



This is a repository copy of *Weighted min-cut: sequential, cut-query, and streaming algorithms*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/200975/>

Version: Accepted Version

Proceedings Paper:

Mukhopadhyay, S. orcid.org/0000-0002-3722-4679 and Nanongkai, D. orcid.org/0000-0003-4468-2675 (2020) Weighted min-cut: sequential, cut-query, and streaming algorithms. In: Makarychev, K., Makarychev, Y., Tulsiani, M. and Kamath, G., (eds.) STOC 2020: Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing. STOC '20: 52nd Annual ACM SIGACT Symposium on Theory of Computing, 22-26 Jun 2020, IL, Chicago, USA. Association for Computing Machinery (ACM) , pp. 496-509. ISBN 9781450369794

<https://doi.org/10.1145/3357713.3384334>

© ACM 2020. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in STOC 2020: Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing <http://dx.doi.org/10.1145/3357713.3384334>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Weighted Min-Cut: Sequential, Cut-Query and Streaming Algorithms

Sagnik Mukhopadhyay ^{*} Danupon Nanongkai [†]

Abstract

Consider the following *2-respecting min-cut* problem. Given a weighted graph G and its spanning tree T , find the minimum cut among the cuts that contain at most two edges in T . This problem is an important subroutine in Karger’s celebrated randomized near-linear-time min-cut algorithm [STOC’96]. We present a new approach for this problem which can be easily implemented in many settings, leading to the following randomized min-cut algorithms for weighted graphs.

- An $O\left(m \frac{\log^2 n}{\log \log n} + n \log^6 n\right)$ -time sequential algorithm: This improves Karger’s long-standing $O(m \log^3 n)$ and $O\left(m \frac{(\log^2 n) \log(n^2/m)}{\log \log n} + n \log^6 n\right)$ bounds when the input graph is not extremely sparse or dense. Improvements over Karger’s bounds were previously known only under a rather strong assumption that the input graph is *simple* (unweighted without parallel edges) [Henzinger, Rao, Wang, SODA’17; Ghaffari, Nowicki, Thorup, SODA’20]. For unweighted graphs (possibly with parallel edges) and using bit operations, our bound can be further improved to $O\left(m \frac{\log^{1.5} n}{\log \log n} + n \log^6 n\right)$.
- An algorithm that requires $\tilde{O}(n)$ *cut queries* to compute the min-cut of a weighted graph: This answers an open problem by Rubinfeld, Schramm, and Weinberg [ITCS’18], who obtained a similar bound for simple graphs. Our bound is tight up to polylogarithmic factors.
- A *streaming* algorithm that requires $\tilde{O}(n)$ space and $O(\log n)$ passes to compute the min-cut: The only previous non-trivial exact min-cut algorithm in this setting is the 2-pass $\tilde{O}(n)$ -space algorithm on simple graphs [Rubinfeld et al., ITCS’18] (observed by Assadi, Chen, and Khanna [STOC’19]).

Our approach exploits some cute structural properties so that it only needs to compute the values of $\tilde{O}(n)$ cuts corresponding to removing $\tilde{O}(n)$ pairs of tree edges, an operation that can be done quickly in many settings. This is in contrast to the techniques used by Karger and Lovett-Sandlund to solve 2-respecting min-cut where information about many more cuts is computed, stored in and accessed from sophisticated data-structures.

^{*}KTH Royal Institute of technology, Sweden, sagnik@kth.se

[†]KTH Royal Institute of technology, Sweden, danupon@kth.se

Contents

1	Introduction	1
1.1	Technical Overview	4
1.2	Organization	8
2	Preliminaries	9
3	A schematic algorithm for 2-respecting min-cut	9
3.1	When spanning tree is a path	10
3.2	Interesting edges and paths	15
3.3	Handling general spanning tree	16
4	Karger’s tree packing	19
5	Cut-query & streaming algorithms	21
5.1	Graph cut-query upper bound	21
5.2	Streaming upper bound	24
6	Sequential algorithm for 2-respecting min-cut	26
6.1	Range counting, sampling & searching	26
6.2	Main algorithm: Proof of Theorem 6.1	30
7	Open problems	37
A	Proof of Claim 3.1	40

1 Introduction

Min-Cut. Given a weighted graph G , a *cut* is a set of edges whose removal disconnects the graph. The *minimum cut* or *min-cut* is the cut with minimum total edge weight. The min-cut problem—finding the min-cut—is a classic graph optimization problem with countless applications. A long line of work spanning over many decades in the last century was concluded by the STOC’95 $O(m \log^3 n)$ -time randomized algorithm of Karger [Kar00]. (Throughout, we n and m denote the number of nodes and edges respectively.) Karger’s approach remains the only approach to solve the problem in near-linear time for general graphs. With more tricks, he improved the running time further to $O(m(\log^2 n) \log(n^2/m)/\log \log n + n \log^6 n)$, suggesting the possibility to improve the logarithmic factors further. He also suggested a few approaches that might improve the running time by a $\log n$ factor. Nevertheless, after more than two decades no further improvement was known.

The only improvements over Karger’s bound known to date are under a rather strong assumption that the input graph is *simple*, i.e. it is unweighted and contains no parallel edges. The first such bound was by the $O(m \log^2 n \log \log^2 n)$ -time deterministic algorithm of Henzinger, Rao and Wang [HRW17] (following the breakthrough $O(m \log^{12} n)$ -time deterministic algorithm of Kawarabayashi and Thorup [KT19]). More recently Ghaffari, Nowicki and Thorup [GNT20] improved this bound further to $O(m \log n)$ and $O(m + n \log^3 n)$ with randomized algorithms.

By restricting to simple graphs, progress has been also made in other settings: (I) In the *cut-query* setting, we are allowed to query for the value of a cut specified by a set of nodes. The goal is to compute the minimum cut value. Naïvely, one can make $O(n^2)$ cut queries to reconstruct the graph itself. Rubinfeld, Schramm and Weinberg [RSW18] showed a randomized algorithm that only needs $\tilde{O}(n)$ queries, but their algorithm works only on simple graphs; here \tilde{O} hides polylogarithmic factors. (II) In the *multi-pass semi-streaming* setting, an algorithm with $\tilde{O}(n)$ space reads the input graph in multiple passes, where in each pass it reads one edge at a time in an adversarial order. Assadi, Chen and Khanna [ACK19] observed that the algorithm of Rubinfeld *et al.* can be adapted to solve the problem in two passes. When it comes to non-simple graphs, no non-trivial algorithms were known in both settings.

2-Respecting Min-Cut. The main bottleneck in obtaining similar results on non-simple graphs is the lack of efficient algorithms for the *2-respecting cut problem*. In this problem, we are given a spanning tree T of G and have to find the minimum cut in G among the cuts that contain at most two edges in T (we say that such cuts *2-respect* T).

Solving this problem is a core subroutine and the bottleneck in Karger’s algorithm. Using tree packing, Karger proved that if $T_{rsp}(m, n)$ is the time needed to find the 2-respecting min-cut, then the min-cut can be found in

$$\begin{aligned} T_{cut}(m, n) &= O(T_{rsp}(m, n) \log n + m + n \log^3 n) \text{ and} \\ T_{cut}(m, n) &= O\left(T_{rsp}(m, n) \frac{\log n}{\delta \log \log n} + n \log^6 n + m \log^{1+\delta} n\right) \end{aligned} \quad (1)$$

time for any small $\delta > 0$. His $O(m \log^3 n)$ bound for min-cut was obtained by using a sophisticated dynamic programming algorithm to find 2-respecting min-cut in $T_{rsp}(m, n) =$

Reference	Complexity	Remark
$m \geq n \log^4 n$		
[Kar00]	$O(m \frac{\log^2 n \log(n^2/m)}{\log \log n} + n \log^6 n)$	Old record
Here, [GMW20]	$O(m \frac{\log^2 n}{\log \log n} + n \log^6 n)$	New record
Here	$O(m \frac{\log^{3/2} n}{\log \log n} + n \log^6 n)$	Unweighted
$m \leq n \log^4 n$		
[Kar00]	$O(m \log^3 n)$	Old record
[GMW20]	$O(m \log^2 n)$	New record
Simple graphs (unweighted, no parallel edges)		
[GNT20]	$O(\min(m + n \log^3 n, m \log n))$	Old record
[GMW20]	$O(\min(m + n \log^2 n, m \log n))$	New record

Table 1: Results for sequential min-cut

$O(m \log^2 n)$ time. (A simplification of this algorithm is presented in [LS19].) With more tricks he showed that $T_{rsp}(m, n) = O(m(\log n) \log(n^2/m))$, leading to the bound of $T_{cut}(m, n) = O\left(m \frac{\log^2 n \log(n^2/m)}{\delta \log \log n} + n \log^6 n + m \log^{1+\delta} n\right)$.¹ Using efficient cut sparsifier algorithms that exist in the cut-query and streaming setting (e.g. [AGM12b, AGM12a, GKP12, RSW18]), it is also not hard to adapt Karger’s arguments to the cut-query and streaming settings, leaving the 2-respecting min-cut as the main bottleneck.

Our Results. We show a new way to solve the 2-respecting min-cut problem which can be turned into efficient algorithms in many settings. In the sequential setting, our algorithm takes $T_{rsp}(m, n) = O(m \log n + n \log^4 n)$ time and is correct with high probability. For unweighted graphs (possibly with parallel edges), our bound can be improved to $T_{rsp}(m, n) = O(m\sqrt{\log n} + n \log^4 n)$ when bit operations are allowed. Plugging this in Equation (1), we achieve the first improvement over Karger’s bounds on general graphs.

Theorem 1.1. *A min-cut can be found w.h.p. in $O(m \frac{\log^2 n}{\log \log n} + n \log^6 n)$ time. For unweighted graphs (possibly with parallel edges) and with bit operations, the bound can be improved to $O(m \frac{\log^{3/2} n}{\log \log n} + n \log^6 n)$.*²

Independent work and comparisons: The recent independent work of Gawrychowski, Mozes, and Weimann [GMW20] shows that the 2-respecting min-cut and global min-cut can be

¹It is to be noted that Karger’s result states $T_{cut}(m, n) = O(T_{rsp}(m, n) \log n / \log \log n + n \log^6 n + m \log^2 n)$. But, with a more careful analysis, the $m \log^2 n$ factor in the T_{cut} expression can be made $m \log^{1+\delta} n$ for arbitrarily small positive δ without severely affecting other terms of the expression. To achieve this, simply set the threshold for α to be $1/(4 \log^\delta n)$ in the proof of Lemma 9.1 of [Kar00]. It can be checked that the tree-packing can be done in time $O(n \log^6 n)$ with these new parameters.

²Recall that with high probability (w.h.p.) means with probability at least $1 - 1/n^c$ for an arbitrarily large constant c .

Reference	Complexity	Type of graphs
Cut-query algorithm		
[RSW18]	$\tilde{O}(n)$	Simple
Here	$\tilde{O}(n)$	Weighted
Dynamic streaming algorithm		
[RSW18, ACK19]	$\tilde{O}(n)$ -space, 2-pass	Simple
Here	$\tilde{O}(n)$ -space, $O(\log n)$ -pass	Weighted

Table 2: Our results in cut-query and streaming model and comparison with other works

solved in $O(m \log^2 n)$ time. This also improves the $O(m + n \log^3 n)$ bound of Ghaffari et al. [GNT20] to $O(m + n \log^2 n)$. Table 1 compares Gawrychowski et al.’s result with ours and Karger’s. For most values of m , their and our results are similar and improve Karger’s result. For unweighted graphs, our bound can be improved further slightly. When the graph is very sparse, Gawrychowski et al.’s bound improves Karger’s (our result does not). When $m = \Omega(n^2)$, both results do not improve Karger’s result, except for unweighted graphs where we provide a $\sqrt{\log n}$ improvement.

It is quite easy to implement our algorithm to find a 2-respecting min-cut using $\tilde{O}(n)$ cut queries in the cut-query model, and $O(\log n)$ passes with $\tilde{O}(n)$ internal memory in the semi-streaming settings respectively. Extending these results to find a min-cut (using Karger’s tree packing approach) does not increase any additional overhead. Note that our streaming algorithm works even in the *dynamic streaming* setting where there are both edge insertions and deletions.

Theorem 1.2. *A min-cut can be found w.h.p. using $\tilde{O}(n)$ cut queries.*

Theorem 1.3. *A min-cut can be found w.h.p. by an $O(\log n)$ -pass $\tilde{O}(n)$ -space dynamic streaming algorithm.*

Rubinfeld *et al.* [RSW18] asked whether it is possible to compute the min-cut on weighted graphs with $o(n^2)$ cut queries. (They showed an $\tilde{O}(n)$ bound for simple graphs (unweighted without parallel edges).) Our cut-query bound answers their open problem positively. Using a connection to communication complexity pointed out by Rubinfeld *et al.* [RSW18], we can prove an $\tilde{\Omega}(n)$ lower bound for the number of cut queries, making our result in Theorem 1.2 tight up to polylogarithmic factors. The $\tilde{O}(n)$ space bound in Theorem 1.3 is also tight [FKM⁺05]. As mentioned earlier, the only non-trivial algorithm in the streaming setting (even with insertions only) is the $\tilde{O}(n)$ -space 2-pass algorithm [RSW18, ACK19]. Note that there are some $(1 + \epsilon)$ -approximation $\tilde{O}(n)$ -space 1-pass algorithms in the literature (e.g. [KLM⁺17, AGM12a, AGM12b]). Computing min-cut exactly in a single pass with $\tilde{O}(n)$ space is however impossible [Zel11].

Updates since our work. A recent result [GMW21] shows a deterministic implementation of our schematic algorithm for 2-respecting cut with running time $T_{rsp}(m, n) = O(m \log n + n \log^2 n)$ in the sequential model. By slightly altering the *range search* data-structure that

we use in this paper (explained in the Section 1.1 and 6.1), they also obtain a runtime of $T_{rsp}(m, n) = O(\frac{m}{\epsilon} + \frac{n^{1+2\epsilon} \log n}{\epsilon^2} + n \log^2 n)$.³ Below we summarize the state of the art in the sequential model when this result is taken into account.

- When $m = O(n \log^2 n)$, we have $O(m \log^2 n)$ runtime [GMW20].
- When $m = \omega(n \log^2 n)$, we have $O(\frac{m \log n}{\epsilon} + \frac{n^{1+2\epsilon} \log^2 n}{\epsilon^2} + n \log^3 n)$ runtime (this paper together with [GMW21]).⁴

Our framework has also been implemented in the distributed and parallel settings, leading to a near-optimal distributed algorithm in the CONGEST model [DEM21] and the first work-optimal parallel algorithm [LMN21].

1.1 Technical Overview

Our approach exploits some cute structural properties so that we can limit the number of pairs of tree edges we include in our search space. This is in contrast to the techniques used by Karger [Kar00] and Lovett-Sandlund [LS19] to solve 2-respecting min-cut where information about many more cuts is computed—either using dynamic programming or using heavy-light decomposition—and stored in and accessed from sophisticated data-structures (Lovett-Sandlund additionally show that augmented binary search trees suffice, at the cost of an extra $O(\log n)$ factor).

We now explain the main ideas behind our algorithm. Let G and T be the input graph and spanning tree respectively. For any pair of edges e and f in T , let $\text{cut}(e, f)$ be the number of edges (u, v) where the unique uv -path in T contains *exactly one* of e and f . It is well-known that if a 2-respecting min-cut C contains e and f , then it has weight $\text{cut}(e, f)$. This is because the end-vertices of e has to be in different connected components when we remove edges in C ; otherwise, we can get a smaller cut by removing e from C . (The same applies for f .)

Our algorithm exploits the fact that $\text{cut}(e, f)$, for a pair of tree-edges (e, f) , can be computed efficiently in many settings. For example, in the streaming setting, we only need to make a pass over all edges using $O(\log n)$ -bit space (in addition to the space to keep T). We also require only one cut query since we know how nodes are partitioned (e.g. end-vertices of e should be in the different connected components after the graph is cut). In the sequential setting, we can use the *2-d orthogonal range counting* data structure: We assign nodes with integers based on when they are visited in a post-order traversal on T , and map each edge to a two-dimensional point based on the numbers assigned to their end-vertices. It is not hard to see that we can compute $\text{cut}(e, f)$ if we know the number of points in a few rectangles. This information can be provided by the 2-d orthogonal range counting data structure. Using Chazelle’s semigroup range search data structure [Cha88], this takes $O(m \log m)$ preprocessing time and $O(\log n)$ amortized time to compute each $\text{cut}(e, f)$. With

³This needs a small change from [GMW21], where we need to store T_S in [GMW21] as a compacted trie. We thank Paweł Gawrychowski for the clarification.

⁴When $m = n \log^2(n)g(n)$ for some growing function g , we have $\frac{m \log n}{\epsilon} + \frac{n^{1+2\epsilon} \log^2 n}{\epsilon^2} + n \log^3 n < m \log^2 n$ by setting ϵ to $c(\log \log(g(n)))/\log n$ for some constant c .

bit operations, the pre-processing time can be improved to $O(m\sqrt{\log m})$ using, e.g., Chan and Patrascu’s range counting data structure [CP10].

By the above facts, we can naïvely find the 2-respecting min-cut by computing $\text{cut}(e, f)$ for all n^2 tree-edge pairs (e, f) . By exploiting some structural properties explained below, we can reduce the number of pairs in our search space to $\tilde{O}(n)$. Below we focus on the number of *probes*—the number of tree-edge pairs (e, f) that we have to compute $\text{cut}(e, f)$ for. Since all other steps can be performed efficiently in all settings we consider, we obtain our results.

When T is a path. To explain the structural properties that we use, let us consider some extreme cases. The first case is when T is a path. It is already very unclear how to solve this case efficiently. (In fact, the starting point of this work was the belief that this case requires $\tilde{\Omega}(n^2)$ probes, which is now proven wrong.) The key insight is the structure of the *special case* of the problem, where we assume that the tree-edge pair (e, f) that minimizes $\text{cut}(e, f)$ is on the *different sides* of a given node r . More precisely, assume that, in addition to the path T , we are given a node r . For convenience, we use e'_i and e_i to denote the i^{th} edges of T to the left and right of r respectively (see the graph in Figure 1 for an example). Now consider computing $\min_{(e'_i, e_j)} \text{cut}(e'_i, e_j)$. To do this, define a matrix M whose entry at the i^{th} row and j^{th} column is $M[i, j] = \text{cut}(e'_i, e_j)$ (see the matrix in Figure 1 for an example). Computing $\min_{(e'_i, e_j)} \text{cut}(e'_i, e_j)$ becomes computing $\min_{i,j} M[i, j]$. The key to solve this problem is the following property.




Figure 1: Example of when T is a path. In the graph, $e_1, \dots, e_5, e'_1, \dots, e'_5$ are edges in T and have weight 0. Numbers on other edges indicate the edge weights.

Observation 1.4. *Assume for simplicity that the minimum entry in each column of M is unique. For any j , let $r^*(j) = \arg \min_i M[i, j]$. Then, $r^*(j) \geq r^*(j - 1)$ for any j . That is, r^* is a non-decreasing function.*

For example, the minimum entries of each column of the matrix in Figure 1 are in solid (blue) rectangles. Observe that the positions of these entries do not move up when we scan the columns from left to right.

We postpone proving Observation 1.4 to later. We now show how to exploit it to develop a divide-and-conquer algorithm. Our algorithm probes all entries in the ℓ -th column, where ℓ denote the index of the middle column of M ($\ell = 3$ in Figure 1). It can then compute $r^*(\ell)$. By Observation 1.4, it suffices to recurse the problem on the sub-matrices $M_1 = M[1, \dots, r^*(\ell); 1, \dots, \ell - 1]$ and $M_2 = M[r^*(\ell), \dots; \ell + 1, \dots]$; i.e., entries of M_1 are those in rows $i \leq r^*(\ell)$ and columns $j < \ell$ and entries of M_2 are those in rows $i \geq r^*(\ell)$ and

columns $j > \ell$ (for example, the shaded (orange) areas in Figure 1). It is not hard to see that this algorithm only probes $\tilde{O}(n)$ values in M .

The above algorithm can be easily extended to the case where we do not have r : Pick a middle node of the path T as r and run the algorithm above. Then, make a recursion on all edges on the left of r and another on all edges on the right side of r . It is not hard to see that this algorithm requires $\tilde{O}(n)$ probes.

When T is a star-graph. Now we consider another extreme case, where T is a star-graph; i.e., the spanning tree T has $n - 1$ many disjoint edges of the form $e_i = (r, i)$ where r is the root of the tree. In this case, we exploit the fact that a cut that contains only one edge of T is a candidate for the 2-respecting min-cut too. Let $\deg(u)$ denote the degree of node u and $C(u, v)$ denote the weight of edge (u, v) .

Observation 1.5. *Assume that no 2-respecting min-cut contains only one tree-edge. Let $(e_i, e_j) = \arg \min_{(e_i, e_j)} \text{cut}(e_i, e_j)$; i.e. the 2-respecting min-cuts have weight $\text{cut}(e_i, e_j)$. Then, $d(i) < 2C(i, j)$ and $d(j) < 2C(i, j)$.*

Proof. Note that $\text{cut}(e_i, e_j) = \deg(i) + \deg(j) - 2C(i, j)$. Consider the cut that separates node i from other nodes. This cut has weight $\deg(i)$. Since this is not a 2-respecting min-cut (by the assumption), $\deg(i) > \text{cut}(e_i, e_j) = \deg(i) + \deg(j) - 2C(i, j)$; thus, $d(j) < 2C(i, j)$. Similarly, $d(i) < 2C(i, j)$. \square

Observation 1.5 motivates the following notion. We say that a tree-edge e_i is *interested* in a tree-edge e_j if $\deg(i) < 2C(i, j)$. Note that each tree-edge e_i can be interested in at most one tree-edge: if e_i is interested in two distinct edges e_j and $e_{j'}$, then the degree of e_i is at least $C(i, j) + C(i, j') > (\deg(i)/2) + (\deg(i)/2)$, a contradiction. It is not hard to efficiently find the tree-edge that e_i is interested in, for every tree-edge e_i . (We only need to consider the maximum-weight edge incident to each node.) Our algorithm is now the following. For each edge e_i , let e_j be the edge that it is interested in. Compute $\text{cut}(e_i, e_j)$. Output the minimum value among (I) the values of all the computed 2-respecting cuts and (II) the minimum degree. By Observation 1.5, if (II) does not give the 2-respecting min-cut value, then (I) does. See Figure 2 for an example.

Handling general tree T . Since it is easy to find a 2-respecting min-cut if it contains only one edge in T , we assume from now that all 2-respecting min-cut contains exactly two edges in T . Let $(e^*, f^*) = \arg \min_{(e, f)} \text{cut}(e, f)$; i.e. e^* and f^* are tree-edges such that the 2-respecting min-cuts have weight $\text{cut}(e^*, f^*)$. For simplicity, let us assume further that e^* and f^* are *orthogonal* in that e^* does not lie in the unique path in T between f^* and the root of T , and vice versa.

The ideas behind Observation 1.5 can be naturally extended to a general tree T as follows. For any node u , let u^\downarrow denote the set of nodes in the sub-tree of T rooted at u . Let $\deg(u^\downarrow)$ be the total weight of edges between u^\downarrow and $V(G) \setminus u^\downarrow$. For any nodes u and v , let $C(u^\downarrow, v^\downarrow)$ be the total weight of edges between u^\downarrow and v^\downarrow . Consider any tree-edges (x, y) and (x', y') , where x (and x') is the parent of y (and y' respectively) in T . We say that (x, y) is *interested* in (x', y') if

$$\deg(y^\downarrow) < 2C(y^\downarrow, y'^\downarrow). \quad (2)$$




Figure 2: An example where T is a star-graph. Dashed edges (in blue) are tree edges. Numbers on edges indicate their weights. (All but bold edges have weight 1.) The dashed curve (in red) indicates the min-cut of the graph. This min-cut contains two tree-edges, namely e_3 and e_4 . Observe that the following pairs of tree-edges are interested in each other (e_1, e_2) , (e_3, e_4) and (e_5, e_6) . Our algorithm outputs the minimum among $\min_i(\deg(i))$, $\text{cut}(e_1, e_2)$, $\text{cut}(e_3, e_4)$ and $\text{cut}(e_5, e_6)$.

It is not hard to show that the tree-edges e^* and f^* defined above are interested in each other. Thus, like in the case of star-graphs, we can find the 2-respecting min-cut by computing $\text{cut}(e, f)$ for all pairs of e and f that are interested in each other. However, this does not imply that we need to compute only $\tilde{O}(n)$ many values of $\text{cut}(e, f)$, since an edge can be interested in *many* other edges. An additional helpful property is this:

Observation 1.6. *For any tree-edge $e = (x, y)$, edges that e is interested in form a path in T between the root and some node v .*

Proof. Throughout the proof, let $e' = (x', y')$ and $e'' = (x'', y'')$ be tree-edges where x' (and x'') is the parent of y' (and y'' respectively) in T . If e' lies in the path between the root and e'' and e is interested in e'' , then e is also interested in any e' because $C(y^\downarrow, y'^\downarrow) \geq C(y^\downarrow, y''^\downarrow)$. Now we show that if e' and e'' are orthogonal (i.e. $y' \notin y''^\downarrow$ and vice versa), then e cannot be interested in both of them. Since e' and e'' are orthogonal, the set of edges between y^\downarrow and y'^\downarrow is disjoint from the set of edges between y^\downarrow and y''^\downarrow . So, $\deg(y^\downarrow) \geq C(y, y'^\downarrow) + C(y, y''^\downarrow)$. Now, if e is interested in both e' and e'' , we have $C(y, y'^\downarrow) + C(y, y''^\downarrow) > \deg(y^\downarrow)$, a contradiction. \square

Our last idea is to exploit the above property by decomposing a tree into paths where each root-to-leaf path contains $O(\log n)$ decomposed paths. A tree decomposition of this type was used in the sequential setting in Karger’s algorithm. Later Lovett and Sandlund used the well-known heavy-light decomposition to simplify Karger’s algorithm. (The heavy-light decomposition was also internally used in some data structures used by Karger.) Both decompositions work equally well for our algorithms. To the best of our knowledge, decompositions of this type have not been used before in query complexity and streaming algorithms.

Roughly, the heavy-light decomposition partitions the tree-edges into a family \mathcal{P} of paths, such that every root-to-leaf path P in the tree T shares edges with at most $O(\log n)$ many paths in \mathcal{P} . Thus, Observation 1.6 implies that a tree-edge e is interested in edges lying in $O(\log n)$ many paths in \mathcal{P} . This motivates the following algorithm: We say that a tree-edge e is *interested in a path P* in \mathcal{P} if it is interested in some edge in P . For any two distinct paths

P and Q in \mathcal{P} , do the following. Imagine that we contract all-tree edges except those in P that are interested in Q and those in Q that are interested in P . Suppose that we are left with a path P' . We execute the path algorithm on P' . (Note that in reality we might be left with the case where every node has degree at most two except one node whose degree is three. This can be handled similarly. Also, when we implement this algorithm in different settings we do not actually have to contract edges. We only have to simulate the path algorithm on the edges in P' .)

Recall that when we run the path algorithm on each path P' above we need to compute $\text{cut}(e, f)$ for $\tilde{O}(|E(P')|)$ many pairs of tree-edges (e, f) , where $|E(P')|$ is the number of edges in P' . Moreover, each tree-edge participates in $O(\log n)$ such paths since it is interested in $O(\log n)$ many paths in \mathcal{P} . Thus, the total number of cut values $\text{cut}(e, f)$ that we need to compute for the above algorithm is $\tilde{O}(n)$ in total. To see why the algorithm finds a 2-respecting min-cut, consider when P and Q contains e^* and f^* respectively. After the contractions e^* and f^* remain in P' and thus the path algorithm executed on P' finds $\text{cut}(e^*, f^*)$.

Finding paths that an edge is interested in. Most steps described above can be implemented quite easily in all settings. The step that is sometimes tricky is finding paths in the heavy-light decomposition \mathcal{P} that a tree edge e is interested in. In the cut-query and streaming settings, we can compute this from a *cut sparsifier* which can be computed efficiently (e.g. [AGM12b, AGM12a, GKP12, RSW18]). Since a cut sparsifier preserves all cuts approximately and since we are looking for (y, y') such that $C(y^\downarrow, y'^\downarrow)$ is large compared to $\deg(y^\downarrow)$ (as in Equation (2)), we can use the cut sparsifier to identify (y, y') that “might” satisfy Equation (2). This means that for each edge e , we can identify a set of “potential” edges such that some of these edges will actually interest e . These edges might not form a path as in Observation 1.6, but we can show that they form only $O(1)$ paths.

For our sequential algorithm, let us assume for simplicity that the input graph is unweighted (possibly with parallel edges). For any node y we find y' that *might* satisfy Equation (2) by sampling $\Theta(\log n)$ edges among edges between y^\downarrow and $V(G) \setminus y^\downarrow$. For y' that satisfies Equation (2) there is a sampled edge between y^\downarrow and y'^\downarrow with high probability because $C(y^\downarrow, y'^\downarrow) > \deg(y^\downarrow)/2$. Thus, it suffices to check, for every $O(\log n)$ sampled edge e' and every path P in \mathcal{P} that overlaps with the tree-path from the root to e , whether e is interested in P or not. (Checking this requires computing $\text{cut}(u, v)$ for $O(1)$ many pairs of nodes (a, b) .) To sample the edges, we use the *range sampling* data structure in a way similar to how we use the range counting data structure to compute the cut size as outlined above. We build a range sampling data structure using the *range reporting* data structure by Overmars and Chazelle [Ove88, Cha88]. We need $O(m \log m)$ pre-processing time and $O(\log n)$ amortized time to report or sample each point (edge). (With bit operations, the preprocessing time becomes $O(m\sqrt{\log m})$ using [BGKS15, CP10, MNV16].)

1.2 Organization

First, we provide the preliminaries required in Section 2. In Section 3, we provide a schematic algorithm for the minimum 2-respecting cut problem. We first provide the schematic algorithm when the spanning tree is a path in Section 3.1. Later, we extend it the general case in Section 3.3 with the help of notions developed in Section 3.2. We combine this algorithm

with Karger’s greedy tree packing algorithm to give a weighted min-cut algorithm schematic in Section 4. In Section 5, we provide the implementation of this algorithm in two models: In Section 5.1, we provide a graph cut-query implementation (and prove Theorem 1.2), and in Section 5.2 we provide a semi-streaming implementation (and prove Theorem 1.3). Finally, in Section 6, we detail the sequential implementation of this algorithm and prove Theorem 1.1.

2 Preliminaries

Notation. We denote a spanning tree of a graph G by T . We also generally assume that T is rooted, except when T is a simple path. For a rooted tree T , we denote the subtree rooted at vertex u as u^\downarrow . If the edge e is the parent edge of u , *i.e.*, between two vertices of e , u is the farthest from the root node, then we sometimes denote the subtree rooted as u as e^\downarrow . If two vertices u and v do not belong to the same root-to-leaf path of T , we denote them as $u \perp v$. Similarly, if two edges e_1 and e_2 of T does not belong to the same root-to-leaf path, we denote them as $e_1 \perp e_2$. In a graph G , for two disjoint sets of vertices S and T , $C_G(S, T)$ denoted the total weight of the edges each of whose one edge point belongs to S and the other endpoint belongs to T . By $\mathcal{C}_G(S, T)$, we denote the set of these edges, *i.e.*, the set of edges going across from S to T . When the graph G is clear from the context, we drop the subscript G and denote is as $C(S, T)$ and $\mathcal{C}(S, T)$ respectively. We use $\deg(S)$ to denote the total weight of the edges whose only one end-point belongs to S . To denote the set of such edges, we use the notation $\mathcal{C}(S)$.

Reservoir sampling [Vit85]. We will look at a special randomized sampling technique that will be used in this work, named the *reservoir sampling*. This sampling technique aims to answer the following question:

Question 2.1. *Suppose we see a(n infinite) sequence of items $\{a_1, \dots\}$ and we want to keep ℓ many items in the memory with the following invariant: At any point i of the sequence, the ℓ items in the memory are sampled uniformly at random from the set $\{a_1, \dots, a_i\}$. What is a sampling technique that achieves this?*

The answer to this question is to use the following sampling method.

- Store the first ℓ items in memory.
- From ℓ -th time period on wards for every $i > \ell$:
 - Select the i -th item with probability ℓ/i , and replace a random item from the memory with this new item.
 - Reject the i -th item with probability $(1 - \ell/i)$.

This makes sure that every item in memory is chosen with probability ℓ/i .

3 A schematic algorithm for 2-respecting min-cut

In what follows, we first provide a schematic algorithm for finding a 2-respecting weighted min-cut of a graph G which is oblivious to the model of implementation. We then proceed

to complete the algorithm for finding min-cut on a weighted graph with an application of cut-sparsification and greedy tree packing. In the subsequent sections, we discuss the complexity of this algorithm when we implement it in different models of computation. Of course, those models need to compute a cut-sparsifier of a given weighted graph efficiently as well—we will show that this is indeed true. We start with a restricted case, where the underlying tree T is a path, and devise an algorithm which handles such graphs. Subsequently we show how we can use this algorithm as a subroutine to achieve an algorithm where there is no assumption of the structure of the underlying spanning tree T .

3.1 When spanning tree is a path

We look at a slightly different formulation of the problem of finding a 2-respecting weighted min-cut on a graph G where the underlying spanning tree T is a path. We denote it as the *Interval problem*, which is defined below.

Interval problem. Consider $n - 1$ points $\{1, \dots, (n - 1)\}$, and order them on a line from left to right. An interval $I = (s, t)$ in \mathcal{I} (where $s \leq t$) is said to cover a point i if $s \leq i \leq t$ in the path ordering. Given the set of intervals \mathcal{I} , the cost of a pair of point (i, j) , denoted as $\text{Cost}(i, j)$, is the number of intervals in \mathcal{I} which covers either i or j , but not both. The goal is to find a pair of point (i, j) such that $\text{Cost}(i, j)$ is minimized.

Claim 3.1. *The Interval problem is equivalent to the problem of finding a 2-respecting min-cut of G where the underlying spanning tree T is a path of length n .*




Figure 3: Equivalence between Interval problem and 2-respecting min-cut when the underlying tree is a path of length $n = 6$. There are 5 points in the Interval problem, and 5 corresponding edges in the path. The green edges on the RHS corresponds to intervals in the LHS. We are interested in $\text{Cost}(2, 5)$ (marked in red).

We defer the proof of this claim to Appendix A.

Cost matrix. The cost matrix M_G of G with respect to T is defined as a matrix of dimension $(n - 1) \times (n - 1)$ where the (i, j) -th entry of M_G is the weight of a 2-respecting cut of G which respects the i -th and the j -th edge of T . For the Interval problem, $M_G(i, j) = \text{Cost}(i, j)$. To reiterate, the goal of the Interval problem is to find the smallest entry in the cost matrix M_G .

To this end, we formulate a restricted version of the Interval problem, denoted as *Bipartite interval problem*, as follows:




Figure 4: Bipartite Interval problem and corresponding cost-matrix. The blue boxes show the positions of the minimum entries in each column.

Bipartite Interval problem. Divide $[n - 1]$ in two equal parts $L = \{1, \dots, \lceil \frac{n-1}{2} \rceil\}$ and $R = \{\lceil \frac{n-1}{2} \rceil + 1, \dots, n - 1\}$. We are interested in finding the pair $(i, j) \in L \times R$ such that $\text{Cost}(u, v)$ is minimum among such pairs. For solving this problem, we provide a schematic algorithm as follows. The model specific implementations are discussed in subsequent sections. In the algorithm, we assume a slightly unusual ordering on L for simplicity of exposition: Restrict the cost matrix M to $L \times R$ (Denoted by $M_{L \times R}$). The first row corresponds to $\lceil \frac{n-1}{2} \rceil$ -th point in L , the second row is $(\lceil \frac{n-1}{2} \rceil - 1)$ -th point, and so on till the 1 -th row which is the first point of L . The columns are ordered normally in an increasing order from the set $\{\lceil \frac{n-1}{2} \rceil + 1, \dots, n - 1\}$. The algorithm is recursive: it is instructive to view the execution of the algorithm as a binary tree of depth $O(\log n)$ where (i) each nodes denotes a recursive call, and (ii) a parent node makes two calls corresponding to two of its children on disjoint sub-problems (or submatrices). In any node v , the algorithm reads the middle column (**mid**) of the associated sub-matrix $M_{L_v \times R_v}$. The minimum value of this column can occur at many rows in this column. Let us denote by i_s^v to be the row with the smallest row index where the minimum occurs, and i_t^v to be the row with the largest row index where the minimum occurs. The node v then issues two recursive calls: The left child u is issued with the submatrix $M_{L_u \times R_u}$ and the right child u' is issued with the submatrix $M_{L_{u'} \times R_{u'}}$. Here L_u is a prefix of L_v of size i_s^v (i.e., the initial rows of L_v including the row i_s^v) and $L_{u'}$ is the suffix of L_v of size $|L_v| - i_t^v + 1$ (i.e., the last few rows of L_v including the row i_t^v). Also, R_u is a prefix of R_v of size $\text{mid} - 1$ (i.e., the initial columns of R_v not including the column **mid**) and $R_{u'}$ is the suffix of R_v of size $\text{mid} - 1$ (i.e., the last few columns of R_v not including the column **mid**). At the leaf nodes, when either L_v (or R_v) is a singleton set, the algorithm read the whole row (or column), and outputs the minimum. Note that, in each depth of the recursion tree, the number of entries of the matrix $M_{L \times R}$ is read is at most $n + 1$ —this will be important in the analysis of the complexity of the algorithm.

Algorithm 3.2 Schematic algorithm for Bipartite interval problem

```
1: procedure BIPARTITE-INTERVAL( $L, R$ )
2:   if  $|L| > 1$  and  $|R| > 1$  then
3:     Let  $\text{mid}$  be the middle column of  $R$ .
4:     Find the  $i_s, i_t \in L$  such that  $\text{Cost}(i_s, \text{mid}) = \text{Cost}(i_t, \text{mid})$  is
       minimum in column  $\text{mid}$ , and  $i_s$  is the first and  $i_t$  is the last such row.
        $\triangleright$  Read the  $\text{mid}$  column of  $M_{L \times R}$ 
5:     Let  $L^1$  be the prefix of  $L$  of length  $i_s$ , and  $L^2$  be the suffix of  $L$ 
       of length  $i_t - 1$ .
6:     Let  $R^1$  be the prefix of  $R$  of length  $\text{mid} - 1$ , and  $R^2$  be the suffix
       of  $R$  of length  $\text{mid} - 1$ .
7:     Run BIPARTITE-INTERVAL( $L^1, R^1$ ) and BIPARTITE-
       INTERVAL( $L^2, R^2$ ).
8:   else
9:     Read the entire  $M_{L \times R}$  and record the minimum.
        $\triangleright$  Requires reading at most 1 column or 1 row of matrix  $M$ 
10:  end if
11:  Read all recorded minimums and output the smallest.
12: end procedure
```

Before providing correctness of Algorithm 3.2, we show how, given Algorithm 3.2, we can solve the Interval problem.

Algorithm 3.3 Schematic algorithm for Interval problem

```
1: Set  $L = \{1, \dots, \lceil \frac{n-1}{2} \rceil\}$ ,  $R = \{\lceil \frac{n-1}{2} \rceil, \dots, n-1\}$ .
2: procedure INTERVAL( $L, R$ )
3:   if  $|L| > 1$  and  $|R| > 1$  then
4:     Run BIPARTITE-INTERVAL( $L, R$ ) and record the output.
        $\triangleright$  Requires reading matrix  $M$ 
5:     Divide  $L = L^1 \cup L^2$  in two equal parts where  $L^1$  is the prefix
       of  $L$  and  $L^2$  is the suffix of  $L$ , both of length  $|L|/2$ . Similarly, divide
        $R = R^1 \cup R^2$ .
6:     Run INTERVAL( $L^1, R^1$ ) and INTERVAL( $L^2, R^2$ ).
7:   else
8:     Read the entire  $M_{L \times R}$  and record the minimum.
        $\triangleright$  Requires reading constant many entries of  $M$ 
9:   end if
10:  Read all recorded outputs and output the smallest.
11: end procedure
```

Correctness of Algorithm 3.3. This follows from the observation that, for any $(i, j) \in [n-1] \times [n-1]$, there is a call to the INTERVAL procedure where $i \in L$ and $j \in R$. At the ℓ -th level of the recursion tree, the division of $L = L^1 \cup L^2$ and $R = R^1 \cup R^2$ happens depending on the ℓ -th most significant bit in the binary representation of i and j . This means

that i and j will be separated at the ℓ -th level of recursion (*i.e.*, depth ℓ in the associated recursion tree) if the most significant bit where i and j differs is the ℓ -th bit.

At this point, we make a general claim about the complexity of Algorithm 3.3. This claim will be used later to compute the complexity of Algorithm 3.3 in different models of computation.

Claim 3.4. *Let the complexity of finding i_s and i_t in Line 4 of Algorithm 3.2 on n points be $\tilde{O}(n)$ in some measure of complexity. Then the complexity of Procedure INTERVAL is $\tilde{O}(n)$.*

Proof. Let us denote the complexity of Procedure BIPARTITE-INTERVAL on input L and R where $|L| = \ell, |R| = r$ be $T(\ell, r)$. Then, we can write the following recursion:

$$T(n, n) = \tilde{O}(n) + T(n_0, n/2) + T(n_1, n/2),$$

where $n_0 + n_1 \leq n + 1$ and $T(1, r) = O(r)$ and $T(\ell, 1) = O(\ell)$. At the beginning, $n_0 = i_s$ and $n_1 = n - i_t + 1$. Solving this recursion, we get $T(n, n) = \tilde{O}(n)$.

Now let us denote the complexity of INTERVAL on n points by $C(n)$. Then we have

$$C(n) = T(n/2, n/2) + 2C(n/2),$$

with $C(1) = O(1)$. Solving this recursion and putting the value of $T(n/2, n/2)$, we get $C(n) = \tilde{O}(n)$. \square

3.1.1 Correctness of Algorithm 3.2

First, we note an important property of the cost-matrix M of the Interval problem on $n - 1$ points. Given M , we define $n - 2$ vectors $\Delta_1, \dots, \Delta_{n-2}$, each of dimension $n - 1$ as follows: $\Delta_j(i) = M(i, j) - M(i, j + 1)$.

Claim 3.5. *For each j , the vector Δ_j is monotonically increasing.*

Given Claim 3.5, we can observe the following property of the cost matrix.⁵ Suppose we are interested in finding the minimum entry in each column.

Claim 3.6. *Consider column j and let i_s be the first row and $i_t (\geq i_s)$ be the last row where minimum occurs in the column j . Then, for any column $j' < j$, a minimum entry of column j' will occur at rows in the set $\{1, \dots, i_s\}$, and for any column $j'' > j$, a minimum of column j'' will occur at rows in the set $\{i_t, \dots, n\}$.*

Given Claim 3.6, it is clear that Algorithm 3.2 records minimum entry of each column of the cost matrix. In the end, the algorithm outputs the smallest among these minimum entries, which is the minimum entry of the whole cost-matrix. Next we prove Claim 3.5 and 3.6.

⁵In literature, such a matrix is referred to as a *Monge matrix* [BKR96]. A few algorithms are known for finding column minima of a Monge matrix, most notably the SMAWK algorithm [AKM⁺87]. We use a divide and conquer approach because it gives better handle for adapting the schematic algorithm in different computational models.

Proof of Claim 3.5. We denote any interval which goes from L to R as *crossing interval*. We also denote any interval contained in L as *L-interval* and any interval contained in R as *R-interval*. We start with an empty cost-matrix (every entry is 0) and will introduce each interval one by one. Fix any j . To start with, the Δ_j is an all-0 vector and hence monotonically increasing. We show that Δ_j will maintain this property when we introduce any of the three kinds of intervals. Figure 5 shows the contribution of different types of intervals in the cost-matrix.




Figure 5: Contribution of each type of Interval (L , R , Crossing) in the cost matrix M .

L-interval. This type of interval does not increase any entries of Δ_j as they increase both $M(i, j)$ and $M(i, j + 1)$ by the same amount.

R-interval. If neither of j and $j + 1$ is covered by the interval, or both of them are covered by the interval, then Δ_j does not change. Else, assume that j is covered by the interval. This means $M(i, j)$ is increased by the same value for all i . So, if Δ_j was monotonically increasing before introducing this interval, then Δ_j remains monotonically increasing. The case where $j + 1$ is covered by the interval can be handled similarly.

Crossing interval. Let us assume that the interval starts from $p \in L$ and end on $q \in R$. If $q \neq j + 1$, then this can be handled similar to L -interval on R -interval. When $q = j + 1$, then $M(i, j + 1)$ increases for all $i \leq p$, $M(i, j)$ increases for all $i \geq p$, and all these increments are by the same amount (i.e., by the weight of the interval). This operation also does not violate the necessary property of Δ_j . \square

Proof of Claim 3.6. First we show that, for any column $j' < j$, a minimum entry of column j' will occur at rows up to i_s . The other case will follow by a similar argument. Let us denote $\Delta_{j',j}$ to be a vector where $\Delta_{j',j} = M(i, j') - M(i, j)$. It is not hard to see that $\Delta_{j',j}(i) = \Delta_{j'} + \dots + \Delta_{j-1}$, and hence is also monotonically increasing (Claim 3.5).

Now, for the sake of contradiction, assume that there is no minimum entry in the column j' in any row in the set $\{1, \dots, i_s\}$. Let the minimum entry at column j' occurs at row $i > i_s$, i.e., $M(i, j') < M(i_s, j')$. We have $M(i, j') = M(i, j) + \Delta_{j',j}(i) \geq M(i_s, j) + \Delta_{j',j}(i) \geq$

$M(i_s, j) + \Delta_{j',j}(i_s)$. Here the first inequality follows from the fact that, for column j , $M(i_s, j)$ is a minimum entry; and the second inequality follows from the fact that $\Delta_{j',j}$ is monotonically increasing. This gives us the necessary contradiction. \square

3.2 Interesting edges and paths

This section deals with the notion of *interesting* edges which we will need later to design an algorithm for the general case. From this point onward, we will make connections between edges of a spanning tree T and vertices of T . We will follow the labeling below unless stated otherwise: For vertices u, v, v' , we denote the edges of T which are parents of these vertices as e, e', e'' respectively.

Definition 3.7 (Cross-interesting edge). *Given an edge $e \in T$, we denote an edge $e' \in T$ ($e \perp e'$) to be cross-interesting with respect to e if $C(u^\downarrow, v^\downarrow) > \frac{\deg(u^\downarrow)}{2}$ where e is the parent edge of u and e' is the parent edge of v .*

Definition 3.8 (Down-interesting edge). *Given an edge e , we denote an edge $e' \in u^\downarrow$ to be down-interesting with respect to e if $C(v^\downarrow, V - u^\downarrow) > \frac{\deg(u^\downarrow)}{2}$ where e is the parent edge of u and e' is the parent edge of v .*

If two edges e and e' are cross-interesting to each other, or if e is down-interesting to e' , then we denote the pair (e, e') as *candidate exact-2-respecting cut*. The reason for denoting so will be clear in the analysis of correctness of Algorithm 3.13. Next we make the following observation about cross and down-interesting edges which we justify subsequently.

Observation 3.9. *Given an edge e , there cannot be two edges e' and e'' such that $e' \perp e''$ and both e' and e'' are cross-interesting or down-interesting with respect to e .*

To recall, e is the parent edge of u , e' is the parent edge of v and e'' is the parent edge of v' .

- **Cross-interesting** . $C(u^\downarrow, v^\downarrow)$ and $C(u^\downarrow, v'^\downarrow)$ are disjoint for $e' \perp e''$, and hence $\deg(u^\downarrow) \leq C(u^\downarrow, v^\downarrow) + C(u^\downarrow, v'^\downarrow)$. If both v and v' are interesting with respect to u , then $C(u^\downarrow, v^\downarrow) + C(u^\downarrow, v'^\downarrow) > \deg(u^\downarrow)$ which is a contradiction.
- **Down-interesting** . $C(v^\downarrow, V - u^\downarrow)$ and $C(v'^\downarrow, V - u^\downarrow)$ are disjoint for $e' \perp e''$ as well, and hence a similar argument goes through.

This means that the set of edges which are cross-interesting w.r.t. e belongs to a root-to-leaf path, and the set of edges which are down-interesting w.r.t. e belongs to a u -to-leaf path.

Note that if v'' is the parent of v , $C(u^\downarrow, v''^\downarrow) \geq C(u^\downarrow, v^\downarrow)$ and $C(v''^\downarrow, V - u^\downarrow) \geq C(v^\downarrow, V - u^\downarrow)$. So we can make the following observation:

Observation 3.10. *Given an edge e , if an edge e' is cross-interesting w.r.t e , then all ancestors of e' are interesting w.r.t e . Similarly, if an edge e'' is down-interesting w.r.t. e , then all ancestors of e'' (up to but not including e) are down-interesting w.r.t e .*

Because of Observation 3.10, we extend Definition 3.8 and 3.7 to the following:

Definition 3.11 (Interesting path). *Given an edge e , we denote a path $p \in T$ to be cross-interesting (or down-interesting) with respect to e if there is an edge $e' \in p$ such that e' is cross-interesting (or down-interesting) with respect to e . Given two paths p_1 and p_2 , we denote p_1 to be cross-interesting (or down-interesting) w.r.t. p_2 if there is an edge e in p_2 such that p_1 is cross-interesting (or down-interesting) w.r.t. e .*

3.3 Handling general spanning tree

We use the heavy-light decomposition from [ST83]. Given any rooted tree T , the heavy-light decomposition splits T into a set \mathcal{P} of edge-disjoint paths such that any root-to-leaf path in T can be expressed as a concatenation of at most $\log n$ sub-paths of paths in \mathcal{P} .

Theorem 3.12 ([ST83]). *For any vertex v , the number of paths in \mathcal{P} that starts from v or any ancestor of v is at most $\log n$. In other words, the number of paths from \mathcal{P} which edge-intersects any root-to-leaf path is at most $\log n$.*

Next, we describe the algorithm for finding a 2-respecting min-cut; the pseudo-code is provided subsequently. We assume that, in all models of computations that we are interested in, the spanning tree T is known (*i.e.*, stored in local memory).

The naïve algorithm will go over all possible 2-respecting cuts to find out the smallest among them—we want to minimize such exhaustive search. Note that, we assume, it is rather efficient to check the minimum among 1-respecting cuts. But the number of 2-respecting cuts is $\Omega(n^2)$ and hence we cannot afford to go over all possible 2-respecting cuts to find the minimum if we want to be efficient. To this end, we will examine only those 2-respecting cuts which *have the potential to be the smallest*. We describe the process of finding such potential 2-respecting cuts next. The algorithm consists of 5 steps.

Step 1. Finding 1-respecting min-cut (Line 1 to 3): The algorithm iterates over every edge e of the spanning tree T , and considers the cut which respects e . The algorithm records the value of the smallest such cut.

Step 2. Heavy-light path decomposition (Line 4): The algorithm uses heavy-light decomposition on T (*viz.* Theorem 3.12) to obtain a set of edge-disjoint paths \mathcal{P} . Each root-to-leaf path of T can be split into at most $\log n$ many subpaths from \mathcal{P} .

Step 3 (Line 5 to 8): In this step, we consider 2-respecting cuts which respects 2 edges of a path $p \in \mathcal{P}$. For each path p , we *collapse* all other edges of T to the vertices of p : The *collapse* operation contracts all other edges of T to form super-vertices, and all edges of G which are incident on any vertex which takes part in a super-vertex will now be considered incident on the super-vertex. This operation does not change the 2-respecting cut value that we are interested in for the following trivial reason: A 2-respecting cut which respects the edges of the path p will not cut any other edge of T . Next we run Algorithm 3.3 on this collapsed (*residual*) graph to find out the minimum cut which 2-respects p .

Step 4. Finding *interesting* edge-pairs (Line 9 to 12): In this step, we find *potential* edge-pairs—a pair of edges which, if respected by a cut, will yield a smaller cut value than the cuts which respects only one edge from the pair. In Step 3, the algorithm may

already have taken care a few of them, especially those pairs of edges which fall on the same path in \mathcal{P} . We are now interested in pairs which fall on different paths of \mathcal{P} . We have already introduced necessary notations: for a candidate pair of edges (e, e') , we call e is interesting w.r.t to e' and vice versa. To enumerate the set of candidate pairs, the algorithm iterates over each edge e of T , and finds a set of other edges which can form a candidate pair with e (interesting w.r.t e). By Observation 3.10, all such edges fall on a single root-to-leaf path. We check only the top (closest to the root) edges of \mathcal{P} (also because of Observation 3.10—if there is an edge in a path which is interesting w.r.t e , then the edge in the path which is closest to the root is also interesting w.r.t. e); and if any of those edges is interesting w.r.t e , we declare the path to be interesting.⁶ By Theorem 3.12, there can be at most $\log n$ many paths in \mathcal{P} which are interesting w.r.t e . We also label e with this interesting set of paths. At the end of this step, we have the following: Consider any pair of paths, p_1 and p_2 , in \mathcal{P} . For each edge $e \in p_1$, for which there is an interesting edge in p_2 , e is labelled with p_2 . Similarly, for each edge $e' \in p_2$, for which there is an interesting edge in p_1 , e' is labelled with p_1 .

Step 5. Pairing (Line 13 to 26): In this step, we pair up paths from \mathcal{P} and look at exact-2-respecting candidate cuts which respects one edge from each path of the pair. Consider a pair (p_1, p_2) : If there is no edge in p_1 which is marked by p_2 , or vice versa, or both, we discard this pair. Otherwise, we mark the edges in p_1 which are interested in p_2 , and mark the edges in p_2 which are interested in p_1 as well. We *collapse* (as in Step 3) every other edge of T so that we are left with a residual graph with edge going across only the marked edges of p_1 and p_2 —this does not change the cut value as we are interested in 2-respecting cuts which implies that no other edge of T takes part in the cut. We run Algorithm 3.3 on the residual graph and record the smallest 2-respecting cut which respects one edge in p_1 and another edge in p_2 .

At the end, we compare the recorded cuts from Step 1, 3 and 5, and output the minimum among them.

⁶In all models of implementation in this work, for every edge e , we will have a few root-to-leaf paths from the previous steps with the guarantee that the actual root-to-leaf which contains all interesting edges w.r.t. e is one of those. At this step, we will verify these paths to figure out which one (if any) is the interesting one. See Lemma 5.5 for example.

Algorithm 3.13 Schematic algorithm for 2-respecting min-cut

```
1: for every edge  $e \in T$  do                                ▷ Finding 1-respecting min-cut
2:   Find out the value of the cut which respects only  $e$  and record it.
3: end for
4: Use heavy-light decomposition on  $T$  to obtain disjoint set of paths  $\mathcal{P}$ .
                                                                      ▷ Theorem 3.12
5: for every  $p \in \mathcal{P}$  do                                  ▷ Finding cuts respecting 2
                                                                      edges in a single  $p \in \mathcal{P}$ 
6:   Collapse all edges in  $T$  which are not in  $p$  and run Algorithm 3.3
   on the residual graph.                                          ▷ Algorithm 3.3 finds mini-
                                                                      mum exact-2-respecting cut
                                                                      when  $T$  is a path
7:   Record the outputs.
8: end for
9: for each edge  $e \in T$  do                                ▷ Finding the set of interesting
                                                                      paths for each edge in  $T$ , See
                                                                      Definition 3.11
10:  Find out the set  $\mathcal{P}_e^{cross} \subseteq \mathcal{P}$  which are cross-interesting w.r.t.  $e$ .
11:  Find out the set  $\mathcal{P}_e^{down} \subseteq \mathcal{P}$  which are down-interesting w.r.t.  $e$ .
12: end for
13: for every distinct pair  $(p_1, p_2) \in \mathcal{P} \times \mathcal{P}$  do                                ▷
   Finding cuts which respects
   edges in different paths in  $\mathcal{P}$ 
14:   if  $p_1$  and  $p_2$  are cross-interested in each other then
15:     In  $p_1$ , mark edges which are cross-interested in  $p_2$ .
16:     In  $p_2$ , mark edges which are cross-interested in  $p_1$ .
17:     Collapse unmarked edges of  $T$  and run Algorithm 3.3 on the
   residual graph.                                                ▷ Algorithm 3.3 finds mini-
                                                                      mum exact-2-respecting cut
                                                                      when  $T$  is a path
18:     Record the output.
19:   end if
20:   if  $p_1$  is down-interested in  $p_2$  then
21:     In  $p_1$ , mark edges which are down-interested in  $p_2$ .
22:     In  $p_2$ , mark all edges.
23:     Collapse all unmarked edges of  $T$  and run Algorithm 3.3 on
   the residual graph.                                          ▷ Algorithm 3.3 finds mini-
                                                                      mum exact-2-respecting cut
                                                                      when  $T$  is a path
24:     Record the output.
25:   end if
26: end for
27: Output the smallest of the recorded minimums.
```

Claim 3.14. *Each edge in T takes part in at most $O(\log n)$ many calls to Algorithm 3.3 inside Algorithm 3.13.*

Proof. For an edge e , the set of edges which are cross-interesting falls on a root-to-leaf path. By Theorem 3.12, there can be at most $\log n$ many paths which are cross-interesting to e .

When $e' \in u^\downarrow$ where e is the parent of u , the set of edges which are down-interesting to e falls to a u -to-leaf path, and by Theorem 3.12 there can be at most $\log n$ many paths which are down-interesting to e . For e' , however, the argument is simpler: The set of e which are down-interested in e' are, by definition, falls on the root-to- e' path—they are all ancestors of e' in T . Hence, the number of paths which can pair up with the path containing e' in calls to Algorithm 3.3 is at most $\log n$ (by Theorem 3.12). \square

3.3.1 Correctness of Algorithm 3.13

First we make the following simple observation:

Observation 3.15. *Consider two vertices u and v such that $u \perp v$. The 2-respecting cut which respects the parent edges of u and v is a candidate for exact-2-respecting min-cut (i.e., has cut value smaller than any 1-respecting cut) if $C(u^\downarrow, v^\downarrow) > \frac{1}{2} \max\{\deg(u^\downarrow), \deg(v^\downarrow)\}$. Similarly, when $v \in u^\downarrow$, the 2-respecting cut which respects the parent edges of u and v is a candidate for 2-respecting min-cut if $C(v^\downarrow, V - u^\downarrow) > \frac{1}{2} \deg(u^\downarrow)$.*

This follows from the fact that, for the 2-respecting cut which respects the parent edges of u and v , when $u \perp v$, to be candidate for 2-respecting min-cut, it needs to happen that $\deg(u^\downarrow) + \deg(v^\downarrow) - 2C(u^\downarrow, v^\downarrow) < \min\{\deg(u^\downarrow), \deg(v^\downarrow)\}$ where the right-hand side of the strict inequality represents the value corresponding 1-respecting cuts which respects the parent edges of u and v respectively. Similarly, when $v \in u^\downarrow$, clearly, $\min\{\deg(u^\downarrow), \deg(v^\downarrow)\} = \deg(v^\downarrow)$. Hence, it needs to happen that $\deg(u^\downarrow) + \deg(v^\downarrow) - 2C(v^\downarrow, V - u^\downarrow) < \deg(v^\downarrow)$.

Claim 3.16. *All candidate exact-2-respecting cuts are considered in Algorithm 3.13.*

Proof. First, note that, any pair of edges e and e' such that (i) e is cross-interesting to e' and vice versa, or (ii) e' is down-interesting to e are considered. For any edge e , all $e' \perp e$ which are candidates are also cross-interesting to e (by Definition 3.7). Any $e' \in u^\downarrow$ (e is the parent edge of u) which is a candidate is also down-interesting to e (be Definition 3.8). \square

4 Karger's tree packing

Here we first define the notion of greedy tree packing and cut-sparsification. This exposition is based on [Tho07] to which readers are advised to refer for a more detailed description.

Definition 4.1 ((Greedy) tree packing). *A tree packing \mathcal{T} of G is a multi-set of spanning trees of G . \mathcal{T} loads each edge $e \in E(G)$ with the number of trees in \mathcal{T} that contains that edge e .*

A tree packing $\mathcal{T} = (T_1, \dots, T_k)$ is greedy if each T_i is a minimal spanning tree with respect to the loads introduced by $\{T_1, \dots, T_{i-1}\}$.

Lemma 4.2 ([Kar00]). *Let C be any cut with at most 1.1λ many edges and \mathcal{T} be a greedy tree packing with $\lambda \ln m$ many trees. Then C 2-respects at least $1/3$ fraction of trees in \mathcal{T} .*

Definition 4.3 (Cut sparsifier). *Given a graph $G = (V, E)$ with weight function $w : E \rightarrow \mathbb{R}$, a sparsifier $H = (V, E')$ of G is graph on the same set of vertices V and weight function $w' : E' \rightarrow \mathbb{R}$ with the following properties:*

1. H has $\tilde{O}(n)$ edges,
2. For every cut S , $C_H(S) \in (1 \pm \varepsilon)C_G(S)$.

Now we are ready to provide the complete algorithm for finding min-cut in a weighted graph G . There are two phases: (i) tree-packing phase, where we compute a cut-sparsifier and pack appropriately many spanning trees, and (ii) cut-finding phase, where we find a 2-respecting min-cut which respects a randomly sampled tree from the first phase. The schematic description of the algorithm is given below.

Algorithm 4.4 Schematic algorithm for weighted min-cut

- 1: Compute a cut-sparsifier H for G .
 - 2: Pack $O(\lambda \ln m)$ many spanning trees \mathcal{T} by greedy tree packing in H .
 - 3: Pick a spanning tree $T \in \mathcal{T}$ uniformly at random and run Algorithm 3.13.
-

The correctness follows from the fact that, by Definition 4.3 in H , min-cut $\lambda' \in (1 \pm \varepsilon)\lambda$. Hence, by Lemma 4.2, if we greedily pack $O(\lambda \ln m)$ many trees, the min-cut of H (and hence the min-cut of G) will 2-respect at least $1/3$ fraction of the packed trees. So, if we pick a tree uniformly at random, the sampled tree will be 2-respected by the min-cut with probability at least $1/3$.

In all three models that we describe subsequently, we will compute the cut-sparsifier H efficiently and store it locally. The greedy packing of spanning trees will be performed also locally, and then we will run an algorithm for finding a 2-respecting min-cut (aka model specific implementation of Algorithm 3.13) on a randomly sampled tree from the set of packed trees.

In order to reducing the value of the min-cut of a graph (and, thereby, reducing the number of trees that are needed to be packed by a greedy tree packing algorithm), we use the following skeleton construction due to Karger [Kar00].

Theorem 4.5 ([Kar00], Theorem 4.1). *Given a weighted graph G , we can construct a skeleton graph H of G such that*

1. H has $m' = O(n\varepsilon^{-2} \log n)$ edges,
2. The minimum cut of H is $O(\varepsilon^{-2} \log n)$,
3. The minimum cut in G corresponds (under the same vertex partition) to a $(1 + \varepsilon)$ -times minimum cut of H .

Thus, a set of $O(\log n)$ spanning trees can be packing in G (by performing greedy tree packing on H) such such that the minimum cut of G 2-respects $1/3$ of them with high probability.

Note that Theorem 4.5 is useful only when the value of min-cut is $\omega(\log n)$. Otherwise, we can greedily pack trees in G (or its sparsifier) itself to obtain a set of $O(\log n)$ spanning trees.

5 Cut-query & streaming algorithms

In this section, we provide model-specific implementation of Algorithm 4.4 in two models: In Section 5.1 we look at graph-query model, and, in Section 5.2, we provide an algorithm in the dynamic streaming model. In all these sections, the arguments follow the following general pattern.

- We first show that the randomized reduction from finding a weighted min-cut to finding a 2-respecting min-cut w.r.t to a given spanning tree can be implemented efficiently (Line 1 to 2 of Algorithm 4.4). [Theorem 5.2 and Claim 5.8]
- Next we show that we can implement Algorithm 3.3 efficiently. This algorithm will be called many times in Algorithm 3.13, and hence we want to make sure that this step is efficient. [Claim 5.4, and Claim 5.10]
- Then we show that each edge $e \in T$ can find the sets \mathcal{P}_e^{cross} and \mathcal{P}_e^{down} efficiently. This is a crucial step because this reduces the search space of 2-respecting cuts, and we want to make sure that this step can be done as quickly as possible. [Lemma 5.5, and Claim 5.11]
- Lastly, we analyse the other steps of Algorithm 3.13 to conclude that the algorithm is efficient. [Lemma 5.6, and 5.12]

5.1 Graph cut-query upper bound

In the graph cut-query model, we have oracle access to G in the following way: We can send a partition (S, \bar{S}) of the vertex set V of G to the oracle and the oracle will reply with the value of the cut corresponding to the given partition. The goal is to minimize the number of such accesses to compute weighted min-cut. Graph cut-queries can be quite powerful as the following claim suggests. Most of these are stated in [RSW18] and are fairly easy to verify.

Claim 5.1. *In $O(\log n)$ cut-queries we can*

1. *learn one neighbor of a vertex $u \in V$ or a set of vertices $U \subset V$,*
2. *sample a random neighbor of a vertex $u \in V$ or a set of vertices $U \subset V$ and*
3. *after performing n initial queries, sample a random edge $e \in E$.*

And in 3 cut-queries, we can find out the total weight of all edges going between two disjoint set of vertices.

Now, we state a result of [RSW18] regarding computing a cut-sparsifier efficiently. Note that, such a sparsifier is computed by using all of the above powers of cut-queries.

Theorem 5.2 ([RSW18]). *Fix any $\varepsilon > 0$. By using at most $\tilde{O}(n/\varepsilon^2)$ cut-queries, we can produce a sparsifier H of G such that:*

1. *H has $O(n \ln n / \varepsilon^2)$ edges,*
2. *Every cut in H is within a $(1 \pm \varepsilon)$ -factor of its value in G .*

At this point, we can perform greedy tree packing in H and can sample a random spanning tree from the packed set of trees. We are left with the job of finding a min-cut which 2-respects the sampled tree⁷. Now we turn towards implementing Algorithm 3.13 in the cut-query model. At first we show that Algorithm 3.3 is efficient. As noted in Claim 3.1, any efficient query protocol for the Interval problem immediately implies an efficient query protocol for the 2-respecting min-cut problem where the underlying spanning tree is a path. In the Interval problem, we have oracle access to the entries of the cost-matrix M : This is because, given any pair of edges e_i and e_j of the spanning tree, we can find out the value of the 2-respecting cut which respects e_i and e_j by a single cut-query—this is the value of $\text{Cost}(i, j)$ in the Interval problem. But, even if we want to check all 2-respecting cuts, we need $\Omega(n^2)$ queries which we cannot afford. The following observation follows from the discussion so far.

Observation 5.3. *Any entry of the cost-matrix M associated with the Interval problem can be known in a single cut-query.*

This immediately implies that we can execute Line 4 of Algorithm 3.2 with $O(n)$ cut-queries. So, using Claim 3.4, we can claim the following:

Claim 5.4. *The cut-query complexity of Algorithm 3.3 on a n vertex graph is $O(n)$.*

Now we come to the implementation of Algorithm 3.13. We make the following claim next.

Lemma 5.5. *We can find out $\mathcal{P}_e^{\text{cross}}$ and $\mathcal{P}_e^{\text{down}}$ for all $e \in T$ with $O(n)$ cut-queries.*

Proof. Let \mathcal{C} denote the set of cut queries we are interested in which are (i) $\mathcal{C}(u^\downarrow)$ for all $u \in V$, (ii) $\mathcal{C}(u^\downarrow, v^\downarrow)$ for all pairs of distinct u, v such that $u \perp v$, and (iii) $\mathcal{C}(v^\downarrow, V - u^\downarrow)$ for all pairs of u, v such that $v \in u^\downarrow$. We can assume that

1. we have made n many initial cut-queries to find out $\deg(u^\downarrow)$ for each $u \in V$, and
2. we also know the edges of the sparsifier H from the tree-packing phase (Line 1 to 2 of Algorithm 4.4).

Consider an edge e' which is cross-interesting to e , i.e., $\deg(u^\downarrow) < 2C(u^\downarrow, v^\downarrow)$ where e is the parent edge of u and e' is the parent edge of v . Now, by Theorem 5.2, we know that in H , $\deg_H(u^\downarrow) \leq (1 + \varepsilon) \deg_G(u^\downarrow) < 2(1 + \varepsilon)C_G(u^\downarrow, v^\downarrow)$.

To connect C_G with C_H , we have work a bit more: Let us partition the vertex set V of G in three parts: u^\downarrow , v^\downarrow and $R = V - (u^\downarrow \cup v^\downarrow)$. Now, we know that $2C_G(u^\downarrow, v^\downarrow) > \deg_G(u^\downarrow) = C_G(u^\downarrow, v^\downarrow) + C_G(u^\downarrow, R)$. Similarly, $2C_G(u^\downarrow, v^\downarrow) > C_G(u^\downarrow, v^\downarrow) + C_G(u^\downarrow, R)$. Combining, we get $2C_G(u^\downarrow, v^\downarrow) > C_G(u^\downarrow, R) + C_G(v^\downarrow, R)$. Now, let us look at the sparsifier H . We have the

⁷Note that we have find 2-respecting min-cut for $O(\log n)$ many sampled trees from the packed set of tress in order to attain very high success probability.

following:

$$\begin{aligned}
& 2C_H(u^\downarrow, v^\downarrow) \\
& \underbrace{\geq}_{(*)} (1 - \varepsilon)(C_G(u^\downarrow, v^\downarrow) + C_G(u^\downarrow, R)) + (1 - \varepsilon)(C_G(u^\downarrow, v^\downarrow) + C_G(v^\downarrow, R)) - (1 + \varepsilon)(C_G(v^\downarrow, R) + C_G(u^\downarrow, R)) \\
& = 2(1 - \varepsilon)C_G(u^\downarrow, v^\downarrow) - 2\varepsilon(C_G(u^\downarrow, R) + C_G(v^\downarrow, R)) \\
& \underbrace{>}_{(**)} 2(1 - \varepsilon)C_G(u^\downarrow, v^\downarrow) - 4\varepsilon C_G(u^\downarrow, v^\downarrow) \\
& = 2C_G(u^\downarrow, v^\downarrow)(1 - 3\varepsilon),
\end{aligned}$$

where $(*)$ follows from the sparsifier guarantee, and $(**)$ follows from what we observed before. Hence, we have

$$\deg_H(u^\downarrow) < 2(1 + \varepsilon)C_G(u^\downarrow, v^\downarrow) < 2\frac{1 + \varepsilon}{1 - 3\varepsilon}C_H(u^\downarrow, v^\downarrow).$$

For small enough ε , we have $\deg_H(u^\downarrow) < 3C_H(u^\downarrow, v^\downarrow)$.

Hence, to find out which edges are cross-interesting w.r.t. e in G , we need to check whether $C_H(u^\downarrow, v^\downarrow) > \deg_H(u^\downarrow)/3$ in H . Extending Definition 3.7, let us call the edge e' which is parent of v to be H -cross-interesting w.r.t. e if $C_H(u^\downarrow, v^\downarrow) > \deg_H(u^\downarrow)/3$. Note that the set of H -cross-interesting edges w.r.t. e is superset of the set of edges which are cross-interesting w.r.t. e . By a similar argument as that of Observation 3.9, we know that there can be at most 2 root-to-leaf paths (none of which contains e) in H which can contain edges that are H -cross-interesting to e . These two root-to-leaf paths will intersect with at most $2 \log n$ paths from \mathcal{P} , and hence we need to check $C(u^\downarrow, v^\downarrow)$ for at most $2 \log n$ many vertices $v \perp u$ to find out which paths are actually cross-interesting to e . We can make $6 \log n$ cut queries—3 queries for each path $p \in \mathcal{P}$ which intersects these two root-to-leaf paths—in the original graph to figure out which paths in \mathcal{P} are interesting w.r.t. e , i.e., the set $\mathcal{P}_e^{\text{cross}}$. A similar argument can be made for the set $\mathcal{P}_u^{\text{down}}$. \square

Now we analyze, as before, the cut-query complexity of Algorithm 3.13

Lemma 5.6. *The cut-query complexity of finding a 2-respecting weighted min-cut is $\tilde{O}(n)$.*

Proof. We count the number of cut-queries required at each line of Algorithm 3.13.

Line 1 to 3: This requires $n - 1$ many cut-queries, one for each edge e .

Line 5 to 8: By Claim 5.4, each call to Algorithm 3.3 requires size of the path many cut-queries. The paths in \mathcal{P} are disjoint, and hence each edge takes part in exactly 1 path in \mathcal{P} . Hence, the total number of queries required is $n - 1$.

Line 9 to 12: This requires the knowledge of the sparsifier graph H , which we can assume to possess because of the tree packing steps. This also requires the knowledge of $\deg(u^\downarrow)$ for every u . We already know this from Line 1 to 3. By Lemma 5.5, this step can be executed with $\tilde{O}(n)$ many queries.

Line 13 to 26: As before, we can use Claim 3.14 and 5.4 to conclude that this step requires $\tilde{O}(n)$ many cut-queries.

Hence, in total, $\tilde{O}(n)$ many cut-queries are required. \square

5.2 Streaming upper bound

In the dynamic streaming model, we assume that the algorithm knows the vertex set V of G and has access to a stream of edge-insertion and edge-deletion instructions: each such instruction declares a pair of vertices of the graph and an optional weight, and mentions whether an edge of the corresponding weight needs to be inserted between the two vertices, or if the edge which is already present between those two vertices should be removed. First we state a result of [KLM⁺17] regarding efficient computability of sparsifier.

Theorem 5.7 ([KLM⁺17]). *There exists an algorithm that processes a list of edge insertions and deletions for a weighted graph G in a single pass and maintains a set of linear sketches of this input in $\tilde{O}(n)$ space. From these sketches, it is possible to recover, with high probability, a (spectral⁸) sparsifier H with $\tilde{O}(n)$ edges.*

Claim 5.8. *A dynamic streaming algorithm can perform the randomized reduction from weighted min-cut problem to a 2-respecting weighted min-cut problem in a single pass with $\tilde{O}(n)$ amount of total memory.*

Proof. Given a weighted graph G , the algorithm first uses Theorem 5.7 to construct a sparsifier H of size $\tilde{O}(n)$ in a single pass. Next, the algorithm locally constructs a skeleton graph H' of H such that the value of the min-cut in H' becomes $O(\log n)$. To this this, the algorithm uses Theorem 4.5 by setting ε to be a very small constant less than 1 (say $\varepsilon = 1/100$). Now, the algorithm can perform greedy tree packing on H' locally. Let us denote, as before, the greedy tree packing by $\mathcal{T} = \{T_1, \dots, T_k\}$. Because of the skeleton construction, $k = O(\log n)$ and the algorithms can store the set \mathcal{T} in its internal memory.⁹ The algorithm does reservoir sampling [Vit85] to pick a random tree T from \mathcal{T} : This way the algorithm can sample a tree T uniformly at random from \mathcal{T} with $\tilde{O}(n)$ memory. By Lemma 4.2, the min-cut 2-respects T with constant probability. \square

Now we turn towards implementing Algorithm 3.13 in the dynamic streaming model. We first note the following: For a given cut, we can find the value of the cut in a dynamic stream with just $O(\log n)$ bits of memory—we maintain a counter and whenever the stream inserts (or deletes) an edge in (or from) the cut, we increase (or decrease) the counter. So the following observation is immediate.

Observation 5.9. *Any given $\tilde{O}(n)$ cut queries can be answered by a dynamic streaming algorithm in a single pass with $\tilde{O}(n)$ space.*

As before, by Claim 3.1, this means that the streaming algorithm can find $\tilde{O}(n)$ entries of the cost-matrix corresponding to the Interval problem in a single pass and $\tilde{O}(n)$ space. We prove that this is enough to implement Algorithm 3.3 efficiently.

Claim 5.10. *A dynamic streaming algorithm can perform Algorithm 3.3 on an n vertex graph in $O(\log n)$ pass with $\tilde{O}(n)$ bits of memory.*

⁸A stronger notion than that of cut-sparsifier.

⁹The skeleton construction was not necessary in the case of the cut-query model because of its assumption of unbounded internal computational power.

Proof. We first look at the pass and memory requirement for performing Algorithm 3.2. Note that we can use Claim 3.4 where we assume that a streaming algorithm can find i_s and i_t in Line 4 of Algorithm 3.2 can be done in a single pass with $\tilde{O}(n)$ bits of memory (by Observation 5.9). Hence, by a similar calculation, it is easy to see that Algorithm 3.2 can be performed in $O(\log n)$ passes and with $\tilde{O}(n)$ memory.

For Algorithm 3.3, we can run $O(\log n)$ many occurrences of previous algorithm simultaneously, one for each level of recursion: The i -th algorithm will run 2^{i-1} many instances of BIPARTITE-INTERVAL, each one on a path (or line) with $n/2^{i-1}$ edges (or points). Clearly, each algorithm can run in $O(\log n)$ passes with $\tilde{O}(n)$ memory, and hence the complexity of the combined algorithm is $O(\log n)$ passes with $\tilde{O}(n)$ memory as well. \square

Now, as in the graph cut-query model, we turn towards the implementation of Algorithm 3.13. We make the following claim.

Claim 5.11. *A dynamic streaming algorithm can find $\mathcal{P}_e^{\text{cross}}$ and $\mathcal{P}_e^{\text{down}}$ for all $e \in T$ in two passes and $\tilde{O}(n)$ memory.*

Proof. The proof is similar to Lemma 5.5. As noted in Theorem 5.7, the algorithm can compute a sparsifier H of G in a single pass with $\tilde{O}(n)$ memory. Once done, by a similar argument as that of Lemma 5.5, each edge $e \in T$ needs to check $O(\log n)$ paths from \mathcal{P} to figure out which paths are cross (down)-interesting w.r.t. e in the original graph G . This requires checking cut-values of $\tilde{O}(n)$ many cuts in total. This, by Observation 5.9, can be done in a single pass with $\tilde{O}(n)$ bits of memory. \square

Now we analyze, as before, the streaming complexity of Algorithm 3.13.

Lemma 5.12. *There exists an algorithm that processes a list of edge insertions and deletions for a weighted graph G on n vertices can perform Algorithm 3.13 in $O(\log n)$ passes and with total memory of $\tilde{O}(n)$ bits.*

Line 1 to 3: There are $n - 1$ cuts to check, and hence, by Observation 5.9, it can be done in a single pass with $\tilde{O}(n)$ bits of memory.

Line 5 to 8: We know, by Claim 5.10, that the interval problem on ℓ many points can be performed in $O(\log \ell)$ passes and $\tilde{O}(\ell)$ bits of memory. As the set \mathcal{P} is edge-disjoint, we can run $|\mathcal{P}|$ many streaming algorithm in parallel, all of which read the same dynamic stream. This needs $O(\log n)$ passes and $\tilde{O}(n)$ memory.

Line 9 to 12: By Claim 5.11, this can be done in 2 passes and with $\tilde{O}(n)$ memory. Note that, the first pass of Claim 5.11 is dedicated to finding a sparsifier H , which will be done in the tree-packing phase (Line 1) of Algorithm 4.4 (See proof of Claim 5.8). So, while implementing Algorithm 4.4, these lines can be executed in a single pass instead of two passes.

Line 13 to 26: We will run one instance of the dynamic streaming algorithm for each call to Algorithm 3.3 in parallel, all of which reads the same stream. This needs $O(\log n)$ many passes. For the memory bound, Claim 3.14 guarantees that each edge is going to be used by at most $O(\log n)$ many such algorithms. As stated before, any such algorithm solving Interval problem on ℓ many points can be performed in $O(\log \ell)$ passes and $\tilde{O}(\ell)$

bits of memory. So these steps can be performed in $O(\log n)$ passes and $\tilde{O}(n)$ total memory.

Hence, Algorithm 3.13 can be implemented by a dynamic streaming algorithm in $O(\log n)$ passes and $\tilde{O}(n)$ memory.

6 Sequential algorithm for 2-respecting min-cut

In the following section, we provide a randomized algorithm for the problem of finding minimum 2-respecting min-cut. The error probability of this algorithm is polynomially small, *i.e.*, n^{-c} for any large enough constant c —we denote this as *with high probability*. Our main result is the following.

Theorem 6.1. *There is a randomized algorithm that, given a spanning tree T of an n -node m -edge (weighted) graph G , can find the 2-respect min-cut with high probability in $O(m \log n + n \log^4 n)$ time. If the graph is unweighted (but can have multiple edges between a pair of vertices), then the 2-respecting min-cut can be found in $O(m\sqrt{\log n} + n \log^4 n)$.*

Theorem 6.1 relies on a certain 2-d range-counting and sampling data-structure that we state in the following section. For the 2-d range counting data-structure and all other necessary data-structures introduced in Section 6.2, we will also look at the time complexity averaged out over a sequence of data-structure operations, *i.e.*, we are interested in *amortized* time complexity of data-structure operations.

6.1 Range counting, sampling & searching

We start this section by introducing data-structure for 2-d orthogonal range counting/sampling. We later introduce data-structure for 2-d semigroup range searching.

Definition 6.2 (2-d Orthogonal Range Counting/Sampling). *An (online) range counting data structure is an algorithm that take m points in the plane to pre-process in time t_p and support a sequence of the following query operations:*

1. *given an axes-aligned rectangle, the algorithm outputs the number of points in the rectangle in time t_c (range counting),*
2. *given an axes-aligned rectangle and an integer $k = O(\log m)$, the algorithm outputs k' distinct random points in R ($k \leq k' = O(k)$) with high probability in time $(k' + 1)t_s$ (range sampling).*

We denote such a data-structure by (t_p, t_c, t_s) -data-structure. We will also be interested in amortized range reporting time t_r : given an axes-aligned rectangle and an integer $k = O(\log m)$, the algorithm outputs k' distinct points in R ($k \leq k' = O(k)$) in time $(k' + 1)t_r$.

Below we present two (t_p, t_c, t_s) -data-structures—one where we do not allow bit operations and one where we allow it—with their corresponding pre-processing, counting and sampling time.

Theorem 6.3. *There are (t_p, t_c, t_s) -data-structures when given m points in a 2-d plane to pre-process have the following complexities:*

<i>(Amortized) time</i>	<i>without operations</i>	<i>bit-</i>	<i>with operations</i>	<i>bit-</i>
$t_p(m)$	$O(m \log m)$		$O(m\sqrt{\log m})$	
$t_c(m)$	$O(\log m)$		$O\left(\frac{\log m}{\log \log m}\right)$	
$t_s(m)$	$O(\log m)$		$O\left(\frac{\log m}{\log \log m}\right)$	

For proving Theorem 6.3, when bit-operations are not allowed, we use the following result by [Cha88].

Theorem 6.4 ([Cha88]). *There is a 2-d range count/report data-structure which can pre-process m points in time $O(m \log m)$ and answers range count/report query in $O(\log m)$ amortized time which does not require bit operations.*

Remark 6.5. Chazelle [Cha88] proved an even stronger result where he showed that reporting k many points takes time $O(k + \log m)$. Also, as we will see in Section 6.2, it is enough for these data-structures to work on 2-d grid (instead of 2-d plane). For this special case, Overmars [Ove88] has shown the existence of a data-structure with $t_p(m) = O(m \log m)$ and reporting k points take time $O(k + \sqrt{\log m})$ (given that the grid is of size $m \times m$).

When we allow bit operations, we can do better. We use the following 1-d range rank/select query data-structure from [BGKS15]: A range rank/select query data-structure, given m numbers in an array $A[1, \dots, m]$ to pre-process, can support the following types of queries:

1. a *rank query* is associated with two indices $i, j \in [m]$ and an integer I , and outputs the number of integers in $A[i \dots, j]$ that are smaller than an integer I ,
2. a *select query* is associated with two indices $i, j \in [m]$ and a number k , and outputs the k -th smallest number in $A[i, \dots, j]$.

The following two theorems are stated for data-structure where bit operations are allowed.

Theorem 6.6 ([BGKS15]). *There is a 1-d range rank/select query data-structure which can pre-process m points in time $O(m\sqrt{\log m})$ and answers range rank/select query in time $O\left(\frac{\log m}{\log \log m}\right)$.*

Theorem 6.7 ([CP10]). *There is a 2-d orthogonal range counting data-structure which can pre-process m points in time $O(m\sqrt{\log m})$ and answers range count query in time $O\left(\frac{\log m}{\log \log m}\right)$.*

Proof of Theorem 6.3. First we show how to construct such a data-structure when bit-operations are not allowed. The main idea is to show how to construct a range-sampling data-structure from a range-reporting data-structure from Theorem 6.4. Subsequently, we show how to construct a more efficient range-reporting data-structure when bit-operations are allowed. We will show how to construct an efficient range-reporting data-structure from Theorem 6.6 and 6.7.

Without bit-operations. We show the proof by constructing the (randomized) data-structure and analysing its correctness. We create a data-structure D as in Theorem 6.4. The time required for pre-processing m points is $t_p(m) = O(m \log m)$ and the time required for range counting is $t_c(m) = O(\log m)$. The rest of the proof is to show $t_s(m) = O(\log m)$ as well.

To this end, we show that, given a 2-d range-reporting data-structure with pre-processing time $t_p(m) = O(m \log m)$ and range-reporting time t_r , we can construct a range-sampling data-structure with pre-processing time $t_p(m) = O(m \log m)$ and $t_s = t_r$. The proof then concludes by observing that $t_r(m) = O(\log m)$ for D by Theorem 6.4. The reduction is described as follows:

- **Pre-processing:** We define the following set of points S_0, S_1, \dots, S_k , where $k = \log m$ in the following way: Let $S_0 = S$. For i starting from 1 to k , we define $S_i \subseteq S_{i-1}$ by copying each point of S_{i-1} to S_i with probability $1/2$. For every $0 \leq i \leq k$, we also do the following: We build a 2-d range reporting data structure D_i for each point set S_i . This takes time $\sum_{i=0}^k t_p(|S_i|)$ where t_p is the pre-processing time of the range reporting data structure. This sums up to $\sum_{i=0}^k O(|S_i| \log |S_i|) = O(m \log m)$.
- **Query(R, k):** Let $i = k$. We query the range reporting data structure D_i for points in $R \cap S_i$ and output the reported points. If less than $c \log n$ points were reported so far, we repeat with $i \leftarrow i - 1$. We stop when all points in R are reported (when $i = 0$). D_i reports each point in $t_r(|S_i|) = O(\log |S_i|)$ amortized time. Note that even when R contains no point, we need an extra $O(\log m)$ time to iterate over all values of i . (This can be avoided using the range counting data structure, but this refinement does not affect our overall running time so we do not include it here.)

Analysis: For any $i \leq k$, let $S_{\geq i} = \bigcup_{j \geq i} S_j$. Observe that, for every i , every point is in $S_{\geq i}$ with probability $1/2^i$. Since we answer the query by reporting all points in $S_{\geq i}$ for some i , every point is output with the same probability. Now we bound the number of points we output. Let m' be the number of points in R in total. Let $c \geq 2$ be a large enough constant whose value is independent from the input size. Let i' be such that $m'/2^{i'} \in (ck, 2ck]$. The expected number of points in R that is in $S_{\geq i'}$ is $E[|R \cap S_{\geq i'}|] = m'/2^{i'} \geq ck \geq c \log m$; so by the standard Chernoff's bound we have that $|R \cap S_{\geq i'}| \in [ck/2, 4ck]$ with (very) high probability. If this is the case, the query algorithm will stop when it reports points in $R \cap S_{\geq i'}$ at the latest, implying that it reports at most $|R \cap S_{\geq i'}| = O(k)$ points. The proof follows by noting $t_r(m) = \max_i O(\log |S_i|) = O(\log m)$.

With bit-operations. Similar to what we did previously, we show the proof by constructing the (randomized) data-structure and analysing its correctness. We will create a data-structure D as in Theorem 6.6 and another data-structure D' as in Theorem 6.7. The time required for pre-processing for both D and D' is $O(m\sqrt{\log m})$. This immediately gives $t_p(m) = O(m\sqrt{\log m})$. Also, for range counting, we use D' . By Theorem 6.7, $t_c(m) = O(\frac{\log m}{\log \log m})$. Previously, we have seen that, given a range-reporting data-structure with reporting time t_r , we can design a range-sampling data-structure with sampling time $t_s = t_r$. A similar argument

also holds here as well¹⁰. Hence, the rest of the proof is to show that $t_r(m) = O\left(\frac{\log m}{\log \log m}\right)$.

Concretely, we show how to construct a range-reporting data-structure with pre-processing time $t_p(m) = O(m\sqrt{\log m})$ and range-reporting time $t_r = O\left(\frac{\log m}{\log \log m}\right)$ from D and D' . Note that D is a 1-d data-structure. Hence we need to map the points 2-d orthogonal plane to a single dimension before D can pre-process it. For this, we create an array \mathcal{A} of size m . For points $e = (e^{(x)}, e^{(y)})$ in the 2-d plane, where $e^{(x)}$ and $e^{(y)}$ are the x and y coordinates of e , we order them with respect to the x -coordinate, and for points with same x -coordinate, we order them by y -coordinate. Given this ordering, we have a one-to-one correspondence between m points in the 2-d plane and the entries of \mathcal{A} : The entry $\mathcal{A}[i]$ corresponds to a unique point e_i in the 2-d plane. We store the y -coordinate of e_i in $\mathcal{A}[i]$, i.e., $\mathcal{A}[i] = e_i^{(y)}$.

Now we demonstrate how to report one point in a given axes-aligned rectangle R . Let us denote the columns of R by $\{R_1, \dots, R_t\}$ where R_i resides in the column C_i of the 2-d plane. We also consider two more (related) axes-aligned rectangles, and for that we need to introduce the following notion of *dominance*. A point e in the 2-d plane *dominates* another point e' if $e^{(x)} \geq e'^{(x)}$ and $e^{(y)} \geq e'^{(y)}$. Extending this notion, a rectangle R_1 *dominates* a disjoint rectangle R_2 if each point in R_1 dominates all points in R_2 . The two other rectangles we consider are as follows: (i) Let \tilde{R} be the ambient rectangle formed by taking union of columns $\{C_1, \dots, C_k\}$, and (ii) \bar{R} to be the biggest sub-rectangle of \tilde{R} which is dominated by and disjoint from R . The following two observations are immediate: (i) The points in the rectangle \tilde{R} corresponds to a contiguous range in \mathcal{A} , and (ii) w.r.t. the ordering of \mathcal{A} , the entries in \mathcal{A} corresponding to the points of \bar{R} are less than that of R , i.e., $\mathcal{A}[i] < \mathcal{A}[j]$ if $e_i \in \bar{R}$ and $e_j \in R$.

1. We query D' to find the count of points in \bar{R} . Let the number of such points is c .
2. We also query D' to find the count of points in $R \cup \bar{R}$. Let this count be c' . This means that $c' - c$ many points reside in R . If $c' - c = 0$, we give up and return 0.
3. Otherwise, we query D in the range corresponding to \tilde{R} to select the smallest $(c + j)$ -th entry of \tilde{R} for all $j \leq c' - c$. Note that these points must belong inside R if $c' - c \neq 0$. Finding the range corresponding to \tilde{R} can be achieved by making 2 count-queries to D' : one to find out the number of points on the left side of \tilde{R} which gives us the starting point of the range; and the other to find out the number of points on the right side of \tilde{R} which gives us the ending point of the range.

Hence, to report k points in R , the total query time is $O\left((k + 4)\frac{\log m}{\log \log m}\right)$. This gives an amortized query time of $t_r(m) = O\left(\frac{\log m}{\log \log m}\right)$ as required. This completes the proof of Theorem 6.3. \square

Remark 6.8. Even though the above proof of Theorem 6.3 can deal with multiple points in a single coordinate of the 2-d plane, we do not need such strong guarantee in our work. In what follows, we will have at most a single point in a single coordinate of the 2-d plane.

Next we introduce semigroup range searching data-structure. This is an extension of range counting data-structure where we can assume that each point in the 2-d plane has a

¹⁰One has to be careful about pre-processing time which is now $\sum_i O(|S_i| \sqrt{\log |S_i|}) = O(m\sqrt{\log m})$ as $\sqrt{\log(x)}$ is an increasing function w.r.t. x .

weight associated with it, and given an axes-aligned rectangle we want to know that total weight of the points inside the rectangle (instead of counting just the number of points inside it).

Definition 6.9 (2-d Semigroup Range Searching). *Consider a function f from points on a 2-d plane to a commutative semi-group $(G, +)$. An (online) semi-group range searching data structure is an algorithm that take m points in the plane to pre-process in time t_p^{sg} and support a sequence of the following query operation:*

- given an axes-aligned rectangle R , the algorithm outputs $\sum_{e \in R} f(e)$ where the summation is the semigroup operation.

We denote the query time as t_c^{sg} , and call such a data-structure as $(t_p^{\text{sg}}, t_c^{\text{sg}})$ -data-structure.

We will make use of the following $(t_p^{\text{sg}}, t_c^{\text{sg}})$ -data-structure (for handling weighted graph). This data-structure does not require bit operations.

Theorem 6.10 ([Cha88]). *There is a $(t_p^{\text{sg}}, t_c^{\text{sg}})$ -data-structure in which can pre-process m points in time $t_p^{\text{sg}}(m) = O(m \log m)$ and answers range search query in $t_c^{\text{sg}}(m) = O(\log^2 m)$ amortized time.*

6.2 Main algorithm: Proof of Theorem 6.1

In this section, we will assume the existence of a (t_p, t_c, t_s) -data-structure as in Theorem 6.3 and a $(t_p^{\text{sg}}, t_c^{\text{sg}})$ -data-structure as in Theorem 6.10, and will prove Theorem 6.1. First, we show a few operations we can perform on a graph which are immediate from such data-structures.

Lemma 6.11. *For any spanning tree T of an n -node m -edge unweighted graph G , there is an algorithm that, after a $t_p(m)$ -time pre-processing, can answer the following queries in $t_c(m)$ amortized time:*

1. for any vertex v , $\deg(v^\downarrow)$,
2. for any two vertices u and v , $C(u^\downarrow, v^\downarrow)$, and
3. for any two vertices u and v , $C(u^\downarrow, V - v^\downarrow)$.

When G is weighted, then these operations need $t_p^{\text{sg}}(m)$ -time for pre-processing and $t_c^{\text{sg}}(m)$ -time to answer.

Proof. We sequentialize the vertices in V by the trace of a post-order traversal of T . Let us denote this order by \prec . We will also use this notation to define $S \prec T$ for set of vertices S and T if all vertices in S occur before any vertex of T in the ordering. A post-order traversal order has the following property: For any vertex v with children $u_1 \prec \dots \prec u_k$, the ordering guarantees that $u_1^\downarrow \prec \dots \prec u_k^\downarrow$. This immediately implies that

1. for all $u \in V$, u^\downarrow is a continuous range in the ordering, and
2. for all $v \in V$, $V - v^\downarrow$ comprises of two continuous ranges in the ordering.

So we plot V in this order along both axes, and, for all edge e in the graph G , put the weight of the edge $e = (u, v)$ at the coordinate (u, v) on the plane. (We need to account for e only once between coordinates (u, v) and (v, u) .) By formulating $\deg(v^\downarrow) = C(v^\downarrow, V - v^\downarrow)$, we see that each of the cut-queries in the claim can be answered by at most 2 range-count queries to the (t_p, t_c, t_s) -data-structure if the edges have weight 1 each, or at most 2 queries to the $(t_p^{\text{sg}}, t_c^{\text{sg}})$ -data-structure if the edges are weighted. \square

Now we look at Algorithm 3.2 and 3.3, and see how they can be implemented efficiently if we have a (t_p, t_c, t_s) -data-structure. We start with the following lemma.

Lemma 6.12. *The exists an algorithm that supports the following operations.*

- *Pre-process(T, G): The algorithm is given a spanning tree T of an n -node m -edge unweighted graph G to pre-process.*
- *Query(A, B): The algorithm is given two disjoint lists (arrays) of edges in T , denoted by $A = \{a_1, \dots, a_x\}$ and $B = \{b_1, \dots, b_y\}$, where a_1, \dots, a_x (respectively b_1, \dots, b_y) appear in a path in T in order; i.e. a_1 appears first in the path and a_x appears last. (Note that, for any i , a_i might not be adjacent to a_{i+1} .) The algorithm then outputs $\min_{(a_i, b_j) \in A \times B} \text{Cost}(a, b)$.*

After t_p pre-processing time, the algorithm takes $O(\ell \log \ell \cdot t_c(m))$ amortized time to answer each query, where $\ell = |A| + |B|$. For weighted graph, the pre-processing time is t_p^{sg} , and query time is $O(\ell \log \ell \cdot t_c^{\text{sg}}(m))$.

Proof. The algorithm follows the schematic of Algorithm 3.2. The only difference is that the set A and B are not contiguous sequence of edges. But we can pretend the edges which does not take part in A or B as *collapsed*. We first look at the case where the graph is unweighted.

1. The algorithm prepares the data-structure from Lemma 6.11,
2. Given any query (A, B) , the algorithm runs Algorithm 3.2. In line 4 of Algorithm 3.2, it needs to issue a total of $O(x)$ many 2-respecting cut queries (i.e., queries of the form $C(u^\downarrow, v^\downarrow)$) in each level of recursion. From Theorem 6.11, such queries can be answered in $t_c(m)$ amortized time. Hence, in each level of recursion, the algorithm needs time $O(x \cdot t_c(m))$. There are at most $\log x$ levels of recursion, which implies the total running time is $O(x \log x \cdot t_c(m))$.

As $\ell > x$ we have the required running time. The case of weighted graph can be handled similarly by using $(t_p^{\text{sg}}, t_c^{\text{sg}})$ -data-structure from Lemma 6.11. \square

We look at a related lemma for implementing the INTERVAL sub-routine, but this time with contiguous set of edges.

Lemma 6.13. *There exists an algorithm which, given an unweighted graph G and a spanning tree T , can pre-process in time $t_p(m)$ and can answer the following question: Given a path $p \in T$ of length ℓ , the algorithm can output cut-value the smallest exact-2-respecting cut which respects two edges from in time $O(\ell \log^2 \ell \cdot t_c(m))$. For weighted graph, the pre-processing time is $t_p^{\text{sg}}(m)$, and query time is $O(\ell \log^2 \ell \cdot t_c^{\text{sg}}(m))$.*

Proof. This algorithm follows the schematic of Algorithm 3.3. As before, we can pretend that edges of T not in p to be *collapsed*. For the sake of efficiency, the algorithm stores the edges of p in sequence in an array. Also, as before, we first look at the case of unweighted graphs.

1. The algorithm prepares the data-structure from Lemma 6.11,
2. Given any path $p \in T$, the algorithm runs Algorithm 3.3 on p . The algorithm needs to issue $O(\ell \log^2 \ell)$ many exact-2-respecting cut queries (*i.e.*, queries of the form $C(u^\downarrow, v^\downarrow)$) to the data-structure. Hence the total running time is $O(\ell \log^2 \ell \cdot t_c(m))$.

The case of weighted graphs can be handled similarly by using $(t_p^{\text{sg}}, t_c^{\text{sg}})$ -data-structure from Lemma 6.11. \square

Now we look at the most crucial part of our sequential implementation: How to find the set $\mathcal{P}_e^{\text{cross}}$ and $\mathcal{P}_e^{\text{down}}$ efficiently for each $e \in T$? To this end, we introduce a random sampling lemma below. But, before stating the lemma, we need to an assumption on the edge weights of G , which is as follows:

Lemma 6.14 (Initial assumption). *By spending linear time, we can assume that the input graph G has no parallel edge and the maximum edge weight is at most 3λ where λ is the min-cut of G .*

Proof. We can assume, without loss of generality, that G has no parallel edge simply by merging parallel edges into one edge. As pointed out by Karger [Kar99, Kar95], one can modify Matula’s algorithm [Mat93] to get a $(2 + \epsilon)$ -approximate value of λ in linear time. So we can assume that we have $\lambda' \in [\lambda, 3\lambda]$. We then contract all edges of weight more than λ' . Observe that the min-cut does not change after the contraction. (Note: Karger [Kar99] also described a simple n^2 -approximation algorithm, which can be used to prove a weaker lemma than here. Such lemma is also enough for our purpose.) \square

Lemma 6.15 (Finding big $C(u^\downarrow, v^\downarrow)$). *There is a data structure that, after pre-processing a weighted m -edge connected graph $G = (V, E)$ and a spanning tree T of G in $O(t_p(m))$ time, can answer the following type of queries:*

- *Query(u): Given a node u , the algorithm returns $O(\log n)$ sets of edges in G , denoted by $F_1(u), F_2(u), \dots$, such that with high probability*

(I) $|F_i(u)| = O(\log n)$ for all i , and

(II) for any node v where $C(u^\downarrow, v^\downarrow) > \deg(u^\downarrow)/2$, there exists i such that the number of edges in $F(u)$ that connect between u^\downarrow and v^\downarrow , *i.e.*, $|F_i(u) \cap C(u^\downarrow, v^\downarrow)|$ is at least $|F_i(u)|/16$.

The algorithm takes $t_s(m) \sum_i |F_i(u)|$ amortized time to answer each query.

Proof. Intuition: If the graph is unweighted, we can simply sample $\Theta(\log n)$ edges from $\deg(u^\downarrow)$ using the (t_p, t_c, t_s) -data structure from Theorem 6.3 (where we convert edges to points as in Lemma 6.11). We can then use the Chernoff’s bound to argue that if $C(u^\downarrow, v^\downarrow) > \deg(u^\downarrow)/2$ then one-third of the sample edges are between u^\downarrow and v^\downarrow (where we view both u^\downarrow and v^\downarrow as ranges in two orthogonal axes as in Theorem 6.11). For weighted graphs, more has to be done, as follows.

Pre-processing: Recall that, by Lemma 6.14, we can assume that the maximum edge weight is at most 3λ and there is no parallel edge. Let G' be a graph obtained by removing all edges in G of weight less than $\epsilon\lambda/n^2$, for a small enough constant ϵ . Thus, the ratio between the minimum and maximum weight of G' is $3n^2/\epsilon$. To simplify the exposition, we assume that the minimum weight is 1. Partition edges of G' into E_0, E_1, \dots, E_ℓ where E_i is the set of edges whose weights are in the range $[2^i, 2^{i+1})$. (Note that $\ell = O(\log n)$.) Convert each edge in G' into a 2-dimensional point by a post-order traversal in T as in Lemma 6.11. Let P_i be the set of points resulting from converting edges in E'_i . Let D_i be a (t_p, t_c, t_s) -data structure that is given points in P_i to pre-process¹¹.

Query(u): For every i , we query D_i to sample $k = c \log n$ many points, for a big enough constant c , where the queried range is defined to cover edges between u^\downarrow and $V(G) \setminus u^\downarrow$ (which can be done according to Lemma 6.11.) Let $F_i(u)$ denote the set of edges corresponding to points returned by D_i . In other words, $F_i(u)$ is a set of random edges in between u^\downarrow and $V(G) \setminus u^\downarrow$ in E'_i .

Analysis: By Theorem 6.3, the processing requires $\sum_i t_p(|P_i|) = O(t_p(m))$ time and the query outputs $O(\ell k) = O(\log^2 n)$ edges with high probability. The number of points sampled is $\sum_i |F_i(u)|$, and hence the time required to sample these many points is $t_s(m) \sum_i |F_i(u)|$. This is the time required to answer Query(u). Additionally these data structures require $O(\log n)$ amortized time to output each point, which becomes an output edge.

Now it is left to show (II). That is, if we let v be any node such that $C_G(u^\downarrow, v^\downarrow) > \deg_G(u^\downarrow)/2$, then for some i

$$|F_i(u) \cap \mathcal{C}(u^\downarrow, v^\downarrow)| \geq |F_i(u)|/16. \quad (3)$$

Let $w(E')$ be the total weight of any set of edges E' . Note that $w(E(G) \setminus E(G')) \leq \epsilon\lambda$. Since $\deg_{G'}(u^\downarrow) \geq \lambda$ for any node u (since it defines a cut), we have

$$C_{G'}(u^\downarrow, v^\downarrow) > (1 - 2\epsilon)C_G(u^\downarrow, v^\downarrow) > (1 - 2\epsilon) \deg_G(u^\downarrow)/2.$$

By the averaging argument we have that there is some i such that $w(E_i \cap \mathcal{C}_{G'}(u^\downarrow, v^\downarrow)) > \frac{1-2\epsilon}{2}w(E_i \cap \mathcal{C}_G(u^\downarrow))$. Since all edges in E_i have weights within a factor of two from each other, we have

$$|E_i \cap \mathcal{C}_{G'}(u^\downarrow, v^\downarrow)| > \frac{1 - 2\epsilon}{4} |E_i \cap \mathcal{C}_G(u^\downarrow)|.$$

Thus, when we sample $k = c \log n$ edges from $E_i \cap \deg_G(u^\downarrow)$ to construct $F_i(u)$, the expected number of edges in $E_i \cap \mathcal{C}_{G'}(u^\downarrow, v^\downarrow)$ is at least $\frac{1-2\epsilon}{4}c \log n$. So, by Chernoff's bound, we have

$$|F_i(u) \cap \mathcal{C}_{G'}(u^\downarrow, v^\downarrow)| > \frac{1 - 2\epsilon}{8} |F_i(u)|$$

with high probability when c is large enough. With small enough ϵ , Equation (3) follows. \square

¹¹Even though the edges in P_i have weights in the range $[2^i, 2^{i+1})$, D_i treats all edges equally by disregarding edge weights. More informally, D_i treats the subgraph induced by P_i as an unweighted simple graph. See Remark 6.8 for relevance.

We have a similar claim for down-interesting edges. The proof is similar to that of Lemma 6.15 and is omitted.

Lemma 6.16 (Finding big $C(v^\downarrow, V - u^\downarrow)$). *There is a data structure that, after pre-processing a weighted m -edge connected graph $G = (V, E)$ and a spanning tree T of G in $O(t_p(m))$ time, can answer the following type of queries:*

- *Query(u): Given a node u , the algorithm returns $O(\log n)$ sets of edges in G , denoted by $F'_1(u), F'_2(u), \dots$, such that with high probability*

(I) $|F'_i(u)| = O(\log n)$ for all i , and

(II) for any node v where $C(v^\downarrow, V - u^\downarrow) > \deg(u^\downarrow)/2$, there exists i such that the number of edges in $F'_i(u)$ that connect between u^\downarrow and v^\downarrow , i.e., $|F'_i(u) \cap C(v^\downarrow, V - u^\downarrow)|$ is at least $|F'_i(u)|/16$.

The algorithm takes $t_s(m) \sum_i |F'_i(u)|$ amortized time to answer each query, which is $O(\log^2 n \cdot t_s(m))$ with high probability.

Next, we need a data-structure for heavy-light decomposition which, on an edge query, will provide the set of paths of decomposition which intersects the root-to-leaf path containing that edge. We use the *compressed tree* data-structure from [HT84]. A minor modification of this data-structure will give the following lemma.

Lemma 6.17 ([HT84]). *Consider \mathcal{P} to be a set of paths obtained by heavy-light decomposition of a tree T . Given a tree T on n vertices, there is a data-structure which can pre-process the tree in $O(n)$ times such that it can answer the following type of queries:*

- *Query(e): Given an edge e of the graph, the algorithm returns at most $\log n$ many edges e'_1, \dots, e'_k and a set of paths $\{p_1, \dots, p_k\} \subseteq \mathcal{P}$ such that (i) each e'_i belongs to the same root-to-leaf path that e belongs to, (ii) each e'_i belongs to a distinct path p_i from \mathcal{P} , and (iii) e'_i is the edge closest to the root in p_i .*

The algorithm takes $O(\log n)$ time to answer each such query.

We need one more data-structure for technical purpose.

Lemma 6.18. *Consider a family \mathcal{S} of n many sets of a universe. There is a data-structure which allows the following operations:*

1. **Insertion.** *Given a tuple (S, T, x) where S and T are sets in \mathcal{S} and x is an element in S , the data-structure inserts (S, T, x) in time $O(\log n)$.*
2. **Sequential access.** *At any point, provides sequential access to a set of tuples of the form $(S, \{x_1, \dots, x_s\}, T, \{y_1, \dots, y_t\})$ in $O(1)$ amortized time where, for each x_i , there was an insertion (S, T, x_i) and, for each y_i , there was an insertion (T, S, y_i) .*

Proof sketch. The data-structure maintains an inherent ordering of the sets in \mathcal{S} . Given any tuple (S, T, x) , where $S \prec T$ w.r.t. that ordering, the data-structure first searches whether there is any entry corresponding to (S, T) in it. This can be searched in $O(\log n)$ time by binary-search. If there is such an entry, the data-structure adds x to the set corresponding

to S in that entry. Otherwise, the data-structure creates an entry corresponding to (S, T) at the lexicographically correct position (w.r.t. the ordering of \mathcal{S}) and adds x to the set corresponding to S in that entry. When $T \prec S$, the data-structure works exactly in the same way, but does the search with the tuple (T, S) instead of (S, T) for not creating double entry. This data-structure can be very efficiently maintained by traditional data-structure, such as, B+ tree which also provides the quick sequential access as claimed. \square

At this point, we look at the implementation of Algorithm 3.13 in the sequential model. The algorithm, given a weighted graph G and a spanning tree T , will construct a few data-structures to efficiently perform the computations, but the schematics will remain more-or-less the same. We explain the steps as follows:

Initial data-structure creation: The algorithm first creates a data-structure D_1 as in Lemma 6.11. This can be done in time $t_p^{\text{sg}}(m)$. For unweighted graph, the time needed is $t_p(m)$. The algorithm also initializes a min-heap **Heap**.

Line 1 to 3: For each edge e of T , the algorithm issues query $\text{deg}(u^\downarrow)$ to D_1 where e is the parent edge of e , and inserts it in **Heap**. This can be done in time $O(n \cdot t_c^{\text{sg}}(m))$ for weighted graphs, and $O(n \cdot t_c(m))$ for unweighted graphs.

Line 4 : This can be done in time $O(n)$. At this point, the algorithm also creates a data-structure D_2 as in Lemma 6.17.

Line 5 to 8: The algorithm uses D_1 , as in Lemma 6.13, to find out the minimum exact-2-respecting cut in each path $p \in \mathcal{P}$ and inserts it into **Heap**. The total amount of time needed, by Lemma 6.13, is $O(n \log^2 n \cdot t_c(m))$ for unweighted graphs, and $O(n \log^2 n \cdot t_c^{\text{sg}}(m))$ for weighted graphs.

Line 9 to 12: 1. The algorithm first creates data-structures D_3^{cross} as in Lemma 6.15 and D_3^{down} as in Lemma 6.16. The algorithm also initiates a data-structures D_4^{cross} and D_4^{down} as in Lemma 6.18.

2. The algorithm iterates over each vertex $u \in V$.

- (a) The algorithm queries u in D_3^{cross} to obtain the set $\{F_i(u)\}_i$. Note that the number of edges in $\bigcup_i F_i(u)$ is at most $O(\log^2 n)$. This requires time $O(\log^2 n \cdot t_s(m))$.
- (b) For each edge $e \in \bigcup_i F_i(u)$, let v be the end-point of e not in u^\downarrow . The algorithm queries v in D_2 to obtain a set of edges $\{e'_1, \dots, e'_k\}$ of size $O(\log n)$ along with the paths $\{p'_1, \dots, p'_k\}$ where $e'_i \in p'_i$. Hence, in total, $O(\log^3 n)$ many edges corresponding to e are collected in this process—we call this set E_e^T . Each query takes time $O(\log n)$, and hence in total $O(\log^3 n)$ time is required.
- (c) The algorithm then inserts $O(\log^3 n)$ many tuples of the form (p, p', e') in D_4^{cross} where $e \in P$, $e' \in E_e^T$, and p' is the path in \mathcal{P} which contains e' . Each insertion takes time $O(\log n)$, and hence $O(\log^4 n)$ time is required in total.
- (d) The algorithm also follows Lemma 6.16 to query D_3^{down} to obtain $O(\log n)$ many sets $\{F'_i(u)\}_i$, and executes exactly similar steps: For every edge $e \in \bigcup_i F'_i(u)$, the algorithm queries D_2 to obtain all potential interesting set of edges and

path (up to u), and inserts all these $O(\log^3 n)$ tuples in D_4^{down} . The algorithm also inserts all tuples of the form (p, p', e') in D_4^{down} where $e \in p$, e' is an ancestor of e and $e' \in p'$. Total time required is, as before, $O(\log^4 n)$.

- (e) As the algorithm iterates over each edge $u \in V$, the total time required in $O(n(\log^4 n + \log^2 n \cdot t_s(m)))$.

Line 13 to 26: The algorithm uses D_1 , as in Lemma 6.12, to find the minimum 2-respecting-cuts for each tuple in D_4^{cross} and D_4^{down} . A similar amortized analysis as that in Section 5.1 or 5.2 will yield a running time of $O(n \log n \cdot t_c(m))$ for unweighted graphs, and $O(n \log n \cdot t_c^{sg}(m))$ for weighted graph.

At the end, the algorithm finds the minimum entry from **Heap** and outputs it. Let us compute the total running time for this implementation when the graph is weighted. The total running time of Algorithm 3.13 is (ignoring lower order terms)

$$O(t_p^{sg}(m) + t_p(m) + n \log^2 n \cdot t_s(m) + n \log^2 n \cdot t_c^{sg}(m) + n \log^4 n),$$

which is $O(m \log n + n \log^4 n)$ if we plug in the value of $t_p^{sg}(m), t_c^{sg}(m), t_p(m), t_s(m)$ from Theorem 6.3 and Theorem 6.10. This implementation does not use any bit operation. Note that we can use the range-sampling data-structure where bit-operations are allowed instead of the one where bit-operations are not allowed—Unfortunately this does not improve the running time as t_p^{sg} dominates t_p . In the unweighted case, however, we can replace the semigroup range searching data structure by the range counting data structure, and so we do not need to use the (t_p^{sg}, t_c^{sg}) -data-structure. Hence the running time in this case is

$$O(t_p(m) + n \log^2 n \cdot t_s(m) + n \log^2 n \cdot t_c(m) + n \log^4 n),$$

which is $O(m\sqrt{\log n} + n \log^4 n)$ if we put values of $t_p(m), t_c(m), t_s(m)$ from Theorem 6.3 and 6.3. Clearly, this case requires bit operations.

We will now look at the correctness of this implementation. The most interesting part of this algorithm is Line 9 to 12. From Lemma 6.15, we know that if v is cross-interesting to u in G , then there is an i such that $|F_i(u) \cap \mathcal{C}(u^\downarrow, v^\downarrow)| \geq |F_i(u)|/16$. Let pick such an edge $e' \in F_i(u) \cap \mathcal{C}(u^\downarrow, v^\downarrow)$ and let v' be the vertex of e' in v^\downarrow . Note that v belongs to the same root-to-leaf path as v' . Hence, all edges of the root-to-leaf path which are cross-interesting w.r.t. e will be discovered by a query to D_2 with v' ¹². As we iterate over all i , the set \mathcal{P}_e^{cross} will be discovered w.h.p. A similar argument can be made for \mathcal{P}_e^{down} using Lemma 6.16. This immediately implies the following claim.

Claim 6.19. *In Line 9 to 12, the algorithm finds, for every edge e , the set \mathcal{P}_e^{cross} and \mathcal{P}_e^{down} .*

The correctness of the rest of the implementation follows from the correctness of the data-structures (Lemma 6.11, 6.12 and 6.13).

¹²In fact, we do not need such a strong guarantee on $F_i(u) \cap \mathcal{C}(u^\downarrow, v^\downarrow)$ for this purpose. We just need that $|F_i(u) \cap \mathcal{C}(u^\downarrow, v^\downarrow)| \neq \emptyset$.

7 Open problems

In this work we show that weighted min-cut can be solved efficiently in many settings. One setting that remains difficult is for dynamic graphs. Improved dynamic algorithms (or conditional lower bounds) for maintaining min-cut exactly and approximately are both very challenging. For exact min-cut, the tree packing can be maintained using existing dynamic minimum spanning tree algorithms (e.g. [HK99, HdLT01, Wul13, KKM13, NSW17]). The challenge is again maintaining the minimum 2-respecting cut. See [Tho07] for the state of the art of approximately maintaining the min-cut. Additionally, it is interesting to see if the number of passes can be improved in the streaming setting.

All efficient min-cut algorithms are *randomized* (there is a near-linear time deterministic algorithm in the sequential setting, but it works only on simple graphs [KT19]). Can we obtain an efficient *deterministic* algorithm in any of the settings considered in this paper? Note that our algorithm is randomized due to the need of a cut sparsifier. A cut sparsifier can be obtained deterministically in almost-linear time using the recent result in [CGL⁺19], but the approximation ratio is too high ($n^{o(1)}$). Improving the approximation ratio is the first important step.

Two problems related to min-cut, namely directed min-cut and vertex connectivity, are still widely open. We leave as a major open problem designing an efficient algorithm for any of these problems either in the sequential, streaming, and cut-query setting, or any other interesting computational model. One model that is of our interest is the two-party communication model where edges are partitioned into two parties. In this setting, the two party can compute the min-cut by communicating $\tilde{O}(n)$ bits. Can they use the same amount of communication for computing the minimum directed cut or vertex connectivity? In fact, this question is also open and interesting for other graph problems such as the maximum bipartite matching.

Acknowledgement

We thank Michal Dori and Sorrachai Yingchareonthawornchai for actively participating in the discussions during their visits at KTH, when we were trying to prove a communication complexity lower bound for the interval problem (which turns out to be impossible due to our results). We would also like to thank Yuval Efron, Joseph Swernofsky and Jan van den Brand for fruitful discussions on different stages of this work.

We thank Timothy Chan for pointing out relevant results on range searching data structures. We also thank Paweł Gawrychowski for pointing out an error in the earlier version, where we used Chazelle’s range counting instead of semigroup range searching data structure. (This change does not affect our results.)

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme under grant agreement No 715672. Authors are also supported by the Swedish Research Council (Reg. No. 2015-04659 and 2019-05622).

References

- [ACK19] Sepehr Assadi, Yu Chen, and Sanjeev Khanna. Polynomial pass lower bounds for graph streaming algorithms. In *STOC*, pages 265–276. ACM, 2019.
- [AGM12a] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Analyzing graph structure via linear measurements. In *SODA*, pages 459–467. SIAM, 2012.
- [AGM12b] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Graph sketches: sparsification, spanners, and subgraphs. In *PODS*, pages 5–14. ACM, 2012.
- [AKM⁺87] Alok Aggarwal, Maria M. Klawe, Shlomo Moran, Peter W. Shor, and Robert E. Wilber. Geometric applications of a matrix-searching algorithm. *Algorithmica*, 2:195–208, 1987.
- [BGKS15] Maxim A. Babenko, Pawel Gawrychowski, Tomasz Kociumaka, and Tatiana Starikovskaya. Wavelet trees meet suffix trees. In *Proceedings of the 26th SODA*, pages 572–591, 2015.
- [BKR96] Rainer E. Burkard, Bettina Klinz, and Rüdiger Rudolf. Perspectives of monge properties in optimization. *Discret. Appl. Math.*, 70(2):95–161, 1996.
- [CGL⁺19] Julia Chuzhoy, Yu Gao, Jason Li, Danupon Nanongkai, Richard Peng, and Thatchaphol Saranurak. A deterministic algorithm for balanced cut with applications to dynamic connectivity, flows, and beyond. *CoRR*, abs/1910.08025, 2019.
- [Cha88] Bernard Chazelle. A functional approach to data structures and its use in multidimensional searching. *SIAM J. Comput.*, 17(3):427–462, 1988.
- [CP10] Timothy M. Chan and Mihai Patrascu. Counting inversions, offline orthogonal range counting, and related problems. In *Proceedings of the 21st SODA*, pages 161–173, 2010.
- [DEMN21] Michal Dory, Yuval Efron, Sagnik Mukhopadhyay, and Danupon Nanongkai. Distributed weighted min-cut in nearly-optimal time. *STOC*, 2021.
- [FKM⁺05] Joan Feigenbaum, Sampath Kannan, Andrew McGregor, Siddharth Suri, and Jian Zhang. On graph problems in a semi-streaming model. *Theor. Comput. Sci.*, 348(2-3):207–216, 2005.
- [GKP12] Ashish Goel, Michael Kapralov, and Ian Post. Single pass sparsification in the streaming model with edge deletions. *CoRR*, abs/1203.4900, 2012.
- [GMW20] Pawel Gawrychowski, Shay Mozes, and Oren Weimann. Minimum cut in $o(m \log^2 n)$ time. In *ICALP*, volume 168 of *LIPICs*, pages 57:1–57:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [GMW21] Pawel Gawrychowski, Shay Mozes, and Oren Weimann. A note on a recent algorithm for minimum cut. In *Symposium on Simplicity in Algorithms (SOSA)*, pages 74–79. SIAM, 2021.

- [GNT20] Mohsen Ghaffari, Krzysztof Nowicki, and Mikkel Thorup. Faster algorithms for edge connectivity via random 2-out contractions. In *SODA*. SIAM, 2020.
- [HdLT01] Jacob Holm, Kristian de Lichtenberg, and Mikkel Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *J. ACM*, 48(4):723–760, 2001. Announced at STOC 1998.
- [HK99] Monika Rauch Henzinger and Valerie King. Randomized fully dynamic graph algorithms with polylogarithmic time per operation. *J. ACM*, 46(4):502–516, 1999. Announced at STOC 1995.
- [HRW17] Monika Henzinger, Satish Rao, and Di Wang. Local flow partitioning for faster edge connectivity. In *SODA*, pages 1919–1938. SIAM, 2017.
- [HT84] Dov Harel and Robert Endre Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM J. Comput.*, 13(2):338–355, 1984.
- [Kar95] David Ron Karger. *Random Sampling in Graph Optimization Problems*. PhD thesis, Stanford, CA, USA, 1995. UMI Order No. GAX95-16851.
- [Kar99] David R. Karger. Random sampling in cut, flow, and network design problems. *Math. Oper. Res.*, 24(2):383–413, 1999. announced at STOC’94.
- [Kar00] David R. Karger. Minimum cuts in near-linear time. *J. ACM*, 47(1):46–76, 2000. announced at STOC’96.
- [KKM13] Bruce M. Kapron, Valerie King, and Ben Mountjoy. Dynamic graph connectivity in polylogarithmic worst case time. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 1131–1142, 2013.
- [KLM⁺17] Michael Kapralov, Yin Tat Lee, Cameron Musco, Christopher Musco, and Aaron Sidford. Single pass spectral sparsification in dynamic streams. *SIAM J. Comput.*, 46(1):456–477, 2017.
- [KT19] Ken-ichi Kawarabayashi and Mikkel Thorup. Deterministic edge connectivity in near-linear time. *J. ACM*, 66(1):4:1–4:50, 2019. announced at STOC’15.
- [LMN21] Andrés López-Martínez, Sagnik Mukhopadhyay, and Danupon Nanongkai. Work-optimal parallel minimum cuts for non-sparse graphs, 2021.
- [LS19] Antonio Molina Lovett and Bryce Sandlund. A simple algorithm for minimum cuts in near-linear time. *CoRR*, abs/1908.11829, 2019.
- [Mat93] David W. Matula. A linear time 2+epsilon approximation algorithm for edge connectivity. In *SODA*, pages 500–504. ACM/SIAM, 1993.
- [MNV16] J. Ian Munro, Yakov Nekrich, and Jeffrey Scott Vitter. Fast construction of wavelet trees. *Theor. Comput. Sci.*, 638:91–97, 2016.

- [NSW17] Danupon Nanongkai, Thatchaphol Saranurak, and Christian Wulff-Nilsen. Dynamic minimum spanning forest with subpolynomial worst-case update time. In *FOCS*, pages 950–961. IEEE Computer Society, 2017.
- [Ove88] Mark H. Overmars. Efficient data structures for range searching on a grid. *J. Algorithms*, 9(2):254–275, 1988.
- [RSW18] Aviad Rubinfeld, Tselil Schramm, and S. Matthew Weinberg. Computing exact minimum cuts without knowing the graph. In *Proceedings of the 9th ITCS*, pages 39:1–39:16, 2018.
- [ST83] Daniel Dominic Sleator and Robert Endre Tarjan. A data structure for dynamic trees. *J. Comput. Syst. Sci.*, 26(3):362–391, 1983.
- [Tho07] Mikkel Thorup. Fully-dynamic min-cut. *Combinatorica*, 27(1):91–127, 2007.
- [Vit85] Jeffrey Scott Vitter. Random sampling with a reservoir. *ACM Transactions on Mathematical Software*, 11:37–57, 1985.
- [Wul13] Christian Wulff-Nilsen. Faster deterministic fully-dynamic graph connectivity. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 1757–1769, 2013.
- [Zel11] Mariano Zelke. Intractability of min- and max-cut in streaming graphs. *Inf. Process. Lett.*, 111(3):145–150, 2011.

A Proof of Claim 3.1

A path T of length n has $n - 1$ edges. We consider the edges $\{e_1, \dots, e_{n-1}\}$ of the path T from left to right. This ordering naturally defines left and right vertex of an edge in the path. Given this ordering, we identify each edge e_ℓ with the point ℓ in the Interval problem. For an edge e (which is not in the path) which connects the left vertex of edge e_i and the right vertex of edge e_j , we identify e with an interval $I_e = (i, j)$.

If we consider a 2-respecting cut that respects edges e_i and e_j ($i \leq j$), then we claim that the cost of this cut is equal to $\text{Cost}(i, j)$ in the Interval problem. This is because:

1. For any edge e to contribute to this cut, the edge has to connect a vertex v , which is between right-side of edge e_i and left-side of edge e_j , with a vertex u which is either on the left-side of edge e_i or right-side of edge e_j . The interval I_e , will cover either i or j , but not both, and hence contributes to $\text{Cost}(i, j)$. In Figure 3, where we are interested in $\text{Cost}(2, 5)$, interval I_3 is such an interval.
2. For any edge e' , which connects the left-side of edge e_i to the right-side of edge e_j , I_e covers both i and j . For any edge e'' , which is contained between the right-side of edge e_i to the left-side of edge e_j , $I_{e'}$ covers none of i and j . In both cases, the corresponding intervals do not contribute to $\text{Cost}(i, j)$. In Figure 3, interval I_1 is the first kind of interval, and I_2 is the second kind. Neither of I_1 and I_2 contributes to $\text{Cost}(2, 5)$.

Note that this argument goes both ways, *i.e.*, $\text{Cost}(i', j')$ represents the cost of a 2-respecting min-cut which respects edge $e_{i'}$ and $e_{j'}$.