Parallel Minimum Cuts in Near-linear Work and Low Depth

Conference Paper**Author(s):**

Geissmann, Barbara; Gianinazzi, Lukas

Publication date:

2021

Permanent link:

<https://doi.org/10.3929/ethz-b-000504069>

Rights / license:

[Creative Commons Attribution 4.0 International](#)

Originally published in:

ACM Transactions on Parallel Computing 8(2), <https://doi.org/10.1145/3460890>

Parallel Minimum Cuts in Near-linear Work and Low Depth

BARBARA GEISSMANN and LUKAS GIANINAZZI, Department of Computer Science,
ETH Zurich

We present the first *near-linear work* and *poly-logarithmic depth* algorithm for computing a *minimum cut* in an undirected graph. Previous parallel algorithms with poly-logarithmic depth required at least quadratic work in the number of vertices.

In a graph with n vertices and m edges, our randomized algorithm computes the minimum cut with high probability in $O(m \log^4 n)$ work and $O(\log^3 n)$ depth. This result is obtained by parallelizing a data structure that aggregates weights along paths in a tree, in addition exploiting the connection between minimum cuts and approximate maximum packings of spanning trees.

In addition, our algorithm improves upon bounds on the number of cache misses incurred to compute a minimum cut.

CCS Concepts: • **Theory of computation** → **Parallel algorithms**;

Additional Key Words and Phrases: Minimum cut; graph algorithms; minimum path data structure; parallel algorithms; cache-oblivious algorithms

ACM Reference format:

Barbara Geissmann and Lukas Gianinazzi. 2021. Parallel Minimum Cuts in Near-linear Work and Low Depth. *ACM Trans. Parallel Comput.* 8, 2, Article 8 (July 2021), 20 pages.
<https://doi.org/10.1145/3460890>

1 INTRODUCTION

Two trends have emerged in microprocessor design in the past two decades: (1) larger caches allow fast access to recently used memory locations, and (2) many processing elements can be placed on the same chip, allowing for massively parallel processing. This has led to interest in both algorithms that take caches into account and parallel algorithm in a variety of different settings.

We consider *shared-memory parallel* algorithms for computing a *minimum cut*—a fundamental subject in graph theory that has many applications in practice, such as in network reliability [18] and cluster analysis [3, 15, 32]. Our algorithm is based on one of the fastest known minimum cut algorithms, Karger’s algorithm [18]. It exploits a random edge-sampling technique and returns the correct result with high probability. Recently, we presented a cache-efficient variant of that algorithm [10]. In this article, we build on that result by parallelizing a key data structure and obtain a parallel minimum cut algorithm with low overhead (poly-logarithmic in the number of vertices) compared to the sequential one.

Authors’ addresses: B. Geissmann and L. Gianinazzi, Department of Computer Science, Universitätsstrasse 6 / CAB, ETH Zurich, Zurich, Switzerland; emails: barbara.geissmann@inf.ethz.ch, lukas.gianinazzi@inf.ethz.ch.

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

© 2021 Association for Computing Machinery.

1539-9087/2021/07-ART8 \$15.00

<https://doi.org/10.1145/3460890>

Fig. 1. The vertex shading indicates the vertex partition of the minimum cut, which has value 2. The crossing edges are dashed.

We identify two main challenges in parallelizing graph algorithms. The first challenge is how to parallelize graph searches, since traversing a graph in parallel is problematic, especially when the graph has large diameter. The second challenge is that many graph algorithms (including those for minimum cuts [19, 34]) employ intricate data structures for good performance. This is also problematic, because repeatedly accessing a data structure creates a sequential bottleneck due to the need to avoid concurrent operations.

Our *randomized* parallel minimum cut algorithm solves the first challenge by computing spanning trees that determine the order in which the edges of the input graph are accessed. In contrast to graphs, spanning trees can be traversed efficiently in parallel. Additionally, using parallel sorting, we rearrange the edges of the input graph to the order dictated by the traversal of the spanning tree, which avoids having to naively search the graph.

To solve the second challenge, we perform many data structure operations at once and in parallel. This works, because the control flow of our algorithm does not depend on the result of the data structure operations until the very end, when the results from all data structure operations are aggregated efficiently in parallel.

1.1 Preliminaries

1.1.1 Graphs. We consider an undirected weighted graph G with vertices V , edges E , and positive edge weights $w: E \mapsto \mathbb{N}^+$. The number of vertices $|V|$ is n and the number of edges $|E|$ is m .

A nonempty proper subset of the vertices V is a *cut* C of the graph G . A cut C induces a partition of the vertices into two nonempty sets C and $\bar{C} = V - C$. An edge $\{u, v\}$ that has endpoints in different parts of the partition ($u \in C$ and $v \in \bar{C}$) *crosses* the cut C : It is a *crossing edge*. The total weight of the crossing edges of a cut C is the *value* of C . A cut with the smallest value is a *minimum cut*. In particular, a disconnected graph has a minimum cut of value 0. See Figure 1 for an example of a minimum cut.

In this work, we will also consider directed trees. For two vertices u and v in a directed tree, we say that v is a *descendant* of u (and u is an *ancestor* of v) if there is a directed path from u to v . In particular, every vertex is its own descendant and its own ancestor. We shall denote the set of all descendants of v as $v^\downarrow$ and the set of all ancestors of v as $v^\uparrow$.

1.1.2 Model of Computation. The well-known **parallel random access machine (PRAM)** model [31] consists of a set of p processors each connected to an unbounded *shared memory*. This shared memory is organized into word-sized addressable locations. In each time step, every processor can read $O(1)$ memory locations and perform an $O(1)$ computable function on those words (this includes basic arithmetic, logic, control-flow, and addressing computations). Then, each processor can write back into $O(1)$ locations in the shared memory. The processors are *synchronous*, which means that they proceed at the same speed and complete a time step together.

The *runtime* of a PRAM algorithm is the number of time steps until the result is available in the shared memory. The runtime is determined by the processor that needs the most time steps.

We allow for multiple processors to read from the same memory location in the same time step, but forbid writing to the same location in the same time step. This is called the **concurrent-read exclusive-write (CREW)** PRAM setting.

The Work-Depth model [2] abstracts further from a concrete machine. In particular, a computation is viewed as a **directed acyclic graph (DAG)**, where each node in the graph corresponds to a constant time operation. The out-edges of a node correspond to the outputs of the operation executed in this node, and its $O(1)$ in-edges correspond to the inputs to the operation. Consequently, the input of an algorithm is given at a designated set of nodes, and the output has to be available at another set of designated nodes.

The *work* of an algorithm is its number of nodes in the computation DAG (not counting the input nodes). The *depth* of an algorithm is the longest path from an input node to an output node.

Observe that the Work-Depth model and the PRAM model are closely related, as every PRAM algorithm can be viewed as generating a computation DAG. Conversely, the work and depth bounds translate to PRAM bounds. An algorithm with work W and depth D takes $O(W/p + D)$ time using p processors in CREW PRAM [2, 4].

1.1.3 Randomization. We assume that in each time step, each processor has access to a uniformly random and independent bit. We distinguish between two types of randomized algorithms: A *Monte Carlo* randomized algorithm returns the correct result *with high probability*. This means that the probability to return the wrong result can be made smaller than $1/n^c$ for any constant c . In particular, increasing c by a constant factor only changes the runtime by a constant factor.

1.2 Related Work

1.2.1 Relation to Maximum Flow. The minimum cut problem is a variant of the minimum s - t cut problem, where the two designated vertices s and t must be in different parts of the partition.

The well-known Minimum-Cut-Maximum-Flow theorem [27] says that the value of a minimum s - t cut equals the value of a maximum s - t flow in the same network. Many maximum s - t flow algorithms exist [7, 13, 22], the best of which obtain $O(mn + n^2 \log n)$ runtime in general [13].

1.2.2 Deterministic Minimum Cut. A minimum cut can be computed by fixing an arbitrary vertex s and computing a maximum s - t flow for all different vertices $t \neq s$. In general, such an approach leads to work $\Omega(mn^2)$ using the known algorithms. However, the ideas from computing maximum flows can be adapted [14] to a sequential algorithm with work $O(mn \log(n^2/m))$.

Slightly better bounds of $O(mn + n^2 \log n)$ are obtained by a relatively simple approach based on a graph search [34]. This approach is a simplification of an approach by Nagomochi and Ibaraki [26].

For *unweighted* graphs, a recent results [16, 23] obtain near-linear work for a deterministic (sequential) minimum cut algorithm.

1.2.3 Randomized Minimum Cut. Randomized algorithms obtain both better work and depth than deterministic algorithms. Karger and Stein [21] give a Monte Carlo algorithm with $O(n^2 \log^3 n)$ work and $O(\log^3 n)$ depth, which is faster than any known maximum flow algorithm when $m = \Omega(n \log^3 n)$. Karger's algorithm [19] takes $O(m \log^3 n)$ work. However, the parallel variant of that algorithm uses $O(n^2 \log n)$ work to obtain $O(\log^3 n)$ depth. Two recent algorithms take $O(m \log^2 n)$ work [9] and $O(m \log^2 n / \log \log n + n \log^6 n)$ work [25].

1.3 Our Contributions

Our main contribution is a randomized parallel minimum cut algorithm that has near-linear work $O(m \log^4 n)$ and low depth $O(\log^3 n)$. It returns the correct result with high probability. Previous parallel algorithms [19, 21] with poly-logarithmic depth have quadratic work $\Omega(n^2 \log n)$ in the

Table 1. Bounds for Computing a Minimum Cut

	Work	Depth
Lowest Work [9]	$\Theta(m \log^2 n)$	$O(m \log^2 n)$
Best Previous Polylog-Depth [19]	$\Theta(n^2 \log n)$	$\Theta(\log^3 n)$
<i>This Paper</i>	$O(m \log^4 n)$	$O(\log^3 n)$

All algorithms are randomized and return the minimum cut with high probability.

number of vertices. Our new algorithm, presented in Section 4, is thus much more work-efficient when the graphs are not too dense, that is, $m \leq o(n^2 / \log^3 n)$. Table 1 compares our results to previous work.

As part of our solution, we present a parallel algorithm to solve a type of constrained minimum cut problem. Given a spanning tree T of the graph, we find the cut of smallest value under the additional constraint that at most 2 crossing edges are part of the spanning tree T . See Figure 2 for an illustration of the problem. Our algorithm has work $O(m \log^2 n)$ and depth $O(\log^2 n)$. The best previous algorithm [19] with poly-logarithmic depth has $\Theta(n^2)$ work and $\Omega(\log n)$ depth.

We solve the constrained minimum cut problem by parallelizing a data structure that maintains aggregates of weights along paths in a tree. In this data structure, we consider a fixed tree where every vertex has a variable weight. The queries find the smallest weight in a path of the tree. The updates add a fixed weight to every vertex in a path of the tree (potentially changing many weights). In Section 3, we show how to answer a batch of k mixed queries and updates on a tree of n vertices with $O(k \log n (\log n + \log k) + n \log n)$ work and $O(\log^2 n + \log n \log k)$ depth. Hence, the average work per query and update is $O(\log^2 k)$, when $k \geq \Omega(n)$.

Our new approach also improves the number of cache misses to compute a minimum cut in the cache-oblivious model [5, 8], where width of a cache line is B and the size of the cache is M . As discussed in Section 5, it incurs $O(\lceil (m \log^4 n) / B \rceil)$ cache misses and takes $O(m \log^4 n)$ computation time. The best previous result [10] incurs $\Theta(\lceil (m (\log^4 n) \log_M n) / B \rceil)$ cache misses and takes $\Theta(m \log^5 n)$ computation time. Hence, the new algorithm improves the number of cache misses by a factor $\Theta(\log n / \log M)$, and the computation time by a factor $\Theta(\log n)$.

2 BACKGROUND

2.1 Karger's Minimum Cut Algorithm

On a high level, Karger's randomized algorithm [19] consists of two main steps:

- (1) Find a set of spanning trees S (with a special property as described below) in a graph G . Each of those trees gives rise to a more constrained optimization problem:
- (2) For each spanning tree $T \in S$, compute the smallest cut C of G that has at most two crossing edges in T (at most two edges of T cross C). See Figure 2 for an illustration.

A key insight of Karger is a randomized procedure to find a set S of only $O(\log n)$ spanning trees such that the smallest cut found in the second step is a minimum cut with high probability.

The set S is constructed using an approximate maximum packing procedure [29]. The tree-packing procedure consists of sampling a sparse subgraph and performing a series of $O(\log^2 n)$ minimum spanning tree computations. It can therefore be parallelized using known algorithms.

LEMMA 2.1 (KARGER [19] §6.2). *In $O(\log^3 n)$ time using $O(m + n \log n)$ processors, a set of $O(\log n)$ spanning trees S of a graph G can be computed with the following property: With high probability, there is a tree $T \in S$ such that a minimum cut of G cuts at most 2 edges of T .*

Fig. 2. Example illustrating the constrained optimization problem. The cuts are illustrated by the vertex shading. The left cut cuts one edge of the spanning tree and the right one cuts two edges of the spanning tree. Cuts that cut more edges of the spanning tree do not have to be considered in the constrained optimization problem.

Fig. 3. Illustration of Minimum Path operations. Node v_i stores value w_i . *Left:* MINPATH(v_4) computes the minimum of the weights of the highlighted nodes. *Right:* ADDPATH(v_8, x) adds x to the weight of the highlighted nodes.

The sequential bottleneck of Karger's algorithm lies in its second step. Although finding the smallest cut that cuts exactly *one edge* of a given spanning tree turns out to be relatively easy, finding the smallest cut that cuts exactly *two edges* of a given spanning tree is challenging. Karger gives a $O(\log^3 n)$ depth algorithm to do so, but it performs $\Theta(n^2 \log n)$ work. And for their faster $O(m \log^3 n)$ work algorithm no efficient parallelization is known.

2.2 Minimum Path

The problem left open by Karger is how to compute the smallest cut that cuts exactly *two edges* of a given spanning tree in parallel. The *sequential* $O(m \log^3 n)$ work minimum cut algorithm builds a data structure called *Minimum Path* on each of the spanning trees obtained in the first step and maintains weights on the vertices of the spanning tree, which correspond to certain estimates of the minimum cut.

Given a rooted tree T with n vertices where each vertex v has a weight $w(v)$, a *Minimum Path* structure supports:

- MINPATH(v): Returns the smallest weight of a vertex on the path from v to the root of T .
- ADDPATH(v, x): Adds x to the weight of all vertices on the path from v to the root of T .

See Figure 3 for an illustration of the two operations.

Fig. 4. An operation in the tree is decomposed into operations on the paths in the decomposition.

This problem is a special case of dynamic trees [33], which take $\Theta(\log n)$ time per query and update. That this is optimal in the pointer machine setting follows from a recent lower bound on dynamic prefix sums [28]. The dynamic tree data structure [33] is difficult to parallelize, because it is based on tree rotations. Therefore, for our *parallel* Minimum Path structure in Section 3, we will take a different approach.

2.3 Monotone Minimum Paths

We proceed with an overview of the data structure underlying the cache-oblivious algorithm minimum cut algorithm [10]. It is designed such that each operation accesses the memory in a monotone order, that is, if an operation accesses location x before location y no operation accesses location y before location x . This enables executing a batch of operations by sweeping through the memory only once, thus incurring a small number of cache misses. Similarly, this *monotonicity* is also crucial for our parallel variant in Section 3.

2.3.1 Tree Decomposition. The first step is to decompose the given tree T into vertex-disjoint paths. To query or update a path P in the tree T , simply perform an operation for each path in the decomposition that intersects the path P .

It is possible to decompose the tree such that each root-to-leaf path is decomposed into $O(\log n)$ parts [10, 33]. Figure 4 illustrates such a decomposition and how a query on the tree corresponds to a set of queries on the paths in the decomposition. We present a parallel procedure to compute a tree decomposition in Section 3.3.

2.3.2 List View. We view each path of the decomposition as a list, where the vertex closest to the root is at the beginning (front) of the list. We call the deductive problem about querying and updating prefixes of lists *Minimum Prefix*. Specifically, we define *MinPrefix* analogously to *MinPath* and *AddPrefix* analogously to *AddPath* on a list L of vertices $(v_1, \dots, v_n)$ with weights $(w_1, \dots, w_n)$:

$\text{MINPREFIX}(v_i)$ returns the smallest weight in the prefix $w_1, \dots, w_i$. The example below illustrates $\text{MinPrefix}(v_6)$ on a list with eight nodes.

Fig. 5. At each inner node b_i , store the difference between the smallest weight in b_i 's right and left subtree. An update changes the differences on a single root-to-leaf path. The example above illustrates an update to the first six elements in the list.

$\text{ADDPREFIX}(v_i, x)$ adds x to the weights of the vertices in the prefix $v_1, \dots, v_i$. The example below illustrates $\text{AddPrefix}(v_4, x)$ on a list with eight nodes.

To perform Minimum Prefix operations on a list l we build a complete binary tree B on top the vertices $v_1, \dots, v_n$, such that the vertices form the leaves of the tree B . This tree holds auxiliary information that allows us to perform the prefix operations quickly. A naive attempt would be to store in each inner node of the tree B the minimum value in its subtree. This approach answers queries in $O(\log n)$ time, but an update takes $\Omega(n)$ time (for example when the prefix covers the whole list).

A better approach is to store only differences of minima: Each inner node stores the difference between the smallest leaf in its right subtree and the smallest leaf in its left subtree. With this approach, only $O(\log n)$ values change when the list is updated, namely those on the path from the root of the binary tree to the leaf corresponding to the last vertex that is updated. See Figure 5 for an illustration.

Moreover, these differences (more specifically their signs) suffice to determine which subtree, and, consequently, which vertex of the list, contains the smallest value.

In the following, we describe how to efficiently maintain these difference values and how to use them to compute the smallest weight in a given prefix of the list.

2.3.3 AddPrefix. For any node b with right child r and left child l , let $\min_i(b)$ be the smallest weight of any descendant of b after the i th update and let $\min_0(b)$ be the smallest weight of any descendant of b in the initial state. Recall that this value is not stored directly for efficiency reasons, but instead, the data structure stores in each node b at time i the value

$$\Delta_i(b) := \min_i(r) - \min_i(l).$$

Let $\text{ADDPREFIX}(v, x)$ be the i th update and assume that we know by how much the update changes the minimum in the right subtree (i.e., $\min_i(r) - \min_{i-1}(r)$) and the minimum in the left subtree (i.e., $\min_i(l) - \min_{i-1}(l)$). Then, we can derive by how much the difference between the two subtree minima changes in this update. Therefore, we define $\phi_i(b) = \min_i(b) - \min_{i-1}(b)$. We have

$$\Delta_i(b) = \Delta_{i-1}(b) + \phi_i(r) - \phi_i(l).$$

Of course, it would be too expensive to compute $\Delta_i(b)$ by recursively computing the values ϕ_i of both children, since all descendants would be traversed. However, we observe that at least one child of every node has a trivial value: If all descendant leaves of a node b' are in the prefix of the list that is updated, then $\phi_i(b') = x$. If no descendant leaves of a node u are in the prefix of the list that is updated, then $\phi_i(u) = 0$. This means that the value of ϕ_i is only non-trivial for the nodes on the path from v to the root (which are also the only nodes where Δ changes). Hence, to perform the i th update $\text{ADDPREFIX}(v, x)$, we walk up along this path in the tree starting in the leaf v and compute $\phi_i(b)$ for each node b along this path. For the leaf, this is $\phi_i(v) = x$. For an interior node b , Figure 6 and Figure 7 illustrate how the value of $\phi_i(b)$ is computed. There are two additional symmetric cases.

2.3.4 MinPrefix. The queries also proceed walking up the path from the last vertex in the prefix to the root. We consider the weights and state of the data structure at a fixed time and for simplicity, we omit the time subscripts.

When computing $\text{MINPREFIX}(v_k)$, it is tempting to directly compute for each node along this path the result of $\text{MINPREFIX}(v_k)$ restricted to the current subtree (namely $\min_{v_i \in b^\downarrow, i \leq k} w_i$). Unfortunately, this is not possible using only the value of $\Delta(b)$. Instead, we compute the difference $d(b)$ between this quantity and the minimum of the current subtree $\min(b)$:

$$d(b) := \left(\min_{v_i \in b^\downarrow, i \leq k} w_i \right) - \min(b).$$

Once this difference is known for the root, we get the result of $\text{MINPATH}(v_k)$ by adding the overall minimum to this difference. See Figure 8, Figure 9, and Figure 10 for an illustration for how $d(b)$ is computed bottom-up. The case from Figure 9 is symmetric to Figure 8.

2.3.5 Running Time. The runtime depends on the way we decompose the tree T into paths. Each path in T is decomposed into at most $O(\log n)$ lists. As we can perform MINPREFIX and ADDPREFIX on each list in the decomposition in $O(\log n)$ time, the overall time to perform ADDPATH and MINPATH is $O(\log^2 n)$.

3 PARALLEL MINIMUM PATH

The monotone Minimum Path structure from Section 2.3 can execute Queries and Updates one-by-one. Directly trying to parallelize such an approach leads to problems of *concurrency*: When many updates try to change the same memory location concurrently, these conflicts need to be resolved. Some general techniques to do so are known in practice, such as locks or lock-free methods [17]. However, these approaches essentially serialize accesses to the same memory location. Thus, locations that are accessed by many updates, such as the root in the minimum path structure, become sequential bottlenecks.

Therefore, we take a different approach. We start with two observations implied by cache-oblivious algorithm [10]:

- (1) The complete sequence of updates and queries is known upfront: it is enough to find a parallel algorithm to perform a *batch* of minimum path operations.
- (2) The data structure operations traverse the memory in a fixed order (namely walking bottom-up in some trees). This allows us to perform all of the updates at once, simulating the execution of all updates at the same time by logically sweeping the trees bottom-up and producing for each memory location all its intermediate states at once.

In the cache-oblivious algorithm, the batch of updates is simulated using a priority queue. This is inherently sequential. Moreover, it leads to a runtime of $\Theta(\log^3 n)$ per minimum path operation.

Fig. 6. AddPrefix (A): In case the minimum stays in the right subtree after the i th update, the change in minimum at node b is given directly by the change of minimum in the right subtree.

Fig. 7. AddPrefix (B): In case the minimum was in the left subtree before the i th subtree, and is in the right subtree after the i th update, the difference between the minimum of the two subtrees before the i th update needs to be taken into account.

In Sections 3.1 and 3.2 we instead provide an explicit schedule for performing a batch of updates and queries *in parallel*, such that the overall work per minimum path operation is $\Theta(\log^2 n)$.

3.0.1 Problem Statement. We consider a *batch* of k MINPATH and ADDPATH operations $o_1, \dots, o_k$. Each ADDPATH operation is of the form $o_i = (i, v(i), x(i))$ for a time i , a vertex $v(i)$, and a weight $x(i)$. Similarly, each MINPATH operation o_i is of the form $o_i = (i, v(i))$ for a time i and a vertex $v(i)$. The goal is to compute the result of all MINPATH operations at once in parallel, as if the operation sequence were executed sequentially.

As in the sequential case, we decompose the tree into a set of directed paths (See Section 3.3) and build a minimum prefix structure for each of those paths. We start with the presentation of *parallel* ADDPREFIX and *parallel* MINPREFIX.

3.1 Parallel AddPrefix

We consider a batch of ADDPREFIX operations for a list of length n , where each operation o_i is of the form $o_i = (i, v(i), x(i))$ for a time i , a vertex $v(i)$, and a weight $x(i)$. Conceptually, we build the same binary tree B on top of the list as in the sequential case (see Section 2.3).

To allow for future queries, we need to produce all intermediate states of a particular node that arose when the updates were executed sequentially. From inspecting the update equation for a particular node b with left child l and right child r ,

$$\Delta_i(b) = \Delta_{i-1}(b) + \phi_i(r) - \phi_i(l),$$

we make the following key observation that it telescopes to:

$$\Delta_i(b) = \Delta_0(b) + \left(\sum_{j=1}^i \phi_j(r) \right) - \left(\sum_{j=1}^i \phi_j(l) \right).$$

Fig. 8. MinPrefix (A): The easy case for $\text{MinPrefix}(v_k)$ is when $\min(b)$ is in the left subtree l (i.e., $\Delta_i(b) > 0$). Then, if the query vertex v_k is also in b 's left subtree, $d(b)$ is copied from the left child l . If the query vertex v_k is in the right subtree r , then both $d(b)$ and $d(l)$ are zero. In any case, we have $d(b) = d(l)$.

Fig. 9. MinPrefix (B): The smallest weight in b 's subtree $\min(b)$ is in the right subtree r and the smallest weight in the query prefix $w^* = \min_{v_i \in b^l, i \leq k} w_i$ is in the right subtree l as well. In that case, $d(b)$ is copied from the right child r . This case occurs when $\Delta_i(b) \leq 0$, the query vertex v_k is in the right subtree l , and $d_i(r) + \Delta_i(b) < 0$.

Fig. 10. MinPrefix (C): The smallest weight in the query prefix $w^* = \min_{v_i \in b^l, i \leq k} w_i$ is in the left subtree l , but the smallest weight in b 's subtree $\min(b)$ is in the right subtree r . The difference $\Delta(b)$ between the minimum in the right subtree and the minimum in the left subtree needs to be taken into account. This case occurs when $\Delta_i(b) \leq 0$ and either the query vertex v_k is in the left subtree l or $d_i(r) + \Delta_i(b) \geq 0$.

Therefore, given the all-prefix-sums for $\phi_1(r), \dots, \phi_k(r)$ and for $\phi_1(l), \dots, \phi_k(l)$, we can compute the values of $\Delta_1(b), \dots, \Delta_k(b)$ in $O(1)$ depth and $O(k)$ work. Moreover, given those values, the values $\phi_1(b), \dots, \phi_k(b)$ each only depend on a constant number of already computed terms. Therefore, they can also be computed in $O(1)$ depth and $O(k)$ work. Unfortunately, this naive

Fig. 11. Example illustrating the definition of H . The set $H(b_i)$ keeps track of the indexes of the updates that are relevant at node b_i . These are exactly those that update a prefix that ends in a descendant of b_i . The figure indicates the value of H for all inner nodes and the leaves v_1, v_2, v_7, v_8 .

interpretation still needs far too many processors: $\Omega(k)$ work at every node b , and hence $\Omega(nk)$ work overall.

Fortunately, not every update is relevant to every node: $\Delta_i(b)$ only changes for an update $o_i = (i, v(i), x(i))$ if $v(i)$ is a descendant of b . We can therefore restrict our attention to the times i of those operations: Let $H(b)$ be the union over all times i where the update is such that $v(i)$ is a descendant of b .

We continue with three important observations.

OBSERVATION 1. $H(l)$ and $H(r)$ are disjoint, and $H(b) = H(l) \cup H(r)$. This implies that we can merge the updates relevant to the children to obtain the updates relevant to b . Note that an update is relevant to $\log n + 1$ nodes of the tree, namely those along the path from the root of B towards the last node in the list prefix that changes. See Figure 11 for an example.

OBSERVATION 2. The Observation 1 allows us leave out all the indices that do not occur in $H(b)$, as they are not relevant to b :

$$\Delta_i(b) = \Delta_0(b) + \left(\sum_{j \in H(b)} \phi_j(r) \right) - \left(\sum_{j \in H(b)} \phi_j(l) \right), \text{ for any } i \in H(b).$$

OBSERVATION 3. In a bottom up traversal of the tree, we get $\phi_i(l)$ for all $i \in H(l)$ and $\phi_i(r)$ for all $i \in H(r)$. This means that we have some “missing” values for the sums in Observation 2, namely the values of $\phi_i(r)$ for all $i \in H(l)$ and the values of $\phi_i(l)$ for all $i \in H(r)$. As explained for the sequential operation (see Section 2.3.3), these values are trivial to get:

- Updates $(i, v(i), x(i))$ with $i \in H(l)$ do not affect the descendants of r . Hence, for such an $i \in H(l)$ we have $\phi_i(r) = 0$.
- Updates $(i, v(i), x(i))$ with $i \in H(r)$ change all descendants of l by the same amount: for such an $i \in H(r)$ we have $\phi_i(l) = x(i)$.

Observations 1–3 suggest a parallel bottom-up procedure that produces for each node b :

- The values of $H(b)$ as an array in sorted order. In the following the i th largest value in $H(b)$ is denoted $H(b)[i]$. This means that $H(b)[i]$ is the index (w.r.t. the batch of updates) of the i th update relevant at node b .
- An array $X(b)$ stores the increment of the updates relevant at b :

$$X(b)[i] = x(j), \quad \text{where } j = H(b)[i].$$

This means that the i th update relevant at b is of the form $\text{AddPath}(v(j), X(b)[i])$ for some descendant $v(j)$ of b .

- An array $\Phi(b)$ storing $\phi_j(b)$ for the updates relevant at b . Specifically, we have

$$\Phi(b)[i] = \phi_j(b), \quad \text{where } j = H(b)[i].$$

Moreover, the procedure produces for each interior node b the array $\Delta(b)$ containing the intermediate states of the data structure at all relevant times:

$$\Delta(b)[i] = \Delta_j(b), \quad \text{where } j = H(b)[i].$$

In the following, we explain how to obtain these values for leafs, inner nodes, and the root of the binary tree B .

3.1.1 Leafs. At the leafs, we just need to group the updates by vertex and keep track of the relevant quantities. First, apply a stable sort by vertex to the operation sequence. Now, the operation sequence is sorted by vertex and tuples with the same vertex are sorted by time. From this, we can easily find which operations belong to a given vertex v (by a binary search for v and its successor). Then, for each vertex v initialize $H(v)$ as the times of the tuples that have v as a vertex and initialize $X(v)$ with the corresponding increments. Set $\Phi(v)$ to $X(v)$ - recall that $\Phi(v)$ records how much the minimum in the subtree changed for every operation relevant at v and at the leafs the minimum changes by the increment.

3.1.2 Inner Nodes. At an inner node b with left child l and right child r , we merge the results from its children, use prefix sums to generate $\Delta(b)$ in parallel, and construct $\Phi(b)$ based on $\Delta(b)$.

Using Observation 1, we merge the two sorted arrays $H(l)$ and $H(r)$ in parallel to receive $H(b)$ in sorted order. Similarly, we merge the update increments $X(l)$ and $X(r)$ to obtain $X(b)$ in sorted order.

Using Observation 3, we reconstruct the missing values for the right child r and merge them with $\Phi(r)$ that we got from the right child. Proceed similarly for the left child l . Now, for all times i relevant to b (i.e., for all $i \in H(b)$), we have $\phi_i(l)$ and $\phi_i(r)$ each in an array sorted by increasing time i .

Using Observation 2, we construct $\Delta(b)$ as follows. We compute (in parallel) the all-prefix-sums over the array that contains the $\phi_i(l)$ for all i in $H(b)$, and the array that contains the $\phi_i(r)$ for all i in $H(b)$. Then, the observation immediately implies $\Delta_i(b)$ for all i in $H(b)$ in parallel.

Finally, we compute each entry of $\Phi(b)$ in parallel:

$$\Phi(b)[i] = \begin{cases} \Phi(l)[i] & \text{if } \Delta(b)[i-1] > 0 \text{ and } \Delta(b)[i] > 0, \\ \Phi(l)[i] - \Delta(b)[i-1] & \text{if } \Delta(b)[i-1] \leq 0 \text{ and } \Delta(b)[i] > 0, \\ \Phi(r)[i] & \text{if } \Delta(b)[i-1] \leq 0 \text{ and } \Delta(b)[i] \leq 0, \\ \Phi(r)[i] + \Delta(b)[i-1] & \text{if } \Delta(b)[i-1] > 0 \text{ and } \Delta(b)[i] \leq 0. \end{cases}$$

3.1.3 Root. For the root ρ of the binary tree, we proceed as for an internal node. Additionally, we generate the overall minimum weight after every update, which will be needed for the parallel

MINPATH queries. Observe that

$$\min_i(\rho) = \min_0(\rho) + \sum_{j=1}^i \phi_j(\rho).$$

Hence, we compute all the values $\min_1(\rho), \dots, \min_k(\rho)$ by a parallel all-prefix-sums computation on $\phi_1(\rho), \dots, \phi_k(\rho)$.

3.1.4 Running Time. By the proceeding described above, we obtain the following bounds for parallel ADDPREFIX operations.

LEMMA 3.1. *Performing a batch of k parallel ADDPREFIX operations on a list of length n takes $O(k(\log n + \log k) + n)$ work and depth $O(\log n \log k)$.*

PROOF. At the leafs, we sort the operation sequence and perform $O(k)$ parallel binary searches. This takes $O(k \log k)$ work and $O(\log k)$ depth. At every inner node b , we perform a constant number of parallel all-prefix-sums operations and parallel merge operations on arrays of size $O(H(b))$. By the bounds of merging sorted arrays [6] and parallel prefix sums [31], the work at an inner node b is $O(H(b) + 1)$ and the depth is $O(\log(H(b)) + 1)$.

All nodes with the same distance to the root can be processed in parallel. Thus, the total work arising from nodes at distance i to the root is $O(k + 2^i)$ and the depth is $O(\log k)$. Here, we used that summing $H(b)$ over all nodes at a fixed distance i gives exactly k , because each leaf is the descendant of exactly one node at distance i . Nodes with different distance from the root need to be processed by decreasing distance (bottom up). Hence, the overall work arising at inner nodes is $O(k \log n + n)$ and the overall depth is $O(\log n \log k)$.

At the root, we perform an additional parallel all-prefix-sums operation on an array of size k . $\square$

3.2 Parallel MinPrefix

The parallel update algorithm (batch of ADDPATH operations) produces all intermediate states for the data structure. If we store for each node in the data structure all the state it ever has (sorted by time), then the value of a cell after the i th update can be determined by doing a binary search on those states, taking $\Theta(\log k)$ time. Each query is then performed independently. Overall, this takes $\Theta(k \log n \log k)$ work and $\Theta(\log k \log n)$ depth, which would be $\Omega(\log k)$ more work compared to the original data structure.

To get rid of this logarithmic factor in work, we also perform the queries in a batch and we use parallel merging to avoid the binary searches. We obtain the following bounds:

LEMMA 3.2. *Performing a batch of k MINPREFIX and ADDPREFIX operations on a list of length n takes $O(k(\log n + \log k) + n)$ work and depth $O(\log n \log k)$.*

The procedure is similar as for the updates. The queries are placed at the leafs and the nodes at the same distance from the root are processed in parallel.

For a single query $(i, v(i))$, the following happens: The query gets processed bottom-up in all nodes on the path P_i from the leaf $v(i)$ to the root, such that every internal node b obtains an intermediate result $d_i(c)$ from its child c , updates this result based on $\Delta_i(b)$ to $d_i(b)$, and passes this result to its parent node. Initially, leaf $v(i)$ sets $d_i(v(i)) = 0$. Then, every internal node b with left child l and right child r updates this result to

$$d_i(b) = \begin{cases} d_i(l) & \text{if } \Delta_i(b) > 0, \\ d_i(r) & \text{if } v(i) \in r^\downarrow, \text{ and } d_i(r) + \Delta_i(b) < 0, \\ d_i(l) - \Delta_i(b) & \text{otherwise,} \end{cases}$$

Fig. 12. The illustrated tree has 4 boughs, indicated by the vertex colors. Each bough starts at a leaf and continues upwards until reaching the first node that has a sibling.

where $d_i(l)$ and $d_i(r)$ have either been already computed, since the child lies on the path P_i , or are equal to zero.

To process all queries in parallel, every leaf now initiates an array that contains all its queries with the corresponding intermediate results sorted by time. Hence, an internal node b obtains such an array L from its left child and R from its right child, which it merges (in parallel, by time) to an array Q that now contains all queries and intermediate results relevant to b .

The last issue is to have the Δ values ready for each relevant query. In particular, we need the value that belongs to the last update before the query.

Thus, we record the time when this value was set by translating the indices of b 's relevant updates to the times that correspond to it: $\Delta(b)[i]$ is mapped to the tuple $(H(b)[i], \Delta(b)[i])$. Then, we merge Q and the relevant Δ -values in parallel and sorted by time. The new array contains a mix of queries and Δ -values, sorted by time, such that now each query just needs to read the last Δ -value on its left. This is achieved by a segmented broadcast, where each Δ -value broadcasts its value to all following queries. A segmented broadcast can be implemented using a variant of the parallel all-prefix-sums algorithm [31].

At the root, each query needs to read the overall minimum at the closest preceding time to the query. This can be achieved similarly using parallel merging and a segmented broadcast.

The overall runtime analysis is very similar to the one for the **ADDPREFIX**, since processing node b has the same cost: It has work $O(H(b) + 1)$ and depth $O(\log(H(b)) + 1)$.

3.3 Tree Decomposition Using Boughs

To solve **MINPATH** and **ADDPATH** for general trees, we need a suitable decomposition of the tree T into paths. The idea is to repeatedly remove certain subpaths that start at a leaf, as follows: We call a path that starts at a leaf and ends at the first vertex that has a sibling on the way up to the root, a *bough* of T . A bough ends at a vertex whose parent has multiple children or, if the tree T is a path, at the root. See Figure 12 for an example.

The algorithm partitions the tree into paths as follows, until no edges remain:

- (1) Identify the boughs of T . Each bough is in the decomposition.
- (2) Remove all vertices that are part of a bough.

Note that shrinking boughs is also used for other parallel graph problems [30]. The algorithm is also related to a parallel tree contraction algorithm [24]. Observe that, in each repetition, the number of leaves is at least halved. This is because each leaf in the new tree had at least two children before the contraction. Hence, there are at most $\log_2 n$ repetitions. This implies that every root-to-leaf path in T is decomposed into at most $\log_2 n$ paths.

There is a work-efficient deterministic algorithm to identify the tree decomposition based on parallel expression tree evaluation [11, 12] [Appendix A].

LEMMA 3.3 (GIANINAZZI AND HOEFLE [11]). *A tree T with n vertices can be decomposed into a set of pairwise vertex-disjoint paths P such that:*

- Each root-to-leaf path in T intersects at most $\log_2 n$ paths in P .
- This takes work $O(n)$ and depth $O(\log n)$.

3.4 Parallel MinPath and AddPath

Each MINPATH and ADDPATH operation corresponds to $O(\log n)$ MINPREFIX and ADDPREFIX operations, respectively, which can be processed in parallel. For the MINPATH operations, the smallest result of the $O(\log n)$ MINPREFIX queries can be found sequentially after they have completed. We conclude:

LEMMA 3.4. *Performing a batch of k MINPATH and ADDPATH operations on a tree with n nodes takes $O(k \log n(\log n + \log k) + n \log n)$ work and depth $O(\log n(\log k + \log n))$.*

4 PARALLEL MINIMUM CUTS

The parallel *Minimum Path* structure is the missing puzzle-piece in a parallelization of Karger's algorithm. This solved, it remains to show how to create the batch of Minimum Path operations and how to combine the results of the Minimum Path structure into a minimum cut, achieving the following overall bounds.

THEOREM 4. *The minimum cut of a graph can be computed with depth $O(\log^3 n)$ and work $O(m \log^4 n)$ (Monte Carlo randomized).*

In the following, we consider a (rooted) spanning tree T of the input graph G and we want to compute the smallest cut that cuts at most two edges of T . We will do so with work $O(m \log^3 n)$ and depth $O(\log^2 n)$. Together with Lemma 2.1 this implies the main result.

Karger already showed a parallel algorithm that computes the smallest cut that cuts *exactly one* edge of a given spanning tree. In fact, the algorithm computes for each vertex v , the value of the cut $v^\downarrow$ that has the descendants of v on one side of the cut (and therefore cuts only the edge from v to its parent in T).

LEMMA 4.1 (KARGER [19]). *A smallest cut that cuts exactly one edge of a given spanning tree can be computed with work $O(m)$ and depth $O(\log m)$.*

We therefore focus on giving a parallel algorithm for the case where *exactly two* edges of the spanning tree are cut.

4.1 Cutting Two Edges of A Spanning Tree

Our parallel algorithm uses ideas from Karger's sequential algorithm [20] to reduce the problem to a set of Minimum Path operations, which we already showed how to perform in parallel in Section 3.

We are given a spanning tree T of G . Assume that the smallest cut C that cuts at most two edges of G cuts the edges (u, v) and (s, t) of T (where u is the parent of v and s the parent of t). Let us focus on the case where u is not an ancestor of s and vice versa. See Figure 13 for an example. The case where u is an ancestor of s is similar.

If we take the value of the cut $t^\downarrow$ and add the value of the cut $v^\downarrow$, then we incorrectly count (twice) the edges that go between the descendants of v and the descendants of t . We will use the Minimum Path structure to keep track of these "extra" edges that go between the two parts of the tree.

In the following, we will explain how the algorithm handles the different possibilities for these four vertices, including for instance if one of v or t is a leaf.

Fig. 13. The weight of the (black, white) cut that cuts (u, v) and (s, t) is given by the dashed edges. This equals the edges with at least one black endpoint minus the edges with two black endpoints.

4.1.1 Handling a Leaf. Assume that v is a leaf of T (see Figure 13 for an illustration). The algorithm proceeds as follows:

- (1) Initialize a Minimum Path Structure where the initial weight of each vertex x is given by the value of the cut $x^\downarrow$. Recall that these values are computed by the algorithm from Lemma 4.1.
- (2) Then, for each edge $e = (v, y)$ incident to v , perform $\text{ADDPATH}(v, -2w(e))$.

Observe that after these steps, the weight in the minimum path structure of a vertex x that is not an ancestor of v is exactly the value of the cut $x^\downarrow$ minus twice the weight of the edges between v and the descendants of x . Moreover, since we assume that C is a smallest cut of G among those that cut at most two edges of T , at least one of v 's neighbors has to be a descendant of t . Otherwise, the cut $t^\downarrow$ would be smaller than the cut C , which leads to a contradiction. This implies that the value of C is given by the value of the cut $v^\downarrow$ plus the minimum weight of a node x , such that x is a neighbor but not an ancestor of v . Thus, we find the value of the cut C as follows:

- (1) Add ∞ to the weight of all ancestors of v by $\text{ADDPATH}(v, \infty)$.
- (2) For each neighbor x of v , call $\text{MINPATH}(x)$ and keep the smallest result.
- (3) Add the value of the cut $v^\downarrow$. This is the value of the cut C .

4.1.2 Handling a Bough. The observations from before can be generalized to also handle boughs. Observe that a bough is a maximal induced subpath that contains a leaf. Similar to the case where v is a leaf, we use the Minimum Path Structure to keep track of the “extra” edges that go between $v^\downarrow$ and $t^\downarrow$. The following procedure handles the case where v is in a bough.

- (1) Initialize the minimum path structure just as in the leaf case.
- (2) Start at the leaf of the bough and walk up the bough. At every node y in the bough:
 - (a) If y is a leaf, then perform $\text{ADDPATH}(y, \infty)$.
 - (b) Moreover, for every edge $e = (y, x)$ incident to y , perform $\text{ADDPATH}(x, -2w(e))$.
 - (c) Afterwards, for every neighbor x of y , perform the query $\text{MinPath}(x)$. Record the smallest result.

Consider the state of the minimum path structure when the above procedure has processed some node y in the bough. Consider another node x that is not an ancestor of any node in the bough. Then, the weight of x is equal to the value of the cut $x^\downarrow$ minus twice the weight of the “extra” edges that exist between $y^\downarrow$ and $x^\downarrow$.

Moreover, by the minimum assumption of C , it must hold that v has a descendant that is a neighbor of a descendant of t .

Putting these two observations together, we conclude that it is indeed enough to perform a $\text{MINPATH}(y)$ query for every neighbor y of a node in the bough and record the smallest result. When we have processed v , the smallest result seen so far plus the value of the cut $v^\downarrow$ gives the value of the cut C .

4.1.3 Handling a General Tree. To handle a general subtree we repeat the procedure for every bough. Then, we contract all edges (of the spanning tree and the overall graph) with at least one endpoint in a bough. Afterwards, we recurse in the new tree. We call such a phase a *bough-phase*. Note that this gives the same decomposition of the tree as in Section 3.3.

In each bough-phase, the boughs can be handled in any order. However, after handling a bough, we need to restore the weights to their initial state. This is done by reversing the order and the sign of all **ADDPATH** operations: We visit the nodes in the bough top-down and replace each **AddPath**(x, w) by **AddPath**($x, -w$). In the end, we return the smallest cut value found.

4.2 Generating the Batch of Minimum Path Operations

We show how to generate the batch of minimum path operations for one bough-phase. The algorithm from Lemma 3.3 computes the boughs of the tree. The remaining difficulty is to compute (in parallel) the order in which the edges are accessed.

Observe that each edge is accessed at most four times: for each of its endpoints once on the way going upwards in the bough containing this endpoint, and once on the way down. We get the order and the operations as follows.

- (1) Order each bough by list-ranking and give each leaf a unique identifier. Then, the order in which a vertex is visited is derived from the number of the leaf of its bough and its position in the list-ranking. See Figure 14 for an example.
- (2) Each leaf creates an **ADDPATH**(v, ∞) at the time of its first visit and an **ADDPATH**($v, -\infty$) at the time of its second visit.
- (3) When a node y in a bough is visited at times t_1 and t_2 , then every neighbor x of y creates an update that corresponds to **ADDPATH**($x, -2w(e)$) at time t_1 and a query that corresponds to **MinPath**(x) at time t_1 . Moreover, it creates an update that corresponds to **ADDPATH**($x, 2w(e)$) at time t_2 (this undoes the operation of the former update). Each edge in a graph is thus accessed at most four times, namely every time one of its endpoints is visited. See Figure 15 for an illustration.
- (4) The queries and updates are sorted according to the visit times, where additionally operations with the same visit time are ordered such that updates come before queries. This gives the operation sequence that handles all the boughs in the tree.

LEMMA 4.2. *Generating the batch of Minimum Path operations from Section 4.1 that handles all the boughs of a tree takes work $O(m \log n)$ and depth $O(\log n)$.*

PROOF. List ranking of the boughs takes $O(\log n)$ depth and $O(m)$ work [1]. The vertices derive their visit time with $O(\log n)$ depth and $O(n)$ work, the edges can then be processed completely in parallel with $O(1)$ depth and $O(m)$ work. Finally, sorting the batch of $O(m)$ operations takes $O(\log n)$ depth and $O(m \log n)$ work [6]. $\square$

4.3 Extracting the Minimum Cut.

To finally find the smallest cut of G that cuts at most two edges of a given spanning tree T , we need to repeatedly apply such batches of operations in every bough-phase:

- (1) For all vertices x , compute the value of the cut $x^\downarrow$ of G that has the descendants of x in T on one side of the cut (using the algorithm from Lemma 4.1). In particular, this yields the smallest cut that cuts exactly one edge of T .
- (2) Find the boughs of T . Contract all edges incident to a node in a bough (contract the edges in T and G at the same time) and recurse until the graph has a single vertex. This generates

Fig. 14. The boughs of the tree are indicated by colors. The node labels indicate a possible order in which they are visited by the algorithm. Each node in a bough is visited twice. Once on a bottom-up traversal of its bough and then on a top-down traversal of its bough. The order between boughs is arbitrary.

Fig. 15. The order in which the edges are visited is given by the order of the vertices. An edge is visited whenever one of its endpoints is visited. Thus, each edge is visited two or four times. The figure indicates the times when the non-tree edges (grey) are visited, based on the example on the left.

a sequence of graphs $G_1, G_2, \dots, G_k$ and corresponding spanning trees $T_1, T_2, \dots, T_k$. Note that $G_1 = G$ and $T_1 = T$.

- (3) In parallel, for each of the pairs T_i and G_i in the sequence, run the algorithm from Section 4.2 to generate the necessary Minimum Path operations.
- (4) For each of the trees T_i , execute the batch of Minimum Path operations in parallel (Section 3.4). Find the smallest value $\text{MinPath}(v) + C(v^\downarrow)$ over all vertices v (and all queries to v), where $C(v^\downarrow)$ is the value of the cut of G that has the descendants of v in T_i on one side of the value. These values have already been computed when finding the smallest cut that cuts one edge of T .
- (5) The smallest value $\text{MinPath}(v) + C(v^\downarrow)$ found overall is the value of the smallest cut that cuts exactly two edges of T .

The algorithm obtains the following bounds. Together with Lemma 2.1, they imply our main result as stated in Theorem 4.

LEMMA 4.3. *Finding the smallest cut of a graph G that cuts at most two edges of a given spanning tree T of G takes work $O(m \log^3 n)$ and depth $O(\log^2 n)$.*

PROOF. Using Lemma 3.3, generating the graphs $G_1, G_2, \dots, G_k$ and corresponding spanning trees $T_1, T_2, \dots, T_k$ takes $O(m \log n)$ work and $O(\log^2 n)$ depth. Since the number of leaves is at least halved in each bough-phase, $k = O(\log n)$. Contracting the edges in the bough is just a matter of mapping each vertex in a bough to the vertex that is the parent of the topmost vertex in the bough, look up the value of the endpoints of each vertex in the mapping and remove loops. It is not necessary to combine parallel edges.

For each graph in $G_1, G_2, \dots, G_k$, generating the queries takes $O(m \log n)$ work and $O(\log n)$ depth (by Lemma 4.2). Executing the batch of $O(m)$ minimum path operations and finding the minimum result takes $O(m \log^2 n)$ work and $O(\log^2 n)$ depth (by Lemma 3.2). As this can be done in parallel for each of the $O(\log n)$ graphs, the algorithm takes $O(m \log^3 n)$ work and $O(\log^2 n)$ depth overall.

The algorithm can be easily adapted to also output the edges of T that define the cut, essentially by recording the edges that generated a Minimum Path query. $\square$

5 CACHE-OBVIOUS ALGORITHM

The parallel minimum cut algorithm gives improved cache miss bounds to compute a minimum cut. The key difference to our previous cache-oblivious algorithm [10] is the implementation of the minimum path structure. The parallel minimum path algorithm from Section 3 uses operations (such as merging sorted lists and computing prefix sums) that are easily made cache-efficient.

The cache-oblivious model [5, 8] considers a fully-associative cache with optimal replacement strategy of size M with cache lines of width B . The parameters B and M of the machine cannot be used in the algorithm description, hence the name *cache-oblivious*.

We obtain the following bounds to compute a minimum cut if we replace the minimum path structure in our previous cache-oblivious algorithm [10] with the data structure from Section 3.

THEOREM 5. *Finding a minimum cut incurs $O(\frac{m \log^4 n}{B} + 1)$ cache misses and takes $O(m \log^4 n)$ time (Monte Carlo randomized).*

6 CONCLUSION

Compared to the best sequential algorithm for sparse graphs [9], our algorithm performs work larger by a factor of $O(\log^2 n)$, namely $O(m \log^4 n)$ work and $O(\log^3 n)$ depth. It remains an open problem to find a work-optimal minimum cut algorithm that has poly-logarithmic depth.

The $\Omega(\log^3 n)$ depth of our algorithm comes from the algorithm to find an approximately maximum tree packing. The packing is the set of spanning trees where for some tree, at most two of its edges cross a minimum cut. The rest of our algorithm has depth $O(\log^2 n)$. Consequently, a lower depth algorithm to find a suitable spanning tree would yield a lower depth minimum cut algorithm.

REFERENCES

- [1] Richard J. Anderson and Gary L. Miller. 1991. Deterministic parallel list ranking. *Algorithmica* 6, 6 (1991), 859–868. <https://doi.org/10.1007/BF01759076>
- [2] Guy E. Blelloch. 1996. Programming parallel algorithms. *Commun. ACM* 39, 3 (Mar. 1996), 85–97. <https://doi.org/10.1145/227234.227246>
- [3] Rodrigo A. Botafogo. 1993. Cluster analysis for hypertext systems. In *Proceedings of the 16th Annual International Conference on Research and Development in Information Retrieval (SIGIR'93)*. 116–125. <https://doi.org/10.1145/160688.160704>
- [4] Richard P. Brent. 1974. The parallel evaluation of general arithmetic expressions. *J. ACM* 21, 2 (1974), 201–206. <https://doi.org/10.1145/321812.321815>
- [5] Gerth St. Brodal. 2004. Cache-oblivious algorithms and data structures. In *Proceedings of the 9th Scandinavian Workshop on Algorithm (SWAT'04)*. 3–13. https://doi.org/10.1007/978-3-540-27810-8_2
- [6] Richard Cole. 1988. Parallel merge sort. *SIAM J. Comput.* 17, 4 (1988), 770–785. <https://doi.org/10.1137/0217049>
- [7] E. A. Dinic. 1970. Algorithm for solution of a problem of maximum flow in a network with power estimation. *Sov. Math Dokl.* 11, 5 (1970), 1277–1280.
- [8] Matteo Frigo, Charles E. Leiserson, Harald Prokop, and Sridhar Ramachandran. 1999. Cache-oblivious algorithms. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science (FOCS'99)*. 285–298. <https://doi.org/10.1109/SFCS.1999.814600>
- [9] Paweł Gawrychowski, Shay Mozes, and Oren Weimann. 2019. Minimum cut in $O(m \log^2 n)$ time. arXiv:1911.01145. Retrieved from <https://arxiv.org/abs/1911.01145>.
- [10] Barbara Geissmann and Lukas Gianinazzi. 2017. Cache oblivious minimum cut. In *Proceedings of the 10th International Conference on Algorithms and Complexity (CIAC'17)*. 285–296. https://doi.org/10.1007/978-3-319-57586-5_24
- [11] Lukas Gianinazzi and Torsten Hoefer. 2020. Parallel planar subgraph isomorphism and vertex connectivity. In *SPAA'20: Proceedings of the 32nd ACM Symposium on Parallelism in Algorithms and Architectures*. 269–280. <https://doi.org/10.1145/3350755.3400259>
- [12] Lukas Gianinazzi and Torsten Hoefer. 2020. Parallel planar subgraph isomorphism and vertex connectivity. arXiv:2007.01199. Retrieved from <https://arxiv.org/abs/2007.01199>.
- [13] Andrew V. Goldberg and Robert E. Tarjan. 1988. A new approach to the maximum-flow problem. *J. ACM* 35, 4 (1988), 921–940. <https://doi.org/10.1145/48014.61051>

- [14] Jianxiu Hao and James B. Orlin. 1994. A faster algorithm for finding the minimum cut in a directed graph. *J. Algor.* 17, 3 (1994), 424–446. <https://doi.org/10.1006/jagm.1994.1043>
- [15] Erez Hartuv and Ron Shamir. 2000. A clustering algorithm based on graph connectivity. *Inf. Process. Lett.* 76, 4-6 (2000), 175–181. [https://doi.org/10.1016/S0020-0190\(00\)00142-3](https://doi.org/10.1016/S0020-0190(00)00142-3)
- [16] Monika Henzinger, Satish Rao, and Di Wang. 2020. Local flow partitioning for faster edge connectivity. *SIAM J. Comput.* 49, 1 (2020), 1–36. <https://doi.org/10.1137/18M1180335>
- [17] Maurice Herlihy and Nir Shavit. 2008. *The Art of Multiprocessor Programming*. Morgan Kaufmann.
- [18] David R. Karger. 1999. A randomized fully polynomial time approximation scheme for the all-terminal network reliability problem. *SIAM J. Comput.* 29, 2 (1999), 492–514. <https://doi.org/10.1137/S0097539796298340>
- [19] David R. Karger. 2000. Minimum cuts in near-linear time. *J. ACM* 47, 1 (2000), 46–76. <https://doi.org/10.1145/331605.331608>
- [20] David R. Karger. 2001. A randomized fully polynomial time approximation scheme for the all-terminal network reliability problem. *SIAM Rev.* 43, 3 (2001), 499–522. <https://doi.org/10.1137/S0036144501387141>
- [21] David R. Karger and Clifford Stein. 1996. A new approach to the minimum cut problem. *J. ACM* 43, 4 (1996), 601–640. <https://doi.org/10.1145/234533.234534>
- [22] Richard M. Karp and Vijaya Ramachandran. 1990. Parallel algorithms for shared-memory machines. In *Handbook of Theoretical Computer Science, Volume A: Algorithms and Complexity (A)*. 869–942.
- [23] Ken-ichi Kawarabayashi and Mikkel Thorup. 2015. Deterministic global minimum cut of a simple graph in near-linear time. In *Proceedings of the 47th Symposium on Theory of Computing (STOC'15)*. 665–674. <https://doi.org/10.1145/2746539.2746588>
- [24] Gary L. Miller and John H. Reif. 1985. Parallel tree contraction and its application. In *Proceedings of the 26th Annual Symposium on Foundations of Computer Science (FOCS'85)*. 478–489. <https://doi.org/10.1109/SFCS.1985.43>
- [25] Sagnik Mukhopadhyay and Danupon Nanongkai. 2020. Weighted min-cut: Sequential, cut-query, and streaming algorithms. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC'20)*. 496–509. <https://doi.org/10.1145/3357713.3384334>
- [26] Hiroshi Nagamochi and Toshihide Ibaraki. 1990. Computing edge-connectivity in multiple and capacitated graphs. In *Proceedings of the International Symposium on Algorithms (SIGAL'90)*. 12–20. https://doi.org/10.1007/3-540-52921-7_51
- [27] Christos H. Papadimitriou and Kenneth Steiglitz. 1998. *Combinatorial Optimization: Algorithms and Complexity*. Courier Corporation, 117–120 pages.
- [28] Mihai Patrascu and Erik D. Demaine. 2006. Logarithmic lower bounds in the cell-probe model. *SIAM J. Comput.* 35, 4 (2006), 932–963. <https://doi.org/10.1137/S0097539705447256>
- [29] Serge A. Plotkin, David B. Shmoys, and Éva Tardos. 1995. Fast approximation algorithms for fractional packing and covering problems. *Math. Operat. Res.* 20, 2 (1995), 257–301. <https://doi.org/10.1287/moor.20.2.257>
- [30] Vijaya Ramachandran. 1997. Parallel algorithms for reducible flow graphs. *J. Algor.* 23, 1 (1997), 1–31. <https://doi.org/10.1006/jagm.1996.0820>
- [31] John H. Reif. 1993. *Synthesis of Parallel Algorithms* (1st ed.). Morgan Kaufmann, San Francisco, CA.
- [32] Roded Sharan and Ron Shamir. 2000. Center CLICK: A clustering algorithm with applications to gene expression analysis. In *Proceedings of the 8th International Conference on Intelligent Systems for Molecular Biology (ISMB'00)*. 307–316.
- [33] Daniel D. Sleator and Robert E. Tarjan. 1983. A data structure for dynamic trees. *J. Comput. Syst. Sci.* 26, 3 (1983), 362–391. [https://doi.org/10.1016/0022-0000\(83\)90006-5](https://doi.org/10.1016/0022-0000(83)90006-5)
- [34] Mechthild Stoer and Frank Wagner. 1997. A simple min-cut algorithm. *J. ACM* 44, 4 (1997), 585–591. <https://doi.org/10.1145/263867.263872>

Received April 2019; revised January 2021; accepted March 2021