster than accessing local persistent storage. The concept of a shared remote memory is in particular useful when machines need to query data adaptively – i.e., deciding what to query next based on the previously queried data – which requires a communication round per query in the MPC model. Behnezhad et al. [3] incorporates this RDMA-like paradigm of remote memory access into the MPC model and introduces AMPC. In the new model, the machines can *adaptively* query from a *distributed hash table*, or a shared read-only memory, during each round. Machines are only allowed to write to shared memory at the end of each round. There is also empirical evidence that AMPC algorithms for several

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Woodstock '18, June 03–05, 2018, Woodstock, NY

© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/10.1145/1122445.1122456>

¹The number of rounds is a main complexity of interest since in practice the bottleneck is often the communication phase.

²Where \tilde{O} hides polylogarithmic factors, i.e., $\tilde{O}(f(n)) = O(f(n) \text{poly} \log(n))$.

problems – including maximal independent set, maximal matching, and connectivity – obtain significant speedups in running time compared to state-of-the-art MPC algorithms [4]. This fact, which stems from the meaningful drop in the number of communication rounds, verifies the practical power of the AMPC model.

In this paper, we provide the first AMPC-specific algorithms for the Min Cut problem. The Min Cut of a given graph $G = (V, E)$ is the minimum number of outgoing edges, $\delta(S)$, among every subset of vertices $S \subseteq V$. The celebrated result of Karger and Stein [15] solves Min Cut by recursively contracting edges in random order. Specifically, it runs two instances of the contraction process with different seeds in parallel. Each instance is run in parallel until the graph size is reduced by a factor of $\frac{1}{\sqrt{2}}$, at which point each instance recurses (thereby creating a parallel split again). They return the minimum of the two returned cuts. The algorithm itself is mainly inspired by another result of Karger [14] for finding the Min Cut using graph contractions. We also extend our approach to the Min k -Cut problem, in which we are given a graph $G = (V, E)$ and an integer k and we want to find a decomposition of V into k subsets V_1, V_2, \dots, V_k so that $\sum_{i=1}^k \delta(V_i)$ is minimized. We utilize the greedy algorithm of Saran and Vazirani [18] which gives an $O(2 - \frac{2}{k})$ -approximation of the Min k -Cut. Gomuri and Hu give an alternative algorithm with the same approximation guarantee with additional features [12].

We study the Min Cut and Min k -Cut problems in the AMPC model. We give $O(\log \log n)$ -round AMPC algorithms for a $(2 + \epsilon)$ -approximation of Min Cut and a $(4 + \epsilon)$ -approximation of Min k -Cut.

1.1 Adaptive Massively Parallel Computation (AMPC)

Massively Parallel Computation (MPC) and Adaptive Massively Parallel Computation (AMPC) both sprung out of an interest in formalizing a theoretical model for the famous MapReduce programming framework. The most common problems in MPC and AMPC are on graph inputs, and since our paper only considers graph problems, we define these two models in terms of problems on graphs. Consider a graph $G = (V, E)$ with $n = |V|$ and $m = |E|$.

In standard MPC [1, 13, 16, 17], we are given a collection of P machines and are allowed to compute the solution to a problem in parallel. As we have already discussed, MPC computation occurs in synchronous rounds, each consisting of local polynomial-time computation and ending with machine-machine communication where all messages sent to and from a machine must fit within its local memory. Fully-scalable algorithms, the strongest memory regime in MPC, require the local memory to be constrained by $O(n^\epsilon)$ for any given $0 < \epsilon < 1$. Additionally, we are primarily interested in algorithms that require at most $O(\log n)$ rounds. However, often sublogarithmic – i.e., $O(\sqrt{\log n})$ or $O(\log \log n)$ – round complexity is much more desirable. In most cases, the total space must be at most $\tilde{O}(n + m)$, though sometimes we allow slightly superlinear total space.

AMPC extends MPC to add functionality while remaining implementable on modern hardware. Formally, in the AMPC model, we are given a set of distributed hash tables $\mathcal{H}_0, \dots, \mathcal{H}_k$ for each

of the k rounds of computation. These hash tables are each limited in size by the total space of the model (i.e., $\tilde{O}(n + m)$). As in MPC, we are given a number of machines and computation proceeds in rounds. In each round, local computations occur and then messages are sent between machines. The distinction in AMPC is that during the local computations, machines are allowed simultaneous read access to the hash table for that round (i.e., \mathcal{H}_{i-1} for round i) and during the messaging phase of the round, they are allowed to write data to the next hash table, \mathcal{H}_i . Reading and writing is limited by machine local memory. The power of the AMPC model over the MPC model is that, at the beginning of a round, the machines do not need to choose all the data they will access during the round. Instead, they can dynamically access the data stored in the hash table over the course of the local computation, thus potentially selecting data based on its own local computation.

It is not too hard to see that AMPC is a strictly stronger model than MPC. In fact, it was formally shown that all MPC algorithms can be implemented in AMPC with the same round and space complexities [3].

1.2 Our Contributions and Methods

This work is the first to study the Adaptive Massively Parallel Computation (AMPC) model for Min Cut problems on graphs. We mainly focus on the standard single Min Cut problem, although we also propose an approximation algorithm for the Min k -Cut problem. Our main result for the Min Cut problem is a $2 + \epsilon$ approximate algorithm that uses sublogarithmic $O(\log \log n)$ rounds.

THEOREM 1. *There is an $O(\log \log n)$ -round AMPC algorithm that uses $\tilde{O}(n + m)$ total memory and $\tilde{O}(n^\epsilon)$ memory per machine which finds a $(2 + \epsilon)$ -approximation of Min Cut with high probability.*

Note that this is a vast improvement over the current state-of-the-art algorithms in MPC by Ghaffari and Nowicki [11], which achieves the same $2 + \epsilon$ approximation in $O(\log n \log \log n)$ rounds. Both our algorithm and that of Ghaffari and Nowicki use Karger’s methods as a general structure for finding the Min Cut. Using this method, the goal is to recursively execute random graph contractions. From the results of Karger, the contraction process either finds a singleton cut that is a $2 + \epsilon$ approximation or preserves a specific Min Cut with probability dependent on the depth of recursion. To leverage this result, at each step of the recursion process, we find the best singleton cut on the existing graph. Once the graph is small enough, the problem can be solved efficiently. Out of all the singleton solutions found during this process and the final Min Cut on the small graph, we simply select the best cut. This is a $2 + \epsilon$ approximate Min Cut with high probability.

To implement this approach in a distributed model, both methods assign random weights to the edges of the input graph and find a minimum spanning tree (MST). Greedily, selecting edges in order of decreasing weight, we contract the graph along the current edge. This process is equivalent to the same greedy random contraction process on the original graph. This step, already, currently requires at least $\Omega(\log n)$ rounds in MPC, but the flexibility of the AMPC model allows us to achieve this step in a constant number of rounds.

It remains to show how can one find the best singleton cuts at each level of recursion. In order to do this, we employ a *low-depth*

tree decomposition on the minimum spanning tree until it becomes a set of separated vertices. On top of this recursive divide-and-conquer process, we design a process to compute and remember the best singleton cut.

The high level idea of recursively partitioning the tree and applying a process on top of that to find the best singleton cut is the same in both our paper and Ghaffari and Nowicki's paper [11]. However, the processes used to do this in MPC do not yield simple improvements in AMPC. Rather, we must use entirely novel techniques that leverage adaptivity to get truly sublogarithmic results. In fact, this must be done in constant rounds to achieve our results, whereas Ghaffari and Nowicki do this in $O(\log n)$ rounds. In order to create a tree decomposition, we consider maximal paths of heavy edges (i.e., edges that go from a parent to its child with the largest subtree). These paths are replaced by binary trees whose leaves are the path and the root connects to the path's parent. Consider labeling the resulting vertices in the graph with their depth. For each internal node in one of these binary trees, which was *not* a vertex in the original tree, we select a specific descendant leaf in the binary tree expansion of the path to send its depth to. The final value a vertex receives is then what we call the "label", which measures at what level of recursion the tree splits at that vertex. An entire labeling of the tree encodes an entire tree decomposition. This is done in constant AMPC rounds.

To compute the singleton cuts at each level, we assign to each singleton cut formed during the contraction process a vertex that has the lowest label. We show that such vertices are well-defined, i.e. there is only one vertex with the lowest label within vertices on the same side of a singleton cut. Because removing vertices of labels lower than i partitions the tree into disjoint subtrees such that each subtree contains at most one vertex with label i , we are able to calculate minimal singleton cuts corresponding to these vertices with label i in parallel in a constant number of AMPC rounds. Since, we constructed the low-depth decomposition such that the range of labels has size $O(\log^2 n)$, thus, by increasing the total memory, we can perform these computations for all different labels in a constant number of AMPC rounds. For more details, we defer to Section 4.

We then show how this work can be leveraged to achieve efficient results for approximate Min k -Cut, generalizing the results from Saran and Vazirani [18]. At a high level, we start by computing a Min Cut. Then we add the edges of the cut to a set D . In all following $k - 1$ iterations, we calculate the Min Cut on the graph without edges in D , and add the new cut edges to D for the next iteration. The set of the first k cuts we compute is our k -cut.

Compared to Saran's and Vazirani's technique, our method uses an *approximate* Min Cut rather than an exact Min Cut on each splitting step. This requires adapted analysis of this general approach. We employ the structure of Gomory-Hu trees (see [12]) for this purpose and show the following result:

THEOREM 2. *Algorithm APX-SPLIT is an $(4 + \epsilon)$ approximation of the Min k -Cut. Furthermore, it can be implemented in the AMPC model with $O(n^\epsilon)$ memory per machine in $O(k \log \log n)$ rounds and $O(m)$ total memory.*

Therefore, for small values of k , we can achieve efficient algorithms for $4 + \epsilon$ approximate Min k -Cut in AMPC. Note that there

are no existing results in the MPC model, however our methods applied to the work of Ghaffari and Nowicki [11] yield:

COROLLARY 1. *There is an algorithm that achieves a $(4 + \epsilon)$ approximation of the Min k -Cut with high probability that can be implemented in the MPC model with $O(n^\epsilon)$ memory per machine in $O(k \log n \log \log n)$ rounds and $O(m)$ total memory.*

Note there is still a logarithmic-in- n improvement in the round complexity in AMPC over MPC no matter the value of k . Due to space constraints both these result are presented in Section A in the appendix.

2 MINIMUM CUT IN AMPC

Karger and Stein [15] proposed a foundational edge contraction strategy for solving Min Cut:

- Create two copies of G , and independently on each, contract edges in a random order until there are at most $\frac{n}{\sqrt{2}}$ vertices.
- Recursively solve the problem on each contracted copy until they have constant size.
- Return the minimum of the cuts found on both copies.

LEMMA 1 ([15]). *The contraction process executed to the point where there are only $\frac{n}{t}$ vertices left preserves any fixed minimum cut with probability $\Omega\left(\frac{1}{t^2}\right)$.*

According to Lemma 1, naively contracting random edges until there are only two vertices remaining preserves at least one minimum cut with probability $\Omega\left(\frac{1}{n^2}\right)$. Thus, we need to repeat the naive contraction process at least $O(n^2 \log n)$ times so that we have a high *probability of success*, i.e., preserving a minimum cut. However, Karger and Stein [15] show that their recursive strategy succeeds with probability $\Omega\left(\frac{1}{\log n}\right)$. In turn, running $O(\log^2 n)$ instances of the recursive strategy is enough to find a minimum cut with high probability.

Roughly speaking, the choice of $t = \sqrt{2}$ as the inverse of the branching factor assures that a minimum cut is preserved with probability $\frac{1}{\sqrt{2}} = \frac{1}{2}$ throughout the contractions in each copy. Thus, the probability of success, say $P(n)$, for n vertices is bounded by:

$$P(n) \geq 1 - \left(1 - \frac{1}{2} \cdot P\left(\frac{n}{\sqrt{2}}\right)\right)^2 \quad (1)$$

Note that the random contractions in two copies are assumed to be independent, and the probability of success for each copy is at least $\frac{1}{2} \cdot P\left(\frac{n}{\sqrt{2}}\right)$ since we recurse on the resulting contracted graph with $\frac{n}{\sqrt{2}}$ vertices. Inequality (1) implies that at the k -th level of recursion (counting from the bottom), the probability of success is $\Omega\left(\frac{1}{k}\right)$, and in particular $\Omega\left(\frac{1}{\log n}\right)$ at the root of recursion. [15].

Let us now give some high-level insight into the approach by Ghaffari's and Nowicki. Ghaffari and Nowicki [11] observed that if we only desire a $(2 + \epsilon)$ approximate cut, we can use a better bound for the probability of preserving a minimum cut, or alternatively, the success probability.

LEMMA 2 ([11, 15]). *On an n -vertex graph G , let C be a minimum cut with weight λ . Fix an arbitrary $\epsilon \in (0, 1)$. The described random*

contraction process that contracts G down to $\frac{n}{t}$ vertices either at some step creates a singleton cut of size at most $(2+\epsilon)\lambda$ or preserves C - i.e., it does not contract any of its edges - with probability at least $\frac{1}{t^{1-\epsilon/3}}$.

A *singleton cut* is a partitioning of graph vertices so that there is only one vertex on one side, i.e., $\delta(S)$ so that $|S| = 1$. Assuming that one is able to verify whether a singleton cut of a small size has been formed during the contraction process, they show that this greater probability of success can boost the recursive process. In short, consider the k -th level of recursion, where level 0 corresponds to the bottom level. Let $\frac{n}{t_k}$ be the size of a single recursive instance at level k , and denote by s_k the total number of instances on this level. For all k , they ensure $s_k = t_k^{1-\epsilon/3}$.

Now, let $x_k^{1-\epsilon/3}$ be the branching factor on level k . That is, the recursion produces $x_k^{1-\epsilon/3}$ copies of the instance at level k , and on each of them independently contracts edges in a random order until the number of vertices is bigger than $\frac{n}{t_k} \cdot \frac{1}{x_k}$. If we have an algorithm that is able to *track* whether a small singleton cut appeared in each of these random processes, we either get a singleton cut that $(2+\epsilon)$ approximates a minimum cut or a minimum cut is preserved with probability $x_k^{1-\epsilon/3}$. Since we made $x_k^{1-\epsilon/3}$ copies, by a similar argument as in Karger's approach, we get that, in the latter case, the probability of preserving a minimum cut is $\Omega\left(\frac{1}{k}\right)$.

Finally, observe that on the k -th level of recursion, the most costly operation is copying a k -th level instance $x_k^{1-\epsilon/3}$ times in order to contract edges in each of these instances. Since the instance has size $\frac{n}{t_k}$ and we have s_k instances, processing these tasks in parallel requires $\frac{n}{t_k} \cdot s_k \cdot x_k^{1-\epsilon/3}$ space. If one want to fit this in $O(n)$ space, then it must be that $x_k \leq t_k^{(\epsilon/3)/(1-\epsilon/3)}$. Anyway, we get that the number of contractions we can make on k -th level is polynomial in the number of contractions we made on higher levels, and if the recurrence is solved, then it follows that it will be $O(\log \log n)$ levels until we reach a graph of a constant size.

Ghaffari and Nowicki [11], use Lemma 2 and the above boosting scheme to show an $O(\log \log n \cdot \log n)$ -round MPC algorithm for Min Cut. The main non-trivial part of their algorithm involves tracking the smallest singleton cut on each recursion level, which they do in $O(\log n)$ rounds because of the divide and conquer nature of their approach. Effectively, they assign all edges random and unique edge weights, and contract all uncontracted edges in decreasing order by edge weight. It can then be shown that all that needs to be done is to compute the MST of this graph and contract these edges accordingly (all other edges will be automatically contracted when another edge is contracted). We reduce the number of rounds for singleton cut tracking down to $O(1)$ rounds in the AMPC model. We aim to prove the following theorem.

THEOREM 1. *There is an $O(\log \log n)$ -round AMPC algorithm that uses $\tilde{O}(n+m)$ total memory and $\tilde{O}(n^\epsilon)$ memory per machine which finds a $(2+\epsilon)$ -approximation of Min Cut with high probability.*

To track singleton cuts, the first step is to find a *low depth decomposition* of the current MST. At a high level, a low depth decomposition of a tree is a labeling of its vertices with values 1 through d , where d is the depth. This label must satisfy the following: for every level $i \in [d]$, the connected components induced on vertices

with label at least i must contain at most one vertex for each i . This defines a recursive splitting process: starting at depth 1, there must be at most one vertex v with the minimum label, so we can split the tree into multiple parts by removing v . Then we simply recurse on each connected component, considering the next set of labels, and knowing the process will always split each connected component once at a time. This is the general idea captured by both this and previous works. However, in order to increase the efficiency of this step, we require a new decomposition structure (see Definition 1) and new methods for finding the decomposition. Notice that it is always true that at each level, each connected component contains at most one vertex at the next level.

In Section 3, we show how to find a low depth decomposition with depth $O(\log^2 n)$ in AMPC in $O(1/\epsilon)$ rounds (Lemma 3) with $O(n^\epsilon)$ space per machine. Roughly speaking, we create a heavy-light decomposition of the MST, where we store "heavy paths" consisting of edges connecting vertices to their children with the largest number of descendants and isolated "light nodes". We replace each heavy path with a complete binary tree whose leaves contain the vertices in the path, which gives us an efficient structure to obtain our labeling. This yields our low depth decomposition.

In the next step, we compute the size of every singleton cut S that is created during the process. Note that the contractions are inherently sequential and the number of contractions we need to make at step k is $x_k \in O(n)$. However, each singleton cut is a connected component on the MST containing a specific edge e , whose contraction - in the increasing order of contracting MST edges - results in subset S , if we only allow the edges that have a smaller weight than e . We partition these connected components based on the vertex in the cut with the lowest level in the heavy-light decomposition of the MST. This way, we can compute every level of the low depth decomposition in parallel with only an $O(\log^2 n)$ blowup in total memory. In Section 4, we show that we can track every singleton cut in the contraction process in $O(1)$ AMPC rounds. A high level pseudocode of the main algorithm is given in Algorithm 1.

Algorithm 1: AMPC-MinCut

(An algorithm that calculates $(2+\epsilon)$ approximation of Min Cut in G . The novel part is underlined.)

Data: A graph $G = (V(G), E(G))$, a parameter k .

Result: $(2+\epsilon)$ approximation of Min Cut.

- 1 **if** $|G| \in n^\epsilon$ **then**
 - 2 | **return** *Min Cut of G calculated on a single machine*
 - 3 **end**
 - 4 Let $\widehat{G}_1, \dots, \widehat{G}_k$ be copies of G with assigned random weight on edges (independently for each copy);
 - 5 In parallel for all $i \in [k]$, $S_i \leftarrow \text{MinSingletonCut}(\widehat{G}_i)$;
 - 6 In parallel for all $i \in [k]$, $G_i \leftarrow \text{copy of } \widehat{G}_i \text{ after first } k \text{ contractions}$;
 - 7 In parallel, $C_i \leftarrow \text{AMPC-MinCut}(G_i)$;
 - 8 **return** $\min(S_1, \dots, S_k, C_1, \dots, C_k)$;
-

Note that `MinSingletonCut` (Algorithm 3) is introduced in Section 4 and it leverages `LowDepthDecomp` (Algorithm 2) from Section 3.

3 GENERALIZED LOW DEPTH TREE DECOMPOSITION

This section and the next address our algorithmic formulation and analysis. Note that all omitted proofs are deferred to the Appendix.

In order to efficiently compute the singleton cuts in parallel, we first need to compute an efficient decomposition of the MST. The low depth tree decomposition Ghaffari and Nowicki [11] introduce is a very specific decomposition with i levels such that at each level ℓ , any connected component of size s on vertices at that level or higher has a single vertex at level ℓ that separates the component into two components with size at least $s/3$ each. Unfortunately, it is unclear how to calculate this precise decomposition efficiently in AMPC. To work around this, we introduce a more generalized version of the low depth tree decomposition, show that it can be computed in AMPC, and later show that we can leverage this to obtain our Min Cut algorithm.

DEFINITION 1. A *generalized low depth tree decomposition* of some tree T is a labeling $\ell : V(T) \rightarrow [h]$ of vertices with levels for decomposition height $h \in O(\log^2 n)$ such that for each level i , the connected components induced on $T^i = \{v \in T : \ell(v) \geq i\}$ have at most one vertex labeled i each.

Notice we do not define how a level is assigned; we simply require it is assigned to satisfy the property on connected components. We describe one way to do that in this section.

To see what such a decomposition looks like, consider a process where at timestep t we look at the subgraph induced on the vertices v with $\ell(v) \geq t$ (i.e., T^t). Consider a connected component C and let v be its minimum level vertex. Then $\ell(v) \geq t$, and it is the only vertex at that level in C . At timestep $\ell(v) + 1$, C becomes separated into multiple components who all contain a vertex adjacent to v . This process defines forests with smaller and smaller trees as time passes, and eventually results in isolated vertices. The completion time of this process depends on the height of the decomposition, which in our case is $O(\log^2 n)$. We will, of course, make this more efficient in Section 4.

It is not that hard to see that Ghaffari and Nowicki's low depth tree decomposition is a specific example of generalized tree decomposition with depth $O(\log n)$. They put a single vertex in the first level and then simply recurse on the two trees in the remaining forest. Note that they require additional properties of this decomposition to obtain their result, specifically that each new component has size at least $\frac{1}{3}$ of the original component, but we will see later that these are not necessary for finding the singleton cuts.

Like Ghaffari and Nowicki in MPC, we prove this can be computed efficiently in, instead, AMPC.

LEMMA 3. *Computing a generalized low depth tree decomposition of an n -vertex tree can be done in $O(1/\epsilon)$ AMPC rounds with $O(n^\epsilon)$ memory per machine and $O(n \log^2 n)$ total memory.*

The rest of this section is dedicated to proving Lemma 3. The formal and complete algorithm is shown in Algorithm 2 and further details and definitions can be found later in this section. At a high level, our algorithm proceeds as follows:

- (1) Root the tree and orient the edges [line 2].
- (2) Contract heavy paths in a *heavy-light decomposition* of T into *meta vertices* to construct a *meta tree*, T_M [lines 3 to 5].
- (3) For each meta vertex, create a *binarized path*, a binary tree whose leaves are the vertices in the heavy path, in order. Expanding meta vertices in this manner yields our *expanded meta tree* [lines 7 to 11].
- (4) Label each vertex according to properties of the expanded meta tree. For all new vertices (i.e., vertices created in step 3) v , label v with the depth of the highest vertex u in the same meta node such that v is the leftmost leaf descending from the rightmost child of u in the binarized path of the meta node [lines 13 to 15].

Each of these steps correspond to the following subsections. For instance, step 1 corresponds to Section 3.1. All relevant terminology related to these steps are additionally found in the corresponding subsections. Lemma 3 is proven at the end of the final subsection.

Algorithm 2: LowDepthDecomp

(Computing a generalized low depth tree decomposition of an input tree in AMPC)

Data: A tree $T = (V(T), E(T))$.

Result: A mapping $\ell : V(T) \rightarrow \mathbb{N}$ of tree vertices to levels.

- 1 Initialize $\ell : V(T) \rightarrow \mathbb{N}$;
 - 2 Root and orient T ;
 - 3 Let $T_H = (V(T), \{e \in E(T) : e \text{ is heavy}\})$;
 - 4 Let \mathcal{P} be the connected components of T_H ;
 - 5 Let $T_M = (\mathcal{P}, \{(P_1, P_2) : P_1, P_2 \in \mathcal{P}, \exists (u_1, u_2) \in V(P_1) \times V(P_2) \text{ such that } (u_1, u_2) \in E(T)\})$;
 - 6 **for** $v \in T_M$ of heavy path P_v in parallel **do**
 - 7 Let $V(T_v)$ be a vertex set of size $2|P_v| - 1$ with associated indices $1, \dots, 2|P_v| - 1$, denoted by i_u ;
 - 8 Let $T_v = (V(T_v), \{(u, p_u) : i_{p_u} = \lfloor i_u/2 \rfloor\})$;
 - 9 Pre-order traverse T and sort P_v accordingly;
 - 10 Pre-order traverse T_v and let L be its sorted leaves;
 - 11 For all $i \in [|P_v|]$, map $P_v[i]$ to $L[i]$;
 - 12 **for** $u \in V(T_v)$ **do**
 - 13 Find path P^u to the root of the expanded meta-tree;
 - 14 Let $u' \in V(T_v) \cap P^u$ be such that u is the leftmost descendant of u' 's right child (otherwise $u' = u$);
 - 15 Label $\ell(u) = d(u')$;
 - 16 **end**
 - 17 **end**
 - 18 **return** ℓ limited to the original vertices in T ;
-

3.1 Rooting the Tree

Like in Ghaffari and Nowicki, the first thing we need to do in line 2 of Algorithm 2 is compute an orientation of the edges. Fortunately




Figure 1: The heavy-light decomposition of an example tree.

this, along with rooting the tree, can be done quickly in AMPC by the results of Behnezhad et al. [3] in their Theorem 7.

LEMMA 4 (BEHNEZHAD ET AL. [3]). *Given a forest F on n vertices, the trees in F can be rooted and edges can be oriented in $O(1/\epsilon)$ AMPC rounds w.h.p. using $O(n^\epsilon)$ local memory and $O(n \log n)$ total space w.h.p.*

Here, w.h.p. means “with high probability.” This completes the first step of our algorithm.

3.2 Meta Tree Construction

We also leverage Ghaffari and Nowicki’s notion of heavy-light decompositions for our AMPC algorithm, which can be found from lines 3 through 5 in Algorithm 2. This process allows us to quickly decompose the tree into a set of disjoint paths of *heavy edges*, which are defined as follows (note that our definition slightly deviates from Ghaffari and Nowicki [11], where the heavy edge must extend to the child with the largest subtree without requiring this subtree to be that large, though it is the same as the definition used by Sleator and Tarjan [19]):

DEFINITION 2 (SLEATOR AND TARJAN [19]). *Given a tree T and a vertex $v \in T$, let $\{u_i\}_{i \in k}$ be the set of children of v where the subtree rooted at u_1 is the largest out of all u_i . If there is no strictly largest subtree, we arbitrarily choose exactly one of the children with a largest subtree. Then (u_1, v) is a **heavy edge** and (u_i, v) is a **light edge** for all $1 < i \leq k$.*

Then the definition of a heavy path follows quite simply.

DEFINITION 3 (GHAFFARI AND NOWICKI [11]). *Given a tree T , a **heavy path** is a maximal length path consisting only of heavy edges in T .*

Ghaffari and Nowicki then make the observation that the number of light edges and heavy paths is highly limited in a tree. This comes from a simple counting argument, where if you consider the path from root r to some vertex v , any time you cross a light edge, the size of the current subtree is reduced by at least a factor of 2. This holds even with our different notion of heavy edges since subtrees rooted at children of light edges are still much smaller compared to the subtree rooted at the parent vertex. This bounds the number of light edges between r and v , where each pair of




Figure 2: The meta-tree of the same tree from Figure 1 is demonstrated in this figure.

light edges are separated by at most one heavy path, and therefore it also bounds the number of heavy paths.

OBSERVATION 1 (GHAFFARI AND NOWICKI [11]). *Consider a tree T oriented towards root r . For each vertex v , there are only $O(\log n)$ light edges and only $O(\log n)$ heavy paths on the path from v to r .*

Using the definition of heavy edge from Sleator and Tarjan [19] instead of from Ghaffari and Nowicki, we get an additional nice property. This is because in our definition, every internal vertex has one descending heavy edge to one child.

OBSERVATION 2 (SLEATOR AND TARJAN [19]). *Given a tree T and an internal vertex $v \in T$, v must be on exactly one heavy path. For a leaf $\ell \in T$, ℓ must be on at most one heavy path.*

Our first goal is to compute what we call a *meta tree*. This is a decomposition of our tree that will allow us to effectively handle heavy edges. It is quite analogous to Ghaffari and Nowicki’s notion of the heavy-light decomposition, which partitions the tree into heavy and light edges.

DEFINITION 4. *Given a tree T , the **meta tree** of T , denoted T_M , comes from contracting all the heavy paths in T . We call the vertices of T **original vertices** and the vertices of T_M **meta vertices**.*

Note that contracting all heavy paths simultaneously is valid because, by Observation 2, all heavy paths must be disjoint. Additionally, all internal meta vertices are contracted heavy paths (as opposed to original vertices), again by Observation 2. We note that in AMPC, since connectivity is easy, it is additionally quite easy to contract the heavy paths of T into single vertices.

LEMMA 5. *Given a tree T , the meta tree T_M can be computed, rooted, and oriented in AMPC in $O(1/\epsilon)$ rounds with $O(n^\epsilon)$ memory per machine and $O(n \log^2 n)$ total space w.h.p.*

This completes the second step of our decomposition algorithm.

3.3 Expanding Meta Vertices

In order to label the vertices, we need a way to handle the heavy paths corresponding to each meta vertex. Let $v \in T_M$ be a meta vertex, and P_v be the heavy path of original vertices in T corresponding to v . Note that we have no stronger bound on the length of a heavy path than $O(n)$. Therefore, a recursive partitioning, or labeling of vertices that has polylogarithmic depth must be able to cleverly divide heavy paths. We can do this with a new data structure.

DEFINITION 5. Given some path P , a **binarized path** is an almost complete binary tree T with $|P|$ leaves where there is a one-to-one mapping between P and the leaves of T such that the pre-order traversal of P and T limited to its leaves agree.

By “agree”, we mean that if a vertex $v \in P$ comes before a vertex $u \in P$ in the pre-order traversal of P , then it also does in the pre-order traversal of T . To characterize this tree, we make a quick observation:

OBSERVATION 3. An almost complete binary tree on n leaves has $2n - 1$ vertices, $\lfloor \log_2 n \rfloor + 1$ max depth, and every layer is full except the last, which has $2n - 2^{\lfloor \log_2 n \rfloor + 1}$ vertices.

Additionally, we can find a relationship between the ancestry of triplets in P based off of the order of the three vertices. While this is not required for expanding meta vertices, it is a property of the binarized path that will be useful when we label vertices later.

OBSERVATION 4. Given a binarized path T of a path P , for any $u, u', u'' \in P$ that appear in that order (or reversed), if v is the lowest common ancestor of u and u' and v' is the lowest common ancestor of u and u'' , then v' is an ancestor of v or $v' = v$.

To create the tree, we do the following for every $v \in T_M$:

- (1) Create an almost complete binary tree T_v with $|P_v|$ leaves, linking children to parents and noting if a vertex is a left or right child [lines 7 and 8].
- (2) Do a pre-order traversal of T_v and P_v and map the vertices in P_v to the leaves of T_v such that the pre-order traversal of P_v and of T_v limited to its leaves agree. [lines 9 to 11].

Next, it is pretty direct to see that the produced tree is a binarized path.

OBSERVATION 5. The process described above produces a binarized path T_v of P_v for all v .

We prove that this can be done in the proper constraints.

LEMMA 6. The heavy paths of a tree can be converted into binarized paths in $O(1/\epsilon)$ AMPC rounds with $O(n^\epsilon)$ local memory and $O(n \log n)$ total space w.h.p.

3.4 Labeling Vertices

Our next goal is to label the vertices with the level they should be split on. Consider, hypothetically, expanding the meta tree T_M such that every heavy path for a meta vertex v is replaced with its binarized path (which is an almost complete binary tree) T_v , and the tree continues at the leaves corresponding to the nodes in the heavy path. Note that only some vertices in the hypothetical tree correspond to vertices in the original tree T . Specifically, the internal nodes of each component subtree T_v are not vertices in T , but the leaves correspond exactly to the vertices in T .

Ultimately, for a vertex $u \in T$ in meta vertex v , let u' be the vertex in T_v such that u is the leftmost leaf-descendant of the right child of u' in T_{u_M} (or if this doesn't exist, $u' = u$). Then we will label $\ell(u) = d(u')$ where d is the depth in the expanded meta tree. Following this, our vertex labeling process will be as follows for each $v \in T_M$ and $u \in T_v$:

- (1) Vertex u finds the path P^u from u to the root of T_M , assuming the meta vertices are expanded [line 13].

- (2) Let u' be the highest vertex in T_v such that u is the leftmost descendant of the right child of u' . If there is no such vertex, let $u' = u$ [line 14].
- (3) Label u with the depth (assuming roots have depth 1) of u' in the expanded T_M [line 15].

We start by making a quick observation that comes directly from Observations 1 and 3.

OBSERVATION 6. The max depth of T_M with meta nodes expanded (“the expanded T_M ”) into binary trees is $O(\log^2 n)$.

This will be greatly helpful in showing the efficiency of our algorithm. We now show that this final part can be implemented efficiently, which is sufficient to prove our main lemma.

LEMMA 7. The process described above finds a generalized low depth tree decomposition of original tree T of height $h \in O(\log^2 n)$ in 1 round with $O(n^\epsilon)$ local memory and $O(n \log^2 n)$ total space.

4 CALCULATING THE SMALLEST SINGLETON CUT

In this section, we show a $O(1/\epsilon)$ round AMPC algorithm that executes a series of contractions and outputs the size of the smallest singleton cut that appeared during the contraction process. That is we prove the following result.

THEOREM 3. There exists an AMPC algorithm that given a graph G with unique weights on edges calculates the minimum singleton cut that appears during the contraction process in $O(1/\epsilon)$ rounds using $O(n^\epsilon)$ local memory and $O((n+m) \log^2 n)$ total space.

4.1 Contraction process

We view the contraction process of a weighted graph $G = (V, E, w : E \rightarrow [n^3])$ as a sequential process in which we iterate over multiple timesteps 0 to n^3 . For a given time i , we contract the edge e having $w(e) = i$ to a single vertex. Let G_0, \dots, G_{n^3} be the sequence of graphs created in the process, where G_0 denotes the graph before any contraction and G_{n^3} denotes the graph after all contractions. Via a quick comparison to Kruskal's algorithm, it is clear that the edges whose contraction changed the topology of the graph must belong to the minimum spanning tree of the weighted graph G (since weights are unique, the MST is unique as well). Let $T = (V, E_T, w : E_T \rightarrow [n^3])$ be the minimum spanning tree of G .

From the previous observation, it is enough to consider only contracting edges from tree T , which we will focus on in the rest of this section. It will also be convenient visualize vertices as simply being grouped instead of fully contracted.

DEFINITION 6. A **bag** of vertex v at time $t \in [n^3]$, which we denote $\text{bag}(v, t)$, is the set of vertices that can be reached from v using only edges of tree T of weight at most t . We denote $\text{nbr_bag}(v, t)$ for set of **neighbors of a bag**, that is set of these vertices u that do not belong to the bag and there exists an edge connecting u and any vertex of the bag of weight greater than t . The **degree of a bag**, denoted $\Delta \text{bag}(v, t)$, is the size of the set $\text{nbr_bag}(v, t)$.

If we proceed with our edge contraction process, where an edge with weight t is contracted at time t , then $\text{bag}(v, t)$ is the set of all vertices that have been contracted with v at time t . The value

$\Delta\text{bag}(v, t)$ is simply the degree of the vertex that corresponds to contracted vertices. Therefore, the following simple observations holds.

OBSERVATION 7. *The value of the minimum singleton cut in the contraction process of the weighted graph G is equal to the following:*

$$\min_{v \in V, t \in [n^3]} \Delta\text{bag}(v, t).$$

4.2 Simulating tree contractions with low depth decomposition

By Observation 7 our goal is to calculate the value of

$$\min_{v \in V, t \in [n^3]} \Delta\text{bag}(v, t).$$

To find this, we could calculate the value $\min_{t \in [n^3]} \Delta\text{bag}(v, t)$ for every vertex v independently in parallel. However, this would require a minimum of $\Omega(n \cdot (n + m))$ total space, which roughly corresponds to replicating the whole graphs for each independent instance. There are two key observations that will allow us to reduce the space complexity. First, bags are determined solely from the topology of tree T . Second, for larger t , it is likely the case that $\text{bag}(u, t) = \text{bag}(v, t)$, so we would like to remove this redundant computation. Therefore, we will exploit tree properties and the low depth decomposition to partition the work and avoid redundancy.

Let $\ell : V \rightarrow [h]$, $h \in O(\log^2 n)$ be the labeling from the generalized low depth decomposition of tree T (see Definition 1). Let us asses to each bag a uniquely chosen vertex.

DEFINITION 7. *The **leader** of a bag, denoted $\text{bagLeader}(v, t)$, is the vertex u with the smallest label $\ell(u)$ among all vertices from $\text{bag}(v, t)$. We define a number $\text{ldr_time}(v)$ to be the greatest number $0 \leq t' \leq n^3$ such that $\text{bagLeader}(v, t') = v$.*

Let us first argue the correctness of the above definitions.

LEMMA 8. *The leader of every bag can be determined uniquely. Also, for every vertex $v \in V$ it holds: the number $\text{ldr_time}(v)$ exists, $\text{ldr_time}(v) \geq 0$, and for every $0 \leq t' \leq \text{ldr_time}(v)$ we have that $\text{bagLeader}(v, t') = v$.*

Using the fact that each bag has exactly one leader, we can reformulate the expression $\min_{v \in V, t \in [n^3]} \Delta\text{bag}(v, t)$ as follows

$$\min_{v \in V, t \in [n^3]} \Delta\text{bag}(v, t) = \min_{v \in V} \min_{0 \leq t \leq \text{ldr_time}(v)} \Delta\text{bag}(v, t).$$

We then will distribute the work needed to calculate the right-hand side of the above equality by requiring each vertex to calculate the minimal degree among bags for which it is the leader:

$$\min_{0 \leq t \leq \text{ldr_time}(v)} \Delta\text{bag}(v, t).$$

Let i be a number in $[[\log^2 n]]$. Let L_i (the i th level) be the set of vertices $v \in V$ with low depth decomposition label $\ell(v) = i$, and $L_{\leq i}$ be that with label $\ell(v) \leq i$ (for convenience we assume that $L_{\leq 0} = \emptyset$). Let T^i be the tree T with $L_{\leq i-1}$ removed. The following observation, derived from the fact that a bag is a connected sub-graph of T and the leader has lowest value $\ell(\cdot)$, relates bag location to the topology of the low depth decomposition.

OBSERVATION 8. *For every $i \in [[\log^2 n]]$, $v \in L_i$, and $0 \leq t \leq \text{ldr_time}(v)$, the set $\text{bag}(v, t)$ belongs to a single connected component of graph T^i . For any two $u, v \in L_i$, sets $\text{bag}(u, \text{ldr_time}(u))$ and $\text{bag}(v, \text{ldr_time}(v))$ belong to different components of graph T^i .*

Recall, that we wanted to calculate the value

$$\min_{0 \leq t \leq \text{ldr_time}(v)} \Delta\text{bag}(v, t)$$

for every $v \in V$, which we rewrote as

$$\min_{v \in V} \min_{0 \leq t \leq \text{ldr_time}(v)} \Delta\text{bag}(v, t).$$

Grouping by vertices in the same layers, we get

$$\begin{aligned} & \min_{v \in V} \min_{0 \leq t \leq \text{ldr_time}(v)} \Delta\text{bag}(v, t) \\ &= \min_{i \in [[\log^2 n]]} \min_{v \in L_i} \min_{0 \leq t \leq \text{ldr_time}(v)} \Delta\text{bag}(v, t). \end{aligned}$$

By Observations 8, we can hope that computing the value

$$\min_{v \in L_i} \min_{0 \leq t \leq \text{ldr_time}(v)} \Delta\text{bag}(v, t),$$

can be done in parallel without exceeding global memory limit of $O(m \log^2 n)$, since for different $v \in L_i$, their bags up to time $\text{ldr_time}(v)$ belong to different components of T^i , thus we might avoid redundant work. The details of computing this value are presented in the next section. Let us now formalize the progress so far.

LEMMA 9. *Given a tree T and a graph $G = (V, E, w : E \rightarrow [n^3])$ as an input, calculating the value $\min_{v \in V, t \in [n^3]} \Delta\text{bag}(v, t)$ can be reduced to $O(\log^2 n)$ instances of calculating values*

$$\min_{v \in L_i} \min_{0 \leq t \leq \text{ldr_time}(v)} \Delta\text{bag}(v, t),$$

for $i \in [[\log^2 n]]$. *The reduction can be implemented in AMPC with $O(1/\epsilon)$ rounds, $O((n+m) \log^2 n)$ total space, and $O(n^\epsilon)$ local memory.*

PROOF. The correctness follows from the above discussion. For the implementation, the generalized low depth decomposition of T can be determined in $O(1/\epsilon)$ rounds with $O(n \log^2 n)$ total space by Lemma 3. Consider now $O(\log^2 n)$ tuples of format (T, ℓ, E, L_i) . Preparing them requires $O((n+m) \log^2 n)$ total space and the above discussion shows that the value

$$\min_{v \in L_i} \min_{0 \leq t \leq \text{ldr_time}(v)} \Delta\text{bag}(v, t)$$

for every $i \in [[\log^2 n]]$ can be computed from the tuple (T, ℓ, E, L_i) , thus the lemma follows. \square

4.3 Resolving the problem for vertices on the same level.

Following Lemma 9, we fix $i \in [[\log^2 n]]$ and set L_i . We calculate:

$$\min_{v \in L_i} \min_{0 \leq t \leq \text{ldr_time}(v)} \Delta\text{bag}(v, t).$$

In this approach, we will frequently query the minimum value over a path in a tree, thus the following result is helpful.

THEOREM 4 (BEHNEZHAD ET AL. [5]). *Consider a rooted, weighted tree T , the heavy-light decomposition of this tree together with an RMQ data structure that supports queries on heavy paths can be computed in $O(1/\epsilon)$ AMPC rounds using $O(n^\epsilon)$ local memory and $O(n \log n)$ total space. If the aforementioned data structures are pre-computed, then obtaining a minimum value on a path of a tree can be calculated with $O(\log n)$ queries to global memory.*

We will also make use of the following theorem.

THEOREM 5 (BEHNEZHAD ET AL. [2]). *For a given sequence of integer numbers S of length n , computing the minimum prefix sum over all prefix sums can be done in $O(1/\epsilon)$ AMPC rounds using $O(n^\epsilon)$ local memory and $O(n \log n)$ total space.*

Finally, we show that the construction of the low depth decomposition provided in Section 3 gives easy access to edges that connect vertices of higher labels with vertices of smaller labels.

LEMMA 10. *For any connected component C^i in T^i , there are at most 2 tree edges between C^i and $V \setminus T^i$ according to the low depth decomposition ℓ given in Lemma 7. Moreover, both edges can be calculated in $O(1/\epsilon)$ AMPC rounds with $O(n^\epsilon)$ memory per machine and $O(n \log^2 n)$ total memory.*

Let us now turn to the proper part of this subsection. First, we show how to compute values $\text{ldr_time}(v)$ for all $v \in L_i$.

LEMMA 11. *Given a tuple (T, ℓ, E, L_i) for tree T , low depth decomposition ℓ , set of weighted edges E , and levels L_i for some $i \in [\lceil \log^2 n \rceil]$, there exists an AMPC algorithm that calculates the value $\text{ldr_time}(v)$ for every $v \in L_i$, in $O(1/\epsilon)$ rounds using $O(n^\epsilon)$ local memory and $O((n+m) \log^2 n)$ global memory.*

PROOF. Consider vertex $v \in L_i$. Vertex v ceases to be the leader of a bag at the first time t when its bag is contracted with another bag containing at least one vertex of the set $L_{\leq i-1}$. According to the tree contraction process, time t is equal to the largest weight of tree edges between v 's connected component in graph T^i and the set of vertices $L_{\leq i-1}$. By Lemma 10, these edges can be extracted with at most $O(\log^2 n)$ queries to the low depth decomposition structure. We then simply find the minimum. Thus, all values $\text{ldr_time}(v)$ for vertices from L_i can be computed in constant number of rounds assumed the conditions stated in the lemma. \square

We can assume that values $\text{ldr_time}(v) \in L_i$ are known. We would like to efficiently compute

$$\min_{0 \leq t \leq \text{ldr_time}(v)} \Delta \text{bag}(v, t),$$

for each $v \in L_i$. For this, we make the following observation.

LEMMA 12. *Consider an edge $(x, y) =: e \in E$ and a vertex $v \in L_i$. All possible values $0 \leq t' \leq \text{ldr_time}(v)$ at which e belongs to set $\text{nbr_bag}(v, t')$ form a consecutive (possibly empty) interval of integers $[a_e, b_e] \subseteq [0, \dots, \text{ldr_time}(v)]$, called also a **time interval** with respect to v .*

PROOF. The lemma follows immediately from the fact that $\text{bag}(v, 0) \subseteq \text{bag}(v, 1) \subseteq \dots \subseteq \text{bag}(v, n^3)$. \square

Additionally, the following observation shows, given edge time intervals, how to derive $\min_{0 \leq t \leq \text{ldr_time}(v)} \Delta \text{bag}(v, t)$ and clarifies the purpose of time intervals.




Figure 3: A sample structure of an MST tree. Firm edges are tree edges, while dotted are non-tree edges. The number inside vertices denote their levels. Different colors symbolize different binarized paths. The numbers underlined are times of contraction of corresponding edges. Next to these number the time intervals of these edges with respect to vertex v are given. Since $\text{ldr_time}(v) = 2$, thus all these intervals are contained in $[0, 2]$.

OBSERVATION 9. *Fix a vertex $v \in V$ and consider time intervals $[a_e, b_e]$ with respect to v , for all $e \in E$. Denote the set of all intervals containing value x by I_x . Then, computing the value*

$$\min_{0 \leq t \leq \text{ldr_time}(v)} \Delta \text{bag}(v, t),$$

is equivalent to computing the minimum over all values $|I_x|$ for x in the range $[0, \text{ldr_time}(v)]$.

Since this task is ‘linear’, it can be computed efficiently in AMPC. We now discuss how to compute the intervals for all edges in E .

LEMMA 13. *Given a tuple (T, ℓ, E, L_i) , there exists an AMPC algorithm that for every vertex $v \in L_i$ and every edge $e \in E$ calculates the maximal, non-empty time interval $[e_a, e_b] \subseteq [0, \dots, \text{ldr_time}(v)]$ of e with respect to v . The algorithm works in $O(1/\epsilon)$ rounds, uses $O(n^\epsilon)$ local memory and $O((n+m) \log^2 n)$ global memory.*

PROOF. The algorithm starts by removing vertices $L_{\leq i-1}$ with all edges adjacent to them from tree T which gives us T^i . Given decomposition ℓ , this can be done in $O(1)$ rounds. By definition 1, vertices $L_i = \{v_1, \dots, v_q\}$ belong to different trees. Next, the algorithm roots these trees that contain vertices from L_i in v_1, \dots, v_q and calculates heavy-light decompositions of each tree together with an RMQ structure on heavy paths. By Theorem 4, this can be done in $O(1/\epsilon)$ within our memory constraints.

Let us now fix an edge $(x, y) =: e \in E$. Importantly, we consider here all edges of the graph G , not only tree edges E_T . Let $r_x \in \{\perp, v_1, \dots, v_q\}$ be the root of this tree in T^i to which the vertex x belongs. If the vertex x does not belong to any tree, that is $x \in L_{\leq i-1}$ since these are the vertices that have been removed, we write $r_x = \perp$. Let $\text{mw}(x)$ be the minimum weight over edges of path that connects vertex x with vertex r_x in graph T^i . Observe, that unless $r_x = \perp$ this value is well defined as T^i is a collection of tree and there is exactly one path connecting x and r_x in this graph. We extend the above definitions on y in the natural way.

By Theorem 4, computing $r_x, r_y, mw(x), mw(y)$ takes $O(\log n)$ queries to the memory for a single edge. Therefore, we can compute these values for all edges $e \in E$ in $O(1)$ round under the conditions assumed in this lemma.

Observe that edge $e = (x, y)$ can have non-empty time intervals only with vertices r_x and r_y . Any other vertex from L_i belongs to a different connected component in graph T_i and therefore its bag cannot contain x nor y while the vertex is the leader of its bag. Thus, all that is left to show is how $mw(x)$ and $mw(y)$ can help determine the time intervals in which edge e belongs to $nbr_bag(r_x)$ and $nbr_bag(r_y)$. We consider the following cases.

Case 1. $r_x = \perp, r_y = \perp$. In this case, edge (x, y) has no effect on degrees of bags of vertices r_x and r_y at any time. The algorithm skips such edges.

Case 2. $r_x = \perp, r_y \neq \perp$, (or symmetrically $r_x \neq \perp, r_y = \perp$). Since T^i is a subset of the minimum spanning tree T , thus the first time when vertex x belongs to r_x 's bag is the time $mw(x)$. Now, y starts to belong to r_x 's bag either at the time being equal to the maximal weight on the path between r_x and y . Observe however, that this path has to contain vertices that does not belong to T^i and therefore the maximal weight has to be greater than $ldr_time(r_x)$. What follows the correct interval in this case is:

$$[mw(x), ldr_time(r_x)],$$

or an empty interval if $mw(x) > ldr_time(r_x)$.

Case 3. $r_x \neq \perp, r_y \neq \perp$. We distinguish two sub-cases:

Subcase a) $r_x \neq r_y$. Since the path between r_x and r_y does not belong to T^i we can proceed analogously to the *Case 2.*. The correct interval for vertex x is

$$[mw(x), ldr_time(r_x)],$$

or an empty interval if $mw(x) > ldr_time(r_x)$, while for vertex y it is

$$[mw(y), ldr_time(r_y)],$$

or an empty interval if $mw(y) > ldr_time(r_y)$

Subcase b) $r_x = r_y$. Since T^i is a subgraph of the minimum spanning tree T , we have that $\min(mw(x), mw(y))$ is the first time when at least one of x and y belongs to r_x 's bag, while the first time when both belong to r_x 's bag is $\max(mw(x), mw(y))$. Thus, the proper time interval for this edge:

$$[\min(mw(x), mw(y)), \max(mw(x), mw(y))] \cap [1, \dots, ldr_time(r_x)]$$

We obtain that for every edge $e \in E$ all non-empty time intervals in which this edge belong to nbr_bag of some vertex v can be computed in $O(\log(n))$ queries to the memory. Therefore, computing these values for all edges can be done in constant number of rounds assumed $O(m \log n)$ total memory. \square

Implementing Observation 9 is purely technical.

LEMMA 14. *There exist an AMPC algorithm that given a set of integer intervals $I = \{[p_1, k_1], \dots, [p_n, k_n]\}, \forall i \in [n] [p_i, k_i] \subseteq [0, R]$ finds the minimal number of intersecting intervals in $O(1/\epsilon)$ rounds using $O(n^\epsilon)$ local memory and $O(n \log^2 n)$ total memory.*

PROOF. First, the algorithm sorts the set $\{p_1, k_1, \dots, p_n, k_n\}$ of all endpoints of these intervals in non-increasing order (ties are resolved with priority for endpoints p_i) obtaining a sequence S . Consider assigning to every endpoint $p_i, i \in [n]$ from sequence S

value $+1$ and to every endpoint $k_i, i \in [n]$ value -1 . This operation leads to a sequence S' of pairs of format (endpoint, value). Finally, let S'' be a sequence constructed from S' in which all consecutive pairs that have the same first coordinate are compressed to a single pair in which the first coordinate is preserved and the second is the sum of second coordinates of contracted pairs. It can be observed that finding the minimal prefix sum of sequence made from second coordinates of pairs in S'' is equivalent to the minimal number of intersecting intervals. The construction of sequence S'' requires only sorting and contracting consecutive pairs which can be implemented in $O(1/\epsilon)$ rounds in AMPC with the memory constraints stated in the lemma. To find the minimal prefix sum we use Theorem 5 which completes the proof. \square

The above discussion is summarized in the following lemma.

LEMMA 15. *There exists an AMPC algorithm that given a tuple (T, ℓ, E, L_i) calculates the value*

$$\min_{v \in L_i} \min_{0 \leq t \leq ldr_time(v)} \Delta bag(v, t)$$

in $O(1/\epsilon)$ rounds using $O(n^\epsilon)$ local memory and $O((n+m) \log^2 n)$ total memory.

PROOF. Using Lemma 11 we are able to calculate value ldr_time for every $v \in L_i$ in constant number of rounds. By Lemma 13 we can calculate time all non-empty time intervals for every $e \in E$ and every $v \in L_i$. This requires $O(m \log^2 n)$ total memory. Each time interval $[a, b]$ can be assigned a vertex v with respect to whom it was calculated. Then, we group time intervals with respect to vertices from L_i they were calculated. This can be done in a single round with $O(m \log^2 n)$ global memory since there are only $O(m)$ non-empty time intervals. Finally, Lemma 14 guarantees that we can compute, for every $v \in L_i$, the minimum number of intersecting intervals in $O(1/\epsilon)$ rounds with total memory proportional to the number of these intervals. Therefore, assumed $O(m \log^2 n)$ global memory we can extend the last computation to a parallel computation for $v \in L_i$ while preserving the round complexity. By Observation 9 this is equivalent to calculating

$$\min_{0 \leq t \leq ldr_time(v)} \Delta bag(v, t),$$

for every $v \in L_i$. Since the minimum of the above values over $v \in L_i$ can be computed in a single round, the lemma is proven. \square

4.4 The final algorithm.

We are now able to prove Theorem 3 and present the final algorithm, `SmallestSingletonCut`, that calculates the smallest singleton cut that appears in the contraction process of G . The pseudocode can be found in Figure 3, while the proof of correctness is below.

PROOF OF THEOREM 3. The correctness follows from Observation 7 and Lemmas 9 and 15. Also the implementation details of lines 3 – 7 are discussed in the these two lemmas. To calculate minimum spanning tree in line 1 we use Lemma 4 while the implementation of the low depth decomposition from the next line is given by Lemma 3. \square

Algorithm 3: SmallestSingletonCut**Data:** Graph $G = (V, E, w : V \rightarrow [n^3])$.**Result:** Size of the smallest singleton cut.

- 1 Compute the minimum spanning tree T of G ;
- 2 Compute the low depth decomposition D_T of T ;
- 3 Prepare $O(\log^2 n)$ tuples $(T, D_T, E, L_i), i \in [[\log^2 n]]$;
- 4 **foreach** tuple (T, D_T, E, L_i) **do**
- 5 | Compute: $lc_i \leftarrow \min_{v \in L_i} \min_{0 \leq t \leq \text{ldr_time}(v)} \Delta \text{bag}(v, t)$;
- 6 **end**
- 7 **return** $\min(lc_1, \dots, lc_{[[\log^2 n]]})$;

Michael J. Franklin, Ali Ghodsi, Joseph Gonzalez, Scott Shenker, and Ion Stoica. Apache spark: a unified engine for big data processing. *Commun. ACM*, pages 56–65, 2016.

REFERENCES

- [1] Alexandr Andoni, Aleksandar Nikolov, Krzysztof Onak, and Grigory Yaroslavtsev. Parallel algorithms for geometric graph problems. In *STOC*, pages 574–583, 2014.
- [2] Soheil Behnezhad, Laxman Dhulipala, Hossein Esfandiari, Jakub Lacki, and Vahab S. Mirrokni. Near-optimal massively parallel graph connectivity. In *FOCS*, pages 1615–1636, 2019.
- [3] Soheil Behnezhad, Laxman Dhulipala, Hossein Esfandiari, Jakub Lacki, Vahab S. Mirrokni, and Warren Schudy. Massively parallel computation via remote memory access. In Christian Scheideler and Petra Berenbrink, editors, *The 31st ACM on Symposium on Parallelism in Algorithms and Architectures, SPAA 2019, Phoenix, AZ, USA, June 22–24, 2019*, pages 59–68. ACM, 2019.
- [4] Soheil Behnezhad, Laxman Dhulipala, Hossein Esfandiari, Jakub Lacki, Vahab S. Mirrokni, and Warren Schudy. Parallel graph algorithms in constant adaptive rounds: Theory meets practice. *Proc. VLDB Endow.*, 13(13):3588–3602, 2020.
- [5] Soheil Behnezhad, MohammadTaghi Hajiaghayi, and David G. Harris. Exponentially faster massively parallel maximal matching. In *FOCS*, pages 1637–1649, 2019.
- [6] Craig Chambers, Ashish Raniwala, Frances Perry, Stephen Adams, Robert R. Henry, Robert Bradshaw, and Nathan Weizenbaum. Flumejava: easy, efficient data-parallel pipelines. In Benjamin G. Zorn and Alexander Aiken, editors, *ACM SIGPLAN*, pages 363–375, 2010.
- [7] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.
- [8] Aleksandar Dragojević, Dushyanth Narayanan, and Miguel Castro. Rdma reads: To use or not to use? *IEEE Data Eng. Bull.*, 40(1):3–14, 2017.
- [9] Aleksandar Dragojević, Dushyanth Narayanan, Miguel Castro, and Orion Hodson. Farm: Fast remote memory. In *{USENIX} ({NSDI})*, pages 401–414, 2014.
- [10] Apache Software Foundation. *Hadoop*. <https://hadoop.apache.org/>.
- [11] Mohsen Ghaffari and Krzysztof Nowicki. Massively parallel algorithms for minimum cut. In *Proceedings of the 39th Symposium on Principles of Distributed Computing*, pages 119–128, 2020.
- [12] R. E. Gomory and T. C. Hu. Multi-terminal network flows. *Journal of the Society for Industrial and Applied Mathematics*, 9(4):551–570, 1961.
- [13] Michael T Goodrich, Nodari Sitchinava, and Qin Zhang. Sorting, searching, and simulation in the mapreduce framework. In *International Symposium on Algorithms and Computation*, pages 374–383. Springer, 2011.
- [14] David R. Karger. Global min-cuts in rnc, and other ramifications of a simple min-cut algorithm. In Vijaya Ramachandran, editor, *Proceedings of the Fourth Annual ACM/SIGACT-SIAM Symposium on Discrete Algorithms, 25–27 January 1993, Austin, Texas, USA*, pages 21–30. ACM/SIAM, 1993.
- [15] David R Karger and Clifford Stein. A new approach to the minimum cut problem. *Journal of the ACM (JACM)*, 43(4):601–640, 1996.
- [16] Howard J. Karloff, Siddharth Suri, and Sergei Vassilvitskii. A model of computation for mapreduce. In *SODA*, pages 938–948, 2010.
- [17] Silvio Lattanzi, Benjamin Moseley, Siddharth Suri, and Sergei Vassilvitskii. Filtering: a method for solving graph problems in mapreduce. In *Proceedings of the twenty-third annual ACM symposium on Parallelism in algorithms and architectures*, pages 85–94, 2011.
- [18] Huzur Saran and Vijay V. Vazirani. Finding k cuts within twice the optimal. *SIAM J. Comput.*, 24(1):101–108, 1995.
- [19] Daniel Dominic Sleator and Robert Endre Tarjan. A data structure for dynamic trees. In *Proceedings of the 13th Annual ACM Symposium on Theory of Computing, May 11–13, 1981, Milwaukee, Wisconsin, USA*, pages 114–122. ACM, 1981.
- [20] Grigory Yaroslavtsev and Adithya Vadapalli. Massively parallel algorithms and hardness for single-linkage clustering under ℓ_p distances. In *ICML*, pages 5596–5605, 2018.
- [21] Matei Zaharia, Reynold S. Xin, Patrick Wendell, Tathagata Das, Michael Armbrust, Ankur Dave, Xiangrui Meng, Josh Rosen, Shivaram Venkataraman,

A AMPC ALGORITHM FOR APPROXIMATED MINIMUM k -CUT

In this section, we show that given an algorithm that calculates $2+\epsilon$ approximation of a min cut, one can construct $4+\epsilon$ approximation of minimum k -cut. Consider the following greedy algorithm, called APX-SPLIT, that extends the classic result of Saran and Vazirani [18]. The algorithm works sequentially. In each iteration, it extends the approximated solution with the smallest non-trivial approximation of the minimum cut of the graph available at a given moment. Being precise, assume that at the beginning i -th iteration the algorithm has split the graph G into ℓ_i connected components (we start with $c_1 = 1$). Then the algorithm calculates $(2+\epsilon)$ -approximation of the minimum cut in each connected component and enlarges the solution by the smallest of these cuts, thereby increasing the number of components by at least 1. The algorithm ends after the first iteration such that the number of connected components after this iteration is at least k . The pseudocode of the algorithm is given in the Algorithm 4.

Algorithm 4: APX-SPLIT

(A greedy algorithm computing an approximation of the minimum k -cut in AMPC)

Data: A graph $G = (V(G), E(G))$ and a parameter k .

Result: A $(4 + \epsilon)$ -approximation of minimum k -cut.

```

1 Initialize  $D \rightarrow \emptyset$ ;
2 while  $G' := (V(G), E(G) \setminus \bigcup_{d \in D} d)$  has less than  $k$ 
   connected components do
3   Let  $C_1, \dots, C_l$  be the set connected of components of
    $G'$ ;
4    $d_i^* \leftarrow \text{AMPC-MinCut}(C_i)$  for all  $i \in [l]$ ;
5    $j \leftarrow \arg \min_{i \in [l]} \text{weight}(d_i)$ ;
6   Add  $d_j^*$  to  $D$ ;
7 end
8 return set of cuts  $D$ ;
```

We will show by generalizing the idea of Saran and Vazirani that the aforementioned greedy algorithm is $(4+\epsilon)$ -approximation minimum k -cut.

THEOREM 2. *Algorithm APX-SPLIT is an $(4 + \epsilon)$ approximation of the Min k -Cut. Furthermore, it can be implemented in the AMPC model with $O(n^\epsilon)$ memory per machine in $O(k \log \log n)$ rounds and $O(m)$ total memory.*

We define $\text{comps}(c_1 \cup \dots \cup c_k)$ the number of components of the graph G after removing all edges from the set $c_1 \cup \dots \cup c_k$.

PROOF. The standard line of proof, proposed in [18] for the case where exact minimum cut is used at each splitting step, is to compare the cut selected by the APX-SPLIT algorithm to the approximated minimum k -cut obtained from the Gomory-Hu tree. The main difficulty is that in our case, we use *only* $(2+\epsilon)$ approximation in each splitting step. This makes our proof different and novel compared to [18]. Let us set a Gomory-Hu tree $H = (V(H), E(H)) = (V(G), E(G))$ of the graph G . The Gomory-Hu tree is defined as follows.

DEFINITION 8 (GOMORY AND HU [12]). *Consider an arbitrary graph G . A weighted tree $H = (V(H), E(H))$ with the set of vertices being equal $V(G)$ is called a Gomory-Hu tree of G , if for every pair of different vertices $s, t \in V(G)$ the minimum weight on the path between s and t in the tree H is equal to the minimum s - t cut in graph G . The existence and construction of Gomory-Hu trees was shown in [12].*

Let us order edges of the tree H (or equivalently cuts in the G graph) with respect to non-decreasing weights. Denote b_1^*, \dots, b_{l-1}^* the sequence of the first $l \leq k-1$ edges (cuts equivalently) from this order such that corresponding cuts split G graph into at least k connected components. Let b_1, \dots, b_{k-1} be a dual sequence of cuts corresponding to that sequence of edges b_1^*, \dots, b_{l-1}^* , with this addition that we put each cut b_i this number of times it increases the number of connected components in G and possibly cut some suffix of such generated sequence to obtain exactly $k-1$ cuts. For such construction we have the following.

OBSERVATION 10 (SARAN AND VAZIRANI [18]). *The sequence of cuts b_1, \dots, b_{k-1} satisfies:*

- (1) *the sequence $|b_1|, \dots, |b_{k-1}|$ is non-decreasing,*
- (2) $\forall i \in [k-1] \text{comps}(b_1 \cup \dots \cup b_i) > i$.

Saran and Vazirani also proved that such selected (and possibly refactored) sequence of cuts is $(2 - \frac{2}{k})$ -approximation of the minimum k -cut.

THEOREM 6 (SARAN AND VAZIRANI [18]). *The cut $\bigcup_{i \in [k-1]} b_i$ is $(2 - \frac{2}{k})$ approximation of the minimum k -cut of G .*

Having established the crucial properties of Gomory-Hu trees and corresponding cuts we can proceed to the proof of correctness of the APX-SPLIT algorithm.

Let d_1, \dots, d_m be the successive cuts selected by the APX-SPLIT algorithm. Note that with each new cut, at least one new component appears in the graph thus $m \leq k-1$. We will show that the sum of these cuts' sizes is not greater than the sum of sizes of cuts b_1, \dots, b_{k-1} . Let $\#c_1, \dots, \#c_m$ be a sequence of numbers where $\#c_i := \text{comps}(d_1 \cup \dots \cup d_i)$. We will show by induction that

$$\forall j \in [m] |d_1 \cup \dots \cup d_j| \leq (2 + \epsilon) |b_1 \cup \dots \cup b_{\min(k-1, \#c_j)}|.$$

The idea behind the induction step defined in the previous line can be explained as follows: inclusion of cuts from d_1 to d_i are at least $(2+\epsilon)$ approximation of cut generated by inclusion of cuts of such prefix of sequence b_1, \dots, b_k that split G on $\#c_i$ for connected components.

For the basis of induction we see that in the first step of the APX-SPLIT algorithm chooses $(2+\epsilon)$ -approximation of the smallest cut in the whole graph G . The b_1 is an s - t cut therefore we have $|d_1| \leq (2+\epsilon)|b_1|$ which implies that $|d_1| \leq (2+\epsilon)|b_1 \cup \dots \cup b_{\min(k-1, \#c_1)}|$.

Now consider $i \in [m-1]$. Since $i < m$, we observe that $\#c_i < k$. Otherwise the algorithm APX-SPLIT would have executed only $m-1$ iterations instead of m . Consider cuts $b_1, \dots, b_{\#c_i+1}$. From the Observation 10 we have that $\text{comps}(b_1 \cup \dots \cup b_{\#c_i+1}) > \#c_i$. On the other hand $\text{comps}(d_1 \cup \dots \cup d_i) = \#c_i$. Since $\text{comps}(b_1 \cup \dots \cup b_{\#c_i+1}) > \text{comps}(d_1 \cup \dots \cup d_i)$ then there must be a cut $b_j, j \in [\#c_i+1]$ that is not covered by the first i cuts from the sequence d_1, \dots, d_m .

Namely, we can choose j such that $b_j \not\subseteq d_1 \cup \dots \cup d_i$. Moreover, since b_j is an s - t cut in the graph G (with all edges included), thus it must split at least one connected component of the graph $G = (V, E \setminus (d_1 \cup \dots \cup d_i))$ into two non-empty parts. Thus this cut is considered in the $i + 1$ -th iteration of the APX-SPLIT algorithm, which implies that $|d_{i+1}| \leq (2 + \epsilon)|b_j| \leq (2 + \epsilon)|b_{\#c_{i+1}}|$. Since $\#c_i + 1 \leq \#c_{i+1}$, we conclude that $|d_1 \cup \dots \cup d_{i+1}| \leq (2 + \epsilon)|b_1 \cup \dots \cup b_{\min(k-1, \#c_{i+1})}|$, which proves the inductive step. Now, we see from Theorem 2 that the solution of b_1, \dots, b_{k-1} is an $(2 - \frac{2}{k})$ -approximation of the minimum k -cut. Thus the solution $d_1 \dots d_{k-1}$ is $(2+\epsilon)(2-\frac{2}{k}) =$ approximation of the minimum k -cut. This proves the correctness of the algorithm.

It remains to be noted that a single iteration of the algorithm can be performed in $O(\log \log n)$ rounds in the AMPC model with $O(n^\epsilon)$ memory per machine and in total memory $O(m)$. The dominant operation is the calculation of a $(2 + \epsilon)$ approximation of the minimum cut in each of the components. Its performance is analyzed in Theorem 1. The calculation of the smallest of all approximated cuts corresponding to different components is a standard operation and can be performed in $O(1)$ rounds. Also, Behnezhad et al. in [4] showed that the number of components of a graph can be determined in $O(1)$ rounds in AMPC with $O(n^\epsilon)$ memory per machine and $O(m)$ total memory. This completes the performance analysis of the algorithm. \square

B MISSING PROOFS FROM SECTION 3

PROOF OF LEMMA 5. First off, we know subtree size can be computed in $O(1/\epsilon)$ low-memory AMPC rounds on trees as shown by Behnezhad et al. [3], and the child of a vertex v with minimum subtree can then be found in $O(1/\epsilon)$ rounds by dividing the children amongst machines and iteratively finding the smallest. Next, consider removing all light edges from T to create a forest F . Run Behnezhad et al.'s AMPC connectivity algorithm [4], which satisfies the round and space constraints, to identify the components and contract them. Add the light edges back in to connect contracted nodes. This clearly is T_M . Additionally, as before, we run Behnezhad et al.'s [3] AMPC algorithm for orienting the tree. This too falls within the constraints. \square

PROOF OF OBSERVATION 4. Consider such a $u, u', u'' \in P_v$ that appear in this order (or reversed), and let v and v' be the lowest common ancestors of u and u' , and u and u'' respectively. Consider traversing from u up the tree from child to parent, and let p be the current vertex. We start with $p = u$ and thus the leaf set of the subtree rooted at p is $L_p = \{u\}$. As we traverse upwards, we add sets of leaves to L_p that are contiguous in P_v . Additionally, one vertex is directly adjacent to a vertex from L_p in P_v because P_v was mapped to the leaves of T_v according to the pre-order traversal. Therefore, L_p is a contiguous chunk of P_v . Thus, when u'' gets added to the subtree (i.e., when $p = v'$) u' must be in L_p too, either because it was added previously or it is being added at the same time. In the former case, v must have happened before v' and thus v' is an ancestor of v , and in the latter case, $v = v'$. \square

PROOF OF LEMMA 6. Correctness of the process described in this section is seen in Observation 5. Thus we simply need to show how to implement it in AMPC. For the first step, we must construct a

generic almost complete binary tree with $|P_v|$ leaves. Call this tree T_v . By Observation 3, this has $2|P_v| - 1$ vertices, $\lceil \log_2 |P_v| \rceil + 1$ max depth, and each layer is full except the last which has $2|P_v| - 2^{\lceil \log_2 |P_v| \rceil + 1}$ vertices. Thus, it is fairly simple to, in parallel, create the set of all vertices in the tree and then connect each vertex to its parent. Each vertex can be given an index: a unique identifier for vertices numbered $1, \dots, 2|P_v| - 1$. This is going to represent the order of the vertices in a breadth-first traversal of the tree. For a vertex with index $1 < i < 2|P_v|$, its parent's index j can be computed as $j = \lfloor i/2 \rfloor$. Whether or not a vertex is a left or right child is simply determined by the parity of its pre-order index.

Note that the size of the tree is $O(|P_v|)$ but each individual processor computation (i.e., computing the size of the tree, and then having each index connect itself to its parent) can be done in constant space and 2 rounds. Thus in 2 rounds, we can create such a tree. Note that we have to construct these trees in parallel, but it is not hard to see that this will only require $O(n \log n)$ total space which can be divided appropriately amongst machines.

For the second step, we can use Behnezhad et al.'s [3] algorithm for pre-order numbering with $O(1/\epsilon)$ AMPC rounds w.h.p. using $O(n^\epsilon)$ local memory and $O(n)$ total space. Let L be a list of the leaves of T_v in pre-order. To map P_v to the leaves, one can simply do a direct map between P_v and L in one round. \square

PROOF OF LEMMA 7. First, note that P^u can be stored entirely on one machine by Observation 6, and additionally, since both T_M and T_v for all $v \in M$ is oriented, it is quite simple to adaptively query the path from u to the root in one round within the space constraints. Assuming the orientations also labels if the vertex is a left or right child, u' can be found simply by searching the path. Finally, the depth of u' , which is the label of u , can also be found quite simply given access to all of P^u .

It is quite simple to show the height is bounded by $O(\log^2 n)$: all labels are depths in the expanded T_M and Observation 6 bounds the max depth. All that is left is to show the connected components induced on $T^i = \{v \in T : \ell(v) \geq i\}$ contain exactly one vertex with label i each.

We show this by induction on the level. At the 1st level, we should only partition the graph once. Let u be a vertex labeled 1 and P^u be its path. For u to be given depth 1, it must have received its label from the root r_M of the expanded T_M , since we are counting depth starting at 1. Thus, it must be the leftmost descendant of the right child of r_M . This is clearly unique, thus u is unique. Therefore, there is exactly one vertex at the 1st level.

Consider a connected component C in $T^i = \{v \in T : \ell(v) \geq i\}$ for some level $i \in h$. Let the "neighborhood" $N(C)$ be all the vertices in $T \setminus C$ that are adjacent to some vertex in C . We will first show that for the largest level j of a vertex in $N(C)$, there is exactly one vertex in $N(C)$ of level j . Note $j < i$, otherwise a vertex of level j in $N(C)$ would actually be included in the component C .

Assume for contradiction there are at least two such vertices, $u, v \in N(C)$ with level j . Let u' and v' be their respective neighbors in C . Since C is a connected component, there is a path P from u' to v' containing only vertices in C . All of these vertices have level i or higher by the definition of C . Tacking on u and v to the start and end of P respectively, there is a path from u to v consisting of vertices $\{u, v\} \cup P$. Since u and v are at level j , that means every

vertex in this path has level j or higher. Therefore, u and v must have been in the same connected component C' in the earlier level T^j . By induction, that component must have had only one vertex at level j . This is a contradiction. Thus $N(C)$ must have exactly one vertex in level j .

Let $v \in N(C)$ be the vertex at level j , and let u be its neighbor in C (note there can only be one since T is a tree). We consider three cases.

Case 1. u is a child of v and they are not on the same heavy path. Let u_M and v_M be the meta vertices in T_M containing u and v respectively, u' and v' be the corresponding nodes found in step 2 for u and v (whose depths are the labels of u and v), and r_{u_M} be the root of the binarized path for u_M . Clearly, u_M is a child of v_M since u is a child of v and $u_M \neq v_M$. Moreover, r_{u_M} must be the child of v in the expanded T_M . Since v' is an ancestor of v , that means its depth in the expanded T_M satisfies $d(v') < d(r_{u_M})$. Since this defines the label of v' , $\ell(v) \leq d(r_{u_M}) - 1$.

Consider any leaf $l \in T_{u_M}$. Its label is the depth of some vertex $w \in T_{u_M}$, which must be deeper than the root r_{u_M} . Thus $\ell(l) \geq d(r_{u_M}) \geq \ell(v) + 1$, then implying all of T_{u_M} is in C . Additionally, since only one $l \in T_{u_M}$ is labeled by r_{u_M} , i.e. $\ell(l) = d(r_{u_M})$, it must have a unique (and smallest) label out of all vertices in P_{u_M} . It turns out this $\ell(l)$ will be our j .

Now, consider l 's placement in the original tree T . It is on a heavy path containing u , and it has the smallest label on the heavy path. Vertex u is also directly adjacent to v . Thus in level i when we consider the tree induced on T^i , it must be in the same component as u since the path from l to u is contained in T^i . Additionally, since T is a tree, the only vertex in $N(C)$ above the component itself is v . Therefore, all of C is a descendant of v , and thus they must also be descendants of u . Thus, in T_M , they must have been in meta vertices at the depth of or below u_M . By a similar logic as before, their label must be strictly greater than l 's label. This implies that $\ell(l) = j$, and it is the only vertex in T_{u_M} with such a label.

Case 2. u is the parent of v and they are not on the same heavy path. Using u_M and v_M as before and with the same logic as before but with reversed roles, we find that $\ell(v) > \ell(u)$ so $\ell(u) < i$. This contradicts that u is in T^i , and thus this case is impossible. Note that this also shows that C cannot contain a vertex in an ancestor u_M of v_M in T_M , because by connectivity, this would imply that there is some u that is the parent of the root of T_{v_M} . A similar argument will hold.

Case 3. u and v are on the same heavy path corresponding to meta vertex u_M whose binarized path T_{u_M} has root r_{u_M} . Let $P \subseteq P_{u_M}$ be the connected subpath of this heavy path that contains u if we remove all vertices of level $i - 1$ or lower. Obviously, $P = C \cap P_{u_M} \subseteq C$. Assume for contradiction that two vertices $p, p' \in P$ have level i , so $\ell(p) = \ell(p') = i$. Let $a = lca_M(p, p')$ be the least common ancestor of p and p' in T_{u_M} . Because a is the least common ancestor, one of its children must contain p and the other p' in its subtree. Without loss of generality, assume p is a descendant of the left child and p' is a descendant of the right child.

Let w and w' be the (possibly internal) vertices of T_{u_M} such that p and p' are labeled with their expanded meta tree depths in step 3 respectively ($\ell(p) = d(w)$ and $\ell(p') = d(w')$). This means

$i = d(w) = d(w')$. Also, notice that w is an ancestor of p and w' is an ancestor of p' , so they are on the respective paths from p and p' to the root of T_{u_M} , call it r_{u_M} . Additionally, a is on both paths, and specifically the paths must merge at a . Since w and w' only pass their depth to exactly one leaf of T_{u_M} each, they must be distinct. Therefore, they cannot be at depth $d_{u_M}(a)$ or lower, else they would both be a . Thus, w and w' are strict descendants of a , so $d(w), d(w') > d(a)$, meaning that $\ell(p), \ell(p') > d(a)$.

Let \bar{p} be the leftmost leaf descendant of the right child of a . Then \bar{p} was labeled by a , thus $\ell(\bar{p}) = d(a)$. Therefore, $\ell(p), \ell(p') > \ell(\bar{p})$. Since p and p' are at level i , $\ell(\bar{p}) < i$. Additionally, note that by the structure of T_{u_M} , \bar{p} is necessarily located between p and p' in the leaf set, thus it must come between them in the pre-order traversal of T_{u_M} by Observation 4. Since T_{u_M} is the binarized path of P_{u_M} , this also implies \bar{p} comes between p and p' in P_{u_M} , and by extension P . Note, however, that $P \subseteq C \subseteq T^i$, and thus its vertices must all have level i or greater. This contradicts that $\ell(\bar{p}) < i$. Thus, it must be the case that there is at most vertex in P that has level i .

Finally, all that needs to be shown is that for every vertex $c \in C \setminus P$, $\ell(c) > i$. For any such c , let c^* be its lowest ancestor in P_{u_M} . Since $\ell(c^*) \geq i$, we can make the same argument as Case 1 to show that c must have label $i + 1$ or larger. \square

PROOF OF LEMMA 3. By combining Lemmas 4, 5, 6, and 7, we have our result. \square

C MISSING PROOFS FROM SECTION 4

PROOF OF LEMMA 8. Consider $\text{bag}(v, t)$ for some $v \in V$, $t \in [n^3]$. Because $v \in \text{bag}(v, t)$, this bag cannot be empty. Let u be the vertex with the smallest label among all vertices from $\text{bag}(v, t)$. Assume for contradiction that there exists another vertex $u' \in \text{bag}(v, t)$ such that $\ell(u) = \ell(u')$. Observe, that there exists only one path between u and u' in tree T and all vertices of this path must be contained in $\text{bag}(v, t)$ since $\text{bag}(v, t)$ always forms a connected component when viewed as a subtree of T , c.f. Definition 6. From the properties of the low depth decomposition, we get that there is a vertex z on the path with label smaller than both $\ell(u)$ and $\ell(u')$ which contradicts with the choice of u . This also proves the uniqueness of the leader. Let us now argue about $\text{ldr_time}(v)$. First, for every $v \in V$ we have that $\text{bag}(v, 0) = v$ as there is no edge with weight smaller than 1. This already proves two first properties. Second, it holds: $\text{bag}(v, 0) \subseteq \text{bag}(v, 1), \dots, \text{bag}(v, n^3)$ as a bag defined for larger time can expand more edges. Therefore, if v is the leader of $\text{bag}(v, t)$ for some $0 \leq t \leq n^3$, it has to be the leader of any subset of this bag. \square

PROOF OF LEMMA 10. Let v be a vertex of this component with lowest value $\ell(v)$. By the construction of the low depth decomposition, the component C_i consists of some connected part P_1 of binarized path P to which v belongs and all binarized paths incident to the part F that contain vertices with larger values ℓ , denote this set P_2 . Note, that each vertex $v \in P_2$ has no edged to the part of graph $V \setminus T^i$. Thus, only vertices from part P_1 can be connected with vertices of smaller labels than i . However, since P_1 forms a path, and we consider only tree edges, then there can be at most two such edges.

The above proof instructs also how to compute these two edges. First, the vertex v can be determined in $O(1/\epsilon)$ rounds, since it requires computing max function only over labels of vertices belonging to C_i . Assume, that for each vertex of a tree, we store not only its value $\ell(v)$, but also its position and the length of the binarized path to which it belongs. Then the first vertex to the left with label smaller than v and the first vertex to the right with label smaller

than v on v 's binarized path can be computed in constant time in local memory of a single machine, since their positions in the binarized path are functions of only the length of the path and the position of v . Assumed, that in the global memory all binarized paths are stored, then the corresponding edges can be found in $O(1)$ queries to the global memory. \square