Near-Optimal Distributed Computation of Small Vertex Cuts

Merav Parter
Weizmann Institute, Rehovot, Israel

Asaf Petruschka
Weizmann Institute, Rehovot, Israel

Abstract

We present near-optimal algorithms for detecting small vertex cuts in the CONGEST model of distributed computing. Despite extensive research in this area, our understanding of the *vertex* connectivity of a graph is still incomplete, especially in the distributed setting. To this date, all distributed algorithms for detecting cut vertices suffer from an inherent dependency in the maximum degree of the graph, Δ. Hence, in particular, there is no truly sub-linear time algorithm for this problem, not even for detecting a *single* cut vertex. We take a new algorithmic approach for vertex connectivity which allows us to bypass the existing Δ barrier.

- As a warm-up to our approach, we show a simple $\widetilde{O}(D)$-round[1] randomized algorithm for computing all cut vertices in a D-diameter n-vertex graph. This improves upon the $O(D + \Delta/\log n)$-round algorithm of [Pritchard and Thurimella, ICALP 2008].
- Our key technical contribution is an $\widetilde{O}(D)$-round randomized algorithm for computing all cut *pairs* in the graph, improving upon the state-of-the-art $O(\Delta \cdot D)^4$-round algorithm by [Parter, DISC '19]. Note that even for the considerably simpler setting of *edge* cuts, currently $\widetilde{O}(D)$-round algorithms are currently known *only* for detecting pairs of cut edges.

Our approach is based on employing the well-known linear graph sketching technique [Ahn, Guha and McGregor, SODA 2012] along with the heavy-light tree decomposition of [Sleator and Tarjan, STOC 1981] . Combining this with a careful characterization of the survivable subgraphs, allows us to determine the connectivity of $G \setminus \{x, y\}$ for every pair $x, y \in V$, using $\widetilde{O}(D)$-rounds. We believe that the tools provided in this paper are useful for omitting the Δ-dependency even for larger cut values.

2012 ACM Subject Classification Theory of computation → Distributed algorithms; Theory of computation → Graph algorithms analysis

Keywords and phrases Vertex-connectivity, Congest, Graph Sketches

Digital Object Identifier 10.4230/LIPIcs.DISC.2022.31

Related Version *Full Version*: `https://www.weizmann.ac.il/math/parter/sites/math.parter/files/uploads/DistributedVertexCutsDISC22.pdf`

Funding *Merav Parter*: Supported by the European Research Council (ERC) No. 949083), and by the Israeli Science Foundation (ISF) No. 2084/18.

1 Introduction and Our Contribution

The vertex connectivity of the graph is a central concept in graph theory and extensive attention has been paid to developing algorithms to compute it in various computational models. Recent years have witnessed an enormous progress in our understanding of vertex cuts, from a pure graph theoretic perspective [36] to many algorithmic applications [30, 28, 36, 20].

[1] Throughout the paper, we use the notation $\widetilde{O}$ to hide poly-logarithmic in n terms.

Despite this exciting movement, our algorithmic toolkit for handling vertex cuts is still somewhat limited. A large volume of the work, in the centralized setting, has focused on fast algorithms for detecting minimum vertex cuts of size at most k, for some small number k. Until recently, near-linear time algorithms where known only for $k \leq 2$ [39, 21]. A sequence of recent breakthrough results [5, 28] provides an almost-linear time sequential algorithm for computing the vertex connectivity (even for large connectivity values).

As we see soon, the situation is considerably worse in distributed settings, where the problem is still fairly open already for $k = 1$. Throughout, we consider the CONGEST model [35]. In this model, each node holds a processor with a unique and arbitrary ID of $O(\log n)$ bits, and initially only knows the IDs of its neighbors in the graph. The execution proceeds in synchronous rounds and in each round, each node can send a message of size $O(\log n)$ to each of its neighbors. The primary complexity measure is the number of communication rounds. For n-vertex D-diameter graphs, Pritchard and Thurimella [37] presented a randomized algorithm for detecting a (single) cut vertex (a.k.a articulation point) within $O(D + \Delta/\log n)$ CONGEST rounds, where Δ is the maximum degree of the graph. [37] conclude their paper by noting:

> [37] *It would be interesting to know if our distributed cut vertex algorithm could be synthesized with the cut vertex algorithm of [40] to yield further improvement. Alternatively, a lower bound showing that no $O(D)$-time algorithm is possible for finding cut vertices would be very interesting.*

No progress on the complexity of this problem has been done since then. For small cut values k, Parter [33] employed the well-known fault-tolerant sampling technique [42, 25] for detecting k vertex cuts in $(\Delta \cdot D)^{\Theta(k)}$ deterministic rounds. Turning to approximation algorithms, for $k = \Omega(\log n)$, Censor-Hillel, Ghaffari and Kuhn [4] provided a $O(\log n)$ approximation for computing the *value* of the vertex connectivity of the graph within $\widetilde{O}(D + \sqrt{n})$ rounds. [4] also presented a lower bound of $\widetilde{\Omega}(D + \sqrt{n/k})$ V-CONGEST rounds. In the V-CONGEST model, each *node* (rather than an edge) is restricted to send only $O(\log n)$ bits, in total, in every round. As shown in this paper, this lower bound does not hold in the standard CONGEST model.

We follow the terminology of [37]: a *cut vertex* is a vertex x such that $G \setminus \{x\}$ is not connected. A *cut pair* is a pair of vertices x, y such that $G \setminus \{x, y\}$ is not connected. For brevity we call these objects, small cuts. Our main results in this paper are near-optimal algorithms for detecting these small cuts, in the sense that for every small cut, there is at least one vertex in the graph that learns it. Our first contribution is in presenting a (perhaps surprisingly) simple randomized algorithm[2] that can detect all cut vertices in the graph in $\widetilde{O}(D)$ rounds. The edge-congestion[3] of the algorithm is $\widetilde{O}(1)$ bits[4].

▶ **Theorem 1.** *There is a randomized algorithm that w.h.p. identifies all single cut-vertices in G within $\widetilde{O}(D)$ rounds. The edge congestion is $\widetilde{O}(1)$. In the output, each vertex $x \in V$ learns if it is a cut vertex.*

[2] As usual, all presented randomized algorithms in this paper have success guarantee of $1 - 1/n^c$, for any given constant $c > 1$.

[3] The edge congestion of a given algorithm is the worst-case bound on the total number of messages exchanged through a given edge e in the graph.

[4] We exploit this bounded congestion for detecting cut pairs.

This settles the question raised by [37]. Our algorithm is based on the well-known *graph-sketching* technique of Ahn et al. [1]. This technique has admitted numerous applications in the context of connectivity computation under various computational settings, e.g., [23, 22, 18, 26, 29, 17, 11, 8]. Yet, to the best of our knowledge, it has not been employed before in the context of CONGEST algorithms for minimum vertex-cut computation.

We then turn to consider the problem of detecting cut pairs. It has been noted widely in the literature that there is a sharp qualitative difference between a single failure and two failures. This one-to-two jump has been accomplished by now for a wide-variety of fault-tolerant settings, e.g., reachability oracles [6], distance oracles [9], distance preservers [32, 19, 34] and vertex-cuts [21, 2, 3, 13]. While it is relatively easy to extend our algorithm of Theorem 1 to detect cut pairs in $\widetilde{O}(D^2)$ rounds, providing a near-optimal complexity of $\widetilde{O}(D)$ rounds, turns out to be quite involved. Our key technical contribution is:

▶ **Theorem 2.** *There is a randomized algorithm that w.h.p. identifies all cut pairs in G within $\widetilde{O}(D)$ rounds. For each cut pair x, y, either x or y learns that fact.*

We observe that even for the simpler problem of edge-connectivity (see Remark below), an $\widetilde{O}(D)$-round algorithm is currently only known for edge cuts of size at most *two* due to [37]. Hence, we are now able to match the complexity of these two problems for small cut values. Our algorithm is based on distinguishing between two structural cases depending on the locations of the cut pair x, y in a BFS tree of G. The first case which we call *dependent* handles the setting where the x and y have ancestry/descendant relations. The second *independent* case assumes that x and y are not dependent, i.e., $\text{LCA}(x, y) \notin \{x, y\}$, where $\text{LCA}(x, y)$ is the lowest/least common ancestor of x and y in the BFS tree. Each of these cases call for a different approach. We believe that the tools provided in this paper should hopefully pave the way towards detecting larger vertex-cuts with no dependency in the maximum degree Δ (as it is the case for the state-of-the-art algorithm by [33]). For a more in-depth technical overview, see Sec. 1.1.

Remark on the Edge-Connectivity. It is widely known that in undirected graphs, vertex connectivity and vertex cuts are significantly more complex than edge connectivity and edge cuts, for which now the following result are known: an $\widetilde{O}(m)$-time centralized exact algorithm [24, 16, 12] and an $\widetilde{O}(D + \sqrt{n})$ exact distributed algorithms [7]. For constant values of edge-connectivity a $\text{poly}(D)$-round algorithm is given in [33].

1.1 Our Approach, in a Nutshell

We provide the key ideas of our algorithms. Our end goal is to simulate a connectivity algorithm in the graph $G \setminus \{x, y\}$ for *every* pair of vertices $x, y \in V$. Note that this is not trivial already for a single x, y pair as the diameter of the subgraph $G \setminus \{x, y\}$ might be as large as $\Omega(\Delta D)$, hence using on-shelf connectivity algorithms lead to a round complexity of $O(\min\{D + \sqrt{n}, \Delta D\})$. We bypass this Δ dependency by using the edges incident to the vertices x, y as *shortcuts*. Then, to minimize the congestion imposed by running possibly n^2 connectivity algorithms in parallel, we employ a preprocessing phase in which we collect *graph-sketch* information (explained next) at each vertex x. This information allows each vertex x to pinpoint at a bounded number of cut-mate suspects. In addition, it allows x, in certain cases, to locally simulate connectivity queries without using further communication. Throughout, let T be a BFS tree rooted at some source s, and denote the T-paths by $\pi(\cdot, \cdot)$.

We start by employing the well-known *heavy-light tree decomposition* technique by Sleator and Tarjan [38]. This classifies the edges of T into light and heavy edges. The useful properties are that each vertex v has $O(\log n)$ light edges on its tree path $\pi(s,v)$, and in addition, each v is the parent of *one* heavy edge, connecting v to its unique heavy child. It is easy to compute this decomposition on T in $\widetilde{O}(D)$ rounds. For a vertex x, let T_x be the subtree of T rooted at x.

Basic Tools: Graph Sketches and Borůvka Algorithm. A *graph sketch* of a vertex v is a randomized string of $\widetilde{O}(1)$ bits that compresses v's edges [1]. The linearity of these sketches allows one to infer, given the sketches of subset of vertices S, an outgoing cut edge $(S, V \setminus S)$ with constant probability. A common approach for deducing the graph connectivity merely from the sketches of the vertices is based on the well-known Borůvka algorithm [31]. This algorithm works in $O(\log n)$ phases, where in each phase, from each growable component an outgoing edge is selected. All these outgoing edges are added to the forest, while ignoring cycles. Each such phase reduces the number of growable components by constant factor, thus within $O(\log n)$ phases a maximal forest is computed. Since this algorithm only requires the computation of outgoing edges it can simulated using $O(\log n)$ independent sketches for each of the vertices. In our algorithms, we aggregate graph sketches over the BFS tree T which allows the vertices x to locally simulate Borůvka in the graph $G \setminus \{x\}$. This is illustrated in our algorithm for detecting a single cut vertex, described next.

Warm Up: Detecting Single Cut Vertices. Our algorithm starts by letting each vertex v locally compute its individual $\mathsf{Sketch}_G(v)$. Then, by aggregating the sketches (using their linearity) from the leaf vertices to the root s over the BFS tree T, each vertex v learns its subtree-sketch $\mathsf{Sketch}_G(V(T_v))$. Once this is completed, it is easy to let each vertex $x \in V$ learn the G-sketch information of all the connected components in $T \setminus \{x\}$. We then show that x can locally modify these G-sketches into $(G \setminus \{x\})$-sketches. At this point, the vertex x can locally apply the Borůvka algorithm in $G \setminus \{x\}$ and deduce if $G \setminus \{x\}$ is connected. The full details (proof of Theorem 1) appear in Appendix A.

We now turn to consider the considerably more challenging task of detecting cut-pairs. We classify these pairs into dependent and independent.

Detecting Dependent Cut Pairs. Our approach for the dependent case is based on designing algorithms $\{\mathcal{A}_y\}_{y \in V}$, where $\mathcal{A}_y$ detects all xy cut pairs of the form $x \in T_y$. We show that each such an algorithm $\mathcal{A}_y$ can be designed in a way that sends a total of $\widetilde{O}(1)$ messages only along edges incident to $V(T_y)$, and runs in $\widetilde{O}(D)$ rounds. The standard random delay technique allows us then to schedule the execution of all n algorithms $\{\mathcal{A}_y\}_{y \in V}$ within $\widetilde{O}(D)$ rounds. At a high level, each algorithm $\mathcal{A}_y$ is based on employing the single-vertex cut algorithm in the graph $G \setminus \{y\}$. Our challenge is then twofold: first, the diameter of the graph $G \setminus \{y\}$ might be as large as $\Omega(\Delta D)$, and second, communication is restricted to use only edges incident to $V(T_y)$. We overcome these challenges by using y as a coordinator, providing global computation services and communication shortcuts that essentially enables efficient simulation (in both dilation and congestion) of the vertex cut algorithm in $G \setminus \{y\}$.

Detecting Independent Cut Pairs. The most technically involved case is where x, y are independent, namely, *do not* have ancestry relations in T. A-priori, the number of such potential cut-mates y for a given vertex x might be even linear in n. To filter out irrelevant options, the algorithm starts by computing at each vertex x a tree $\widehat{T}_x$ that encodes the

connectivity between s and the vertices in $V_x = V(T_x) \setminus \{x\}$ in the graph $G \setminus \{x\}$. Let $\mathcal{C}_x = \{C_1, \ldots, C_k\}$ denote the collection of maximal connected components in the graph $G[V_x]$. The tree $\widehat{T}_x$ consists of k paths of the form $\pi(s, u_C) \circ (u_C, v_C)$ for every component $C \in \mathcal{C}_x$, where v_C is some representative vertex in C. It is then easy to observe that the potential cut mates y must appear on the paths $\{\pi(s, u_C) \mid C \in \mathcal{C}_x\}$. For a given suspect y, we call the $\mathcal{C}_x$-components C for which $y \in \pi(s, u_C)$, *y-sensitive*. Our argument has the following structure.

Multiple xy-Connectivity Algorithms, Under a Promise. For a fixed xy pair, we design an algorithm $A^P_{x,y}$ that determines the connectivity in $G \setminus \{x, y\}$ *given* an x-y path $\Pi_{x,y}$ (on which x, y can exchange messages). The algorithm $A^P_{x,y}$ has the special property that it sends messages either along $\Pi_{x,y}$, or else along edges incident to a restricted subset of vertices in T_x, T_y, defined as follows. Let $\mathsf{LDS}(x, y) \subset V(T_x)$ be the set of all vertices which are descendants of the *light* children of x, and belong to a y-sensitive component in $\mathcal{C}_x$. The set $\mathsf{LDS}(y, x)$ is defined in an analogous manner. The algorithm $A^P_{x,y}$ is then guaranteed to send $\widetilde{O}(1)$ messages only along $\Pi_{x,y}$ and along edges incident to the vertices of $\mathsf{LDS}(x, y) \cup \mathsf{LDS}(y, x)$. This restriction is crucial in order to run multiple $A^P_{x,y}$ algorithms, for distinct x, y, in parallel. Using the properties of the heavy-light tree decomposition and our sensitivity definition, one can show that each vertex $w \in V$ belongs to the $\mathsf{LDS}(x, y)$ sets of at most $\widetilde{O}(D)$ pairs xy. The main challenge is in bounding the overlap between the $\Pi_{x,y}$ paths, cross distinct xy pairs. We show that given a subset $Q \in V \times V$, the collection of $\{A^P_{x,y} \mid (x, y) \in Q\}$ algorithms can be scheduled in parallel in $\widetilde{O}(D)$ rounds, given that following promise holds for Q:

> [**Promise**:] *There is a path collection $\mathcal{P}_Q = \{\Pi_{x,y} \mid (x, y) \in Q\}$ such that each path has length $O(D)$, and each edge appears on $\widetilde{O}(D)$ paths in $\mathcal{P}_Q$.*

One can show, using the properties of heavy-light decomposition, that each vertex belongs to the $\mathsf{LDS}(x, y)$ sets of at most $\widetilde{O}(D)$ pairs x, y. Hence, by combining this fact with the promise, the algorithms for all the Q pairs can be run in parallel, using the random delay approach [27, 14].

On a high level, each algorithm $A^P_{x,y}$ works by letting x and y jointly simulating the Borůvka algorithm in $G \setminus \{x, y\}$. The main challenge is that the communication is restricted to the edges incident to $\mathsf{LDS}(x, y) \cup \mathsf{LDS}(y, x)$, despite the fact that one should also take into account the remaining vertices in T_x, T_y, e.g., descendants of the *heavy* children of x, y. In each Borůvka phase, we maintain the invariant that x, y jointly hold the sketches of connected-subsets (denoted as *parts*) in $G \setminus \{x, y\}$, where we split the responsibility between x, y in a careful manner. We mainly distinguish between parts that contain a heavy child of x, y and the remaining *light* parts that are contained in $\mathsf{LDS}(x, y) \cup \mathsf{LDS}(y, x)$. The merges of the light parts are implemented by using communication between vertices in $\mathsf{LDS}(x, y) \cup \mathsf{LDS}(y, x)$. The merges concerning the heavy parts are implemented by using the direct xy communication over the $\Pi_{x,y}$ path. Each such Borůvka phase is implemented in $\widetilde{O}(D)$ rounds. At the end of the simulation, x, y both learn whether $G \setminus \{x, y\}$ is connected.

Omitting the Promise Based on Classification Into Light and Heavy Independent Pairs. While the promise clearly holds for $\widetilde{O}(D)$ pairs, it clearly does not hold for all n^2 pairs, in general. Our approach is based on classifying the collection of the xy pairs into two classes: *light* and *heavy*. This classification is based on the trees $\widehat{T}_x, \widehat{T}_y$, as well as on the heavy-light decomposition of T. Informally, for a light pair xy, one can define a $\Pi_{x,y}$ that intersects

a light subtree of either x or y. These paths can be shown to have a bounded overlap, hence satisfying the promise. Handling the heavy pairs is more involved. Here we take a mixed approach. We define a special subset of the heavy pairs for which the promise can be satisfied (denoted as *mutual pairs*). This subset is chosen in a careful way that guarantees the following, perhaps surprising, property: the remaining (not mutual) heavy pairs x, y can be decided locally, at either x or y. Our key observation is that for a xy heavy pair, the graph $G \setminus \{x, y\}$ is connected iff one of the heavy children of x, y is connected to s in $G \setminus \{x, y\}$. Hence, it is mainly essential for x, y to collect a sketch information on the components of these heavy children in $\mathcal{C}_x, \mathcal{C}_y$. This information can be then aggregated over T. The formal implementation of this step (completing the proof of Theorem 2) appears in the full version of the paper.

1.2 Preliminaries

Throughout the paper, we fix a connected n-vertex graph $G = (V, E)$, and a BFS tree T for G rooted at some arbitrary source vertex $s \in V$. We denote the unique tree path from u to v by $\pi(u, v, T)$. When the tree T is clear from context, we may omit it and simply write $\pi(u, v)$. We use the $\circ$ operator for path-concatenation. An (undirected) edge between vertices u,v is denoted by (u, v). For $x, y \in V$, a vertex subset $S \subseteq V$ is said to be xy-connected if all the vertices of S belong to the same connected component of $G \setminus \{x, y\}$.

Heavy-Light Tree Decomposition. We now present our heavy-light terminology, the notion of *compressed paths*, and their distributed computation.

▶ **Definition 3** (Heavy-light decomposition). *For a non-leaf vertex $v \in V(T)$, its* heavy child*, denoted v_h, is the child v' of v maximizing*[5] *the number of vertices in its subtree $T_{v'}$. Any other v-child of v is a* light child*. A tree vertex is* heavy *if it is the heavy child of its parent, and* light *otherwise (so the root s is light). A tree edge is* heavy *if it connects a vertex to its heavy child, and* light *otherwise. If (u, u') is a heavy (resp., light) edge in the path $\pi(s, v)$, then u is a* heavy ancestor *(resp.,* light ancestor*) of v, and v is a* heavy descendant *(resp,* light descendant*) of u. (Note that e.g. a "heavy ancestor" need not be a heavy vertex itself.) We denote by $\mathsf{LA}(v)$ (resp., $\mathsf{LD}(v)$) the set of v's light ancestors (resp., descendants). It is easy to show that $\pi(s, v, T)$ contains $O(\log n)$* light *vertices/edges, hence also $|\mathsf{LA}(v)| = O(\log n)$.*

▶ **Definition 4** (Compressed paths). *Let $v \in V(T)$. Let $L = [s = v_0, v_1, \ldots, v_k]$ be the ordered list of the light vertices on the root-to-v path $\pi(s, v, T)$. The* compressed path *of v with respect to T, denoted $\pi^*(s, v, T)$ consists of the list L, along with a table mapping each v_i to the number of heavy vertices appearing between v_i and v_{i+1} in $\pi(s, v, T)$ (where we define $v_{k+1} = v$). Note that the compressed path $\pi^*(s, v, T)$ has bit-length $O(\log^2 n)$.*

Observe that the compressed paths can be used as ancestry labels *in T: Given the compressed path $\pi^*(s, u, T)$ and $\pi^*(s, v, T)$, one can check whether $\pi(s, u, T)$ is a prefix of $\pi(s, v, T)$, and hence determine whether u is an ancestor of v.*

In our context of distributed computation, we have the following lemma. Missing proofs in this section appear in the full version of the paper.

▶ **Lemma 5.** *For every tree T, there is an $\widetilde{O}(D(T))$-rounds $\widetilde{O}(1)$-congestion algorithm letting each vertex v of T learn its heavy/light classification and its compressed path $\pi^*(s, v, T)$.*

[5] Ties are broken arbitrarily and consistently.

Graph Sketches. We now give a formal but brief definition of graph sketches. We follow [8], and refer the reader to Section 3.2.1 therein for a detailed presentation of the subject. Throughout, let $\oplus$ denote the bitwise-XOR operator. The first required ingredients are randomized *unique edge identifiers*:

▶ **Lemma 6** (Modification of Lemma 3.8 in [8]). *Using a random seed $\mathcal{S}_{ID}$ of $O(\log^2 n)$ random bits, one can compute a collection of $M = \binom{n}{2}$ $O(\log n)$-bit identifiers for the pairs in $\binom{V}{2}$, denoted $\mathcal{I} = \{\mathrm{UID}(e_1), \ldots, \mathrm{UID}(e_M)\}$, with the following property: For any nonempty subset $E' \subseteq E$, $\Pr[\oplus_{e\in E'} \mathrm{UID}(e) \in \mathcal{I}] \leq 1/n^{10}$. Furthermore, for any $e = (u, v)$, the identifier $\mathrm{UID}(e)$ can be computed from $\mathrm{ID}(u)$, $\mathrm{ID}(v)$ and the random seed $\mathcal{S}_{ID}$.*

Next, we define the notion of *extended edge identifiers*, formed by augmenting the $\mathrm{UID}(e)$ with the IDs and the T-ancestry labels of the endpoints based on compressed paths, namely $\mathsf{ANC}_T(v) = \pi^*(s, v, T)$. Formally, an edge $e = (u, v)$ we have

$$\mathrm{EID}_T(e) = [\mathrm{UID}(e), \mathrm{ID}(u), \mathrm{ID}(v), \mathsf{ANC}_T(u), \mathsf{ANC}_T(v)] \ . \tag{1}$$

Equipped with these definitions, we are ready to define the sketches. We now follow [10, 11, 8] and use pairwise independent hash functions for this purpose. Choose $L = c \log n$ pairwise independent hash functions $h_1, \ldots, h_L : \{0,1\}^{\Theta(\log n)} \to \{0, \ldots, 2^{\log M} - 1\}$, and for each $i \in \{1, \ldots, L\}$ and $j \in [0, \log M]$ define the edge set $E_{i,j} = \{e \in E \mid h_i(e) \in [0, 2^{\log M - j})\}$. Each of these hash functions can be defined using a random seed of logarithmic length [41]. Thus, a random seed $\mathcal{S}_h$ of length $O(L \log n)$ can be used to determine the collection of all these L functions. For each vertex v and indices i, j, let $E_{i,j}(v)$ be the edges incident to v in $E_{i,j}$. The i^{th} *basic sketch unit* of each vertex v is then given by:

$$\mathsf{Sketch}_{G,i}(v) = [\oplus_{e\in E_{i,0}(v)} \mathrm{EID}_T(e), \ldots, \oplus_{e\in E_{i,\log M}(v)} \mathrm{EID}_T(e)].$$

We extend the sketches to be defined on vertex subsets by XORing. Namely, for every subset of vertices S, we define $\mathsf{Sketch}_{G,i}(S) = \oplus_{v\in S}\mathsf{Sketch}_{G,i}(v)$. The *sketch* of each vertex v is defined by a concatenation of $L = \Theta(\log n)$ basic sketch units:

$$\mathsf{Sketch}_G(v) = [\mathsf{Sketch}_{G,1}(v), \mathsf{Sketch}_{G,2}(v), \ldots \mathsf{Sketch}_{G,L}(v)] \ .$$

Again, we extend this definition to vertex subsets $S \subseteq V$ by $\mathsf{Sketch}_G(S) = \oplus_{v\in S}\mathsf{Sketch}_G(v)$. The main use of graph sketches is in finding outgoing edges:

▶ **Lemma 7** (Modification of Lemma 3.11 in [8]). *For any subset S, given a basic sketch unit $\mathsf{Sketch}_{G,i}(S)$ and the seed $\mathcal{S}_{ID}$ one can compute, with constant probability[6] $\mathrm{EID}_T(e)$ for an outgoing edge e from S in G, if such exists.*

▶ **Lemma 8.** *Let $S \subseteq V$, and let $E' \subseteq E$ be a set of outgoing edges from S. Then, given $\mathsf{Sketch}_G(S)$, the random seed $\mathcal{S}_h$, and the extended identifiers $\mathrm{EID}_T(e)$ of all $e \in E'$, one can compute the $\mathsf{Sketch}_{G\setminus E'}(S)$.*

Distributed Scheduling. The congestion of an algorithm $\mathcal{A}$ is defined by the worst-case upper bound on the number of messages exchanged through a given graph edge when simulating $\mathcal{A}$. Throughout, we make an extensive use of the following random delay approach of [27], adapted to the CONGEST model.

[6] Over the choice of the random seeds $\mathcal{S}_{ID}$ and $\mathcal{S}_h$.

▶ **Theorem 9** ([14, Theorem 1.3]). *Let G be a graph and let $\mathcal{A}_1, \ldots, \mathcal{A}_m$ be m distributed algorithms, each algorithm takes at most $\mathtt{d}$ rounds, and where for each edge of G, at most $\mathtt{c}$ messages need to go through it, in total* over all *these algorithms. Then, there is a randomized distributed algorithm that w.h.p. runs all the algorithms in $\widetilde{O}(\mathtt{c}+\mathtt{d})$ rounds.*

2 Dependent Cut Pairs

In this section we present an $\widetilde{O}(D)$-rounds distributed algorithm for detecting *dependent* cut pairs in G, i.e. pairs xy where x is a descendant of y in the BFS tree T rooted at s. Recall that our approach is based on scheduling the execution of algorithms $\{\mathcal{A}_y\}_{y\in V}$, where $\mathcal{A}_y$ detects all cut pairs xy such that $x \in T_y$ (see Section 1.1). By employing the single-vertex cut algorithm from Section A as a common preprocessing phase prior to the execution of the $\{\mathcal{A}_y\}_{y\in V}$ algorithms, we may assume that there are no 1-vertex cuts in G. Furthermore, by carefully examining the properties of this algorithm, we may assume that every $v \in V$ holds the following preprocessing information:

- The random seeds $\mathcal{S}_{ID}$ and $\mathcal{S}_h$.
- $\mathrm{EID}_T(e)$ for every edge e incident to v.
- $|V(T_v)|$ and $|V(T_{v'})|$ for every T-child v' of v.
- $\mathsf{Sketch}_G(v)$, $\mathsf{Sketch}_G(V)$, $\mathsf{Sketch}_G(V(T_v))$ and $\mathsf{Sketch}_G(V(T_{v_i}))$ for every T-child v' of v.
- An edge set $\widetilde{E}(v) \subseteq E \setminus E(T)$ such that $\widetilde{T}(v) = (T \setminus \{v\}) \cup \widetilde{E}(v)$ is a spanning tree of $G \setminus \{v\}$. For each $e \in \widetilde{E}(v)$, its extended identifier $\mathrm{EID}_T(e)$ is known.

We next describe the algorithms $\mathcal{A}_y$:

▶ **Lemma 10.** *Assuming all vertices know their preprocessing information, there is an $\widetilde{O}(D)$-rounds $\widetilde{O}(1)$-congestion algorithm $\mathcal{A}_y$ that detects all cut pairs xy where $x \in V(T_y)$. The algorithm $\mathcal{A}_y$ sends messages only on edges incident to $V(T_y)$.*

Step 0: Local Computation of Component Tree for $\widetilde{T}(y)$ in y. Throughout, let $\widetilde{E} = \widetilde{E}(y)$ and $\widetilde{T} = \widetilde{T}(y)$, and denote the T-children of y by $y_1, \ldots, y_k$. This preliminary step is executed by local computation in y. It constructs the *component tree* $\widetilde{CT}$ in which every connected component of $\widetilde{T} \setminus \widetilde{E}$ is contracted into a single node. Note that $\widetilde{T} \setminus \widetilde{E} = T \setminus \{y\}$, namely the nodes in $\widetilde{CT}$ correspond to connected components of $T \setminus \{y\}$. More concretely, for every $i = 1, \ldots, k$ the component $C_i = V(T_{y_i})$ is a node of $\widetilde{CT}$, and (unless $y = s$) there is another node for the component $C_0 = V(T) \setminus V(T_y)$. Each edge (C_i, C_j) in $\widetilde{CT}$ correspond to the unique $\widetilde{E}$-edge incident to both C_i and C_j. Observe that the extended edge identifiers known to y by preprocessing contain the T-ancestry labels of all endpoints of $\widetilde{E}$-edges, as well as those of the y_i's. Using these ancestry labels, y can determine the components incident to each edge $e \in \widetilde{E}$, and therefore construct $\widetilde{CT}$.

For clarity of presentation we assume $y \neq s$; the special case $y = s$ is easier, and requires only slight modifications. We set s as the root of $\widetilde{T}$, and accordingly C_0 is the root of $\widetilde{CT}$. For each $i = 1, \ldots, k$, denote by $e_i = (r_i, p_i)$ the unique edge in $\widetilde{E}$ connecting C_i to its parent in $\widetilde{CT}$, where r_i is the endpoint of e_i inside C_i, and p_i is the endpoint lying in the parent component. See Fig. 1 in Appendix B for an illustration.

Step 1: Construction of $\widetilde{T}$. The goal of this step is for each vertex in $V(T_y) \setminus \{y\}$ to learn its parent in $\widetilde{T}$. First, y sends its children their corresponding edges from $\widetilde{E}$, namely each y_i learns $\mathrm{EID}_T(e_i)$. The y_i's then propagate (in parallel) their received edges down their T-subtrees, so that for all $i = 1, \ldots, k$, all the vertices of component C_i know $\mathrm{EID}_T(y_i)$.

Then, a BFS procedure initilized in r_i is executed inside each tree T_{y_i} (in parallel). This completes the step, since the $\widetilde{T}$-parent of each vertex in C_i is its BFS-parent from this last procedure, except for r_i whose $\widetilde{T}$-parent is p_i.

Step 2: Computing $\widetilde{T}$-Ancestry labels. In later steps, we will locally simulate Borůvka's algorithm similarly to Section A, but with the initial components being parts of $\widetilde{T}$. In order to identify which components get merged by the outgoing edges, we will need ancestry labels with respect to the tree $\widetilde{T}$ rather than T. As we are restricted to send messages only on $V(T_y)$-incident edges, we would like the T- and $\widetilde{T}$-labels to coincide for vertices in C_0 (as some of them cannot be informed of new labels). Note that the compressed paths of $v \in C_0$ w.r.t. T and $\widetilde{T}$ are generally different, even though $\pi(s,v,T) = \pi(s,v,\widetilde{T})$, as the these trees have different heavy-light notions. Hence, instead of relying solely on compressed paths in $\widetilde{T}$, we take a hybrid approach and define new labels based on breaking each $\widetilde{T}$-path to a T-part and a strictly $\widetilde{T}$-part, and compressing them accordingly. We still have the challenge of computing (at least part of) the heavy-light decomposition of $\widetilde{T}$. As the diameter of $\widetilde{T}$ might be $\Omega(\Delta D)$, we cannot use simple bottom-up or top-down computations on $\widetilde{T}$. The key for overcoming this is utilizing y as a coordinator, enabling the parts C_i to work in parallel. This yields the following claim. Missing proofs in this section appear in the full version of the paper.

▷ Claim 11. In $\widetilde{O}(D)$-rounds of computation with $\widetilde{O}(1)$ congestion, in which messages are sent only on $V(T_y)$-incident edges, one can compute $\widetilde{T}$-ancestry labels $\mathsf{ANC}_{\widetilde{T}}(\cdot)$ of $\widetilde{O}(1)$ bits, such that *every* vertex v of $\widetilde{T}$ learns $\mathsf{ANC}_{\widetilde{T}}(v)$.

Step 3: Computing Sketches w.r.t. $G \setminus \{y\}$ and $\widetilde{T}$. First, we define new extended edge identifiers for the edges of $G \setminus \{y\}$ based on the spanning tree $\widetilde{T}$. Namely, for an edge $e = (u,v)$ of $G \setminus \{y\}$, let

$$\mathrm{EID}_{\widetilde{T}}(e) = [\mathrm{UID}(e), \mathrm{ID}(u), \mathrm{ID}(v), \mathsf{ANC}_{\widetilde{T}}(u), \mathsf{ANC}_{\widetilde{T}}(v)].$$

Now, for every vertex $v \in V \setminus \{y\}$ we define its sketch $\mathsf{Sketch}^{\widetilde{T}}_{G\setminus\{y\}}(v)$ similarly to $\mathsf{Sketch}_G(v)$, only ignoring edges incident to y in the sampling, and using the $\mathrm{EID}_{\widetilde{T}}$ identifiers for the edges. By this point of the algorithm, computing these new sketches requires $\widetilde{O}(1)$ rounds of communication, in which every $v \in C_1 \cup \cdots \cup C_k$ sends $\mathsf{ANC}_{\widetilde{T}}(v)$ to all its $(G \setminus \{y\})$-neighbors. As the T- and $\widetilde{T}$-ancestry labels coincide on the vertices of C_0, every vertex $v \in V \setminus \{y\}$ can now determine $\mathrm{EID}_{\widetilde{T}}(e)$ for every edge e incident to it in $G \setminus \{y\}$, and use the random seed $\mathcal{S}_h$ to compute $\mathsf{Sketch}^{\widetilde{T}}_{G\setminus\{y\}}(v)$.

3.1: Computing $\widetilde{T}$-Subtree Sketches. Our next goal is for every $x \in C_1 \cup \cdots \cup C_k$ to learn the $(G \setminus \{y\})$-sketch of its $\widetilde{T}$-subtree (not T-subtree), namely $\mathsf{Sketch}^{\widetilde{T}}_{G\setminus\{y\}}(V(\widetilde{T}_x)) = \oplus_{v \in \widetilde{T}_x} \mathsf{Sketch}^{\widetilde{T}}_{G\setminus\{y\}}(v)$. This is done by using y as a coordinator similarly to the $\widetilde{T}$-subtree sum computation of Step 2.1. We start by bottom-up XOR-aggregation of the sketches on each T_{y_i} (in parallel), which produces the component sketches $\mathsf{Sketch}^{\widetilde{T}}_{G\setminus\{y\}}(C_i)$. Next, within $\widetilde{O}(1)$ rounds, the component sketches are all passed to y from its children. Observe that now y can locally compute the $\widetilde{T}$-subtree sketch of each r_i as follows: $\mathsf{Sketch}^{\widetilde{T}}_{G\setminus\{y\}}(V(\widetilde{T}_{r_i})) = \oplus_{j \in J(i)} \mathsf{Sketch}^{\widetilde{T}}_{G\setminus\{y\}}(C_j)$ where $J(i)$ is the set of all indices j such that C_j is the subtree of C_i in the component tree $\widetilde{CT}$. Then y sends each of its children

y_i the $\widetilde{T}$-subtree sketch of r_i, and this information is then propagated down on each T_{y_i} (in parallel), so that each r_i learns its $\widetilde{T}$-subtree sketch. The r_i's then send their $\widetilde{T}$-subtree sketches to their $\widetilde{T}$-parent, which are the p_i's. For each vertex v of $\widetilde{T}$, let

$$\beta_v = \begin{cases} \text{if } v = p_j\text{:} & \mathsf{Sketch}^{\widetilde{T}}_{G\setminus\{y\}}(v) + \mathsf{Sketch}^{\widetilde{T}}_{G\setminus\{y\}}(V(\widetilde{T}_{r_j})) \\ \text{otherwise:} & \mathsf{Sketch}^{\widetilde{T}}_{G\setminus\{y\}}(v) \end{cases}$$

Then by this point of the algorithm, every $v \in C_1 \cup \cdots \cup C_k$ know its β_v value. For $i = 1, \ldots, k$, let $\widetilde{T}^{(i)}$ be the tree induced on C_i by $\widetilde{T}$, where the parents in $\widetilde{T}^{(i)}$ are the same as in $\widetilde{T}$. Equivalently, $\widetilde{T}^{(i)}$ is the tree obtained by rerooting T_{y_i} at the vertex r_i. Each of its leaves is either an original $\widetilde{T}$-leaf or a p_j vertex for some j. The crux is that for each $x \in C_i$ it holds that $\mathsf{Sketch}^{\widetilde{T}}_{G\setminus\{y\}}(x) = \oplus_{v \in \widetilde{T}^{(i)}_x} \beta_v$. That is, the $\widetilde{T}$-subtree sketch of x is equal to the sum-of-β's in its $\widetilde{T}^{(i)}$-subtree. Hence, we complete the computation in this step by executing bottom-up XOR-aggregation of the β_v values in each of the trees $\widetilde{T}^{(i)}$ in parallel.

3.2: Computing the Sketch of $V \setminus \{y\}$. The last required sketch ingredient for the local simulation of Borůvka in the next step is letting all vertices $x \in C_1 \cup \cdots \cup C_k$ to learn the global sum-of-sketches in $G \setminus \{y\}$, i.e. $\mathsf{Sketch}^{\widetilde{T}}_{G\setminus\{y\}}(V \setminus \{y\})$. To this end, we carefully examine the contribution of the vertices in C_0 to this sum, as some of them are not $V(T_y)$-adjacent and cannot participate in the computation. This enables us to transform the global sketch $\mathsf{Sketch}_G(V)$ (known from preprocessing) to the desired global sketch in $G \setminus \{y\}$. The details appear in the full version of the paper, in the proof of the following claim:

▷ Claim 12. In $\widetilde{O}(D)$-rounds of computation with $\widetilde{O}(1)$ congestion, in which messages are sent only on $V(T_y)$-incident messages, each vertex $x \in C_1 \cup \cdots \cup C_k$ can learn $\mathsf{Sketch}^{\widetilde{T}}_{G\setminus\{y\}}(V \setminus \{y\})$.

Step 4: Local Borůvka Simulation In $G \setminus \{x, y\}$. This entire step is executed by local computation in which each $x \in C_1 \cup \cdots \cup C_k$ determines whether it is a cut vertex in $G \setminus \{y\}$, or equivalently if xy is a cut pair in G. This is done by locally simulating Borůvka's algorithms using the sketches of the components of $\widetilde{T} \setminus \{x\}$ (which are known to x by Step 3) in an identical manner to the last step of the (single) cut vertex detection algorithm of Section A, replacing G and T there with $G \setminus \{y\}$ and $\widetilde{T}$. We note that the new ancestry labels, extended identifiers and sketches, *computed with respect to* $\widetilde{T}$, are important for this simulation to follow through exactly as in Section A. This completes the proof of Lemma 10.

We conclude this section by describing the scheduling of the algorithms $\{A_y\}_{y \in V}$:

▶ **Lemma 13.** *The collection of algorithms $\{\mathcal{A}_y\}_{y \in V}$ can be executed simultaneously within $\widetilde{O}(D)$ rounds, w.h.p.*

Proof. The key observation is that every edge e participates in $O(D)$ algorithms. Specifically, since each algorithm $\mathcal{A}_y$ exchanges messages only on edges incident to $V(T_y)$, we get that the algorithms using $e = (u, v)$ are exactly $\{\mathcal{A}_y \mid y \in \pi(s, u, T) \cup \pi(s, v, T)\}$. Therefore, the total number of messages sent through $e = (u, v)$ in the collection of n algorithms $\{A_y\}_{y \in V}$ is at most $\widetilde{O}(1) \cdot (|\pi(s, u, T)| + |\pi(s, v, T)|) = \widetilde{O}(D)$. The proof follows by employing Theorem 9 with congestion and dilation bounds of $\widetilde{O}(D)$. ◀

3 Independent Cut Pairs

We now turn to consider the case where the cut pair xy is independent, i.e., x, y have no ancestor-descendant relations. Throughout this section, for every vertex $x \in V$, let $V_x = V(T_x) \setminus \{x\}$. Recall that we assume that there is no single cut vertex in the graph. Our algorithm is based on the introduced notion of *x-connectivity trees*, $\widehat{T}_x$, computed locally at each vertex x. Let $\mathcal{C}_x = \{C_1, \ldots, C_k\}$ denote the maximal connected components in the induced graph $G[V_x]$. For each $C \in \mathcal{C}_x$, the tree $\widehat{T}_x$ contains a path $\pi_x(s, C) = \pi(s, u_C) \circ (u_C, v_C)$, where (u_C, v_C) is a G-edge such that $v_C \in C$, and $x \notin \pi_x(s, C)$. Therefore, $\widehat{T}_x$ encodes the connectivity of s to V_x in the graph $G \setminus \{x\}$. We next describe the computation of these $\widehat{T}_x$ trees, and later on show how they guide the identification of independent cut pairs. Throughout, we assume that the ID of each vertex v contains also its compressed-path information $\pi^*(s, v)$. For every $v \in V_x$, let $C_{x,v}$ denote the component containing v in $\mathcal{C}_x$. When $v = x_h$, we let $H_x = C_{x,x_h}$ and denote it as the *heavy component* of x.

3.1 Computing x-Connectivity Trees

The computation has two main steps, both are based on the bottom-up aggregation of certain graph sketches over the BFS tree T. The purpose of first step is to allow every $x \in V$ to determine the connected components $\mathcal{C}_x$ in $G[V_x]$ where each such component C is identified by the vertex of largest ID among all the T-children of x in C. In addition, in the output of this step each vertex $u \in V_x$ learns the ID of its component $C_{x,u} \in \mathcal{C}_x$. The second step aggregates a special form of graph sketches that provide x with the required path information in order locally compute $\widehat{T}_x$.

Step 1: Computing Connectivity in $G[V_x]$. For ease of notation, let $D = \text{depth}(T)$ and $d_x = \text{depth}(x)$ denote the depth of x in T. We say that an edge $e = (u, v)$ has depth d if $\text{depth}(\text{LCA}(u, v)) = d$. To locally simulate the connectivity Borůvka algorithm in $G[V_x]$ at every x, it is required for x to learn $\mathsf{Sketch}_{G[V_x]}(V(T_w))$ for each T-child w of x. Observe that the edges of $G[V_x]$ can be identified as G-edges in $V_x \times (V \setminus \{x\})$ of depth at least d_x. For this purpose, the algorithm is based on aggregating the information of D graph sketches, for every depth $d \in \{1, \ldots, D\}$. The computation of d^{th} sketch $\mathsf{Sketch}^d_G(\cdot)$ will be restricted to sampling only edges of depth *at least* d. We thus obtain the following lemma. Missing proofs in this section appear in the full version of the paper.

▶ **Lemma 14.** *There is a randomized $\widetilde{O}(D)$-round algorithm that w.h.p. computes connectivity in each $G[V_x]$ for every $x \in V$ simultaneously. At the end of the execution, each u holds a component-ID in the graph $G[V_x]$ for every $x \in \pi(s, u)$. Moreover, within additional $\widetilde{O}(D)$ rounds, each u can send its entire component-ID information (for every $x \in \pi(s, u)$) to all its neighbors.*

Step 2: Computing x-Connectivity Trees $\widehat{T}_x$ via Path-Sketches. Our next goal is to provide each vertex x with the path information $\pi_x(s, C)$, for every component $C \in \mathcal{C}_x$. Such a path connects a vertex $v_C \in C$ to the source s in $G \setminus \{x\}$. As we assume that x is not a cut vertex, such a path indeed exists. Towards that goal, we augment the identifier of each edge (u, v) with the tree paths $\pi(s, u), \pi(s, v)$. Formally,

$$\text{EID}^P_T(e) = [\text{UID}(e), \text{ID}(u), \text{ID}(v), \mathsf{ANC}_T(u), \mathsf{ANC}_T(v), \pi(s, u), \pi(s, v)] \ . \tag{2}$$

In contrast to the extended-ID of Eq. (1) which have $\widetilde{O}(1)$ bits, the latter $\text{EID}^P_T(e)$ identifiers have $\widetilde{O}(D)$ bits. The sketches obtained with these $\text{EID}^P_T(e)$ IDs are called *path-sketches*, denoted as $\mathsf{Sketch}^P_G(S)$ for $S \subseteq V$. The advantage of these path-sketches is that any detected

outgoing edge (u, v) obtained from $\mathsf{Sketch}^P_G(Q)$ includes the path information $\pi(s, u)$ and $\pi(s, v)$. Note that the path-sketches $\mathsf{Sketch}^P_G(S)$ have $\widetilde{O}(D)$ bits, since the edge IDs have now $\widetilde{O}(D)$ bits.

Our goal is to let each x learn the path-sketches $\mathsf{Sketch}^P_G(C)$ for each component $C \in \mathcal{C}_x$. Since each path-sketch has $\widetilde{O}(D)$ bits, we cannot allow to compute D sketches for each depth $d \in \{1, \ldots, D\}$. Instead we only aggregate the $\mathsf{Sketch}^P_G(u)$ information in a bottom-up manner on T, which allows every vertex x to learn $\mathsf{Sketch}^P_G(V(T_w))$ for each of its T-children w. By combining with the output of the first step, x can then determine $\mathsf{Sketch}^P_G(C)$ for every $C \in \mathcal{C}_x$.

▶ **Lemma 15.** *W.h.p., all vertices x can compute the x-connectivity trees $\widehat{T}_x$ within $\widetilde{O}(D)$ randomized rounds.*

For each $x \in V$ and $C \in \mathcal{C}_x$, we define the compressed path of $\pi_x(s, C)$ as $\pi^*_x(s, C) = \pi^*(s, v_C) \circ (v_C, u_C)$ (hence, $\pi^*_x(s, C)$ has $\widetilde{O}(1)$ bits). We conclude the computation regarding the connectivity trees by letting each vertex v learn the compressed-path $\pi^*_x(s, C_{x,v})$ for each of its ancestors $x \in \pi(s, v)$. Since the compressed-path has $\widetilde{O}(1)$ bits, a vertex is required to receive $\widetilde{O}(D)$ bits of information, which can be done in $\widetilde{O}(D)$ rounds:

▶ **Lemma 16.** *There is an $\widetilde{O}(D)$-round algorithm that allows each vertex v to learn the compressed path $\pi^*_x(s, C_{x,v})$ for each $x \in \pi(s, v)$, as well as the entire path $\pi_x(s, C_{x,v})$ for each $x \in \mathsf{LA}(v)$. In addition, each vertex v can share all of this information with neighbors.*

Proof. We let every vertex x send the full path $\pi_x(s, C_{x,x'})$ to each light child x' of x, and the compressed path $\pi^*_x(s, H_x)$ to its heavy child x_h. This information is propagated towards the leaf vertices of T_x. Since each vertex is required to receive $\widetilde{O}(D)$ bits of information from each of its *light* ancestors, as well as $\widetilde{O}(1)$ bits from each of its heavy ancestors, overall it is required to receive $\widetilde{O}(D)$ bits. This can be done in $\widetilde{O}(D)$ rounds, by standard pipeline techniques. Since each v learns $\widetilde{O}(D)$ bits of information, the learned information can be exchanged between every pair of neighbors within $\widetilde{O}(D)$ rounds, as well. ◀

3.2 Component Classification Based on Sensitivity

We next use the structure of the x-connectivity tree $\widehat{T}_x$ to classify the xy pairs into several types. We also filter-out possibly many irrelevant xy pairs (for which we deduce immediately that xy is not a cut) using the notion of *sensitivity*.

▶ **Definition 17** (Sensitivity Notions of $\mathcal{C}_x$ Components)**.** *Fix an independent pair x, y. A component $C \in \mathcal{C}_x$ is* y-sensitive *if $y \in \pi_x(s, C)$. The y-sensitive components of $\mathcal{C}_x$ are further classified into two types:* pseudo-sensitive *and* fully-sensitive*, as follows. A component $C \in \mathcal{C}_x$ is* pseudo y-sensitive *if the tree path $\pi_x(s, C)$ contains some edge (y, y') such that $x \notin \pi_y(s, C_{y,y'})$, where $C_{y,y'}$ is the component containing y' in $\mathcal{C}_y$. Finally, a y-sensitive component $C \in \mathcal{C}_x$ is* fully y-sensitive *if C is not pseudo-sensitive.*

Hence, in particular a component $C \in \mathcal{C}_x$ is fully y-sensitive if either that last edge of $\pi_x(s, C)$ is incident to y, or that there is an edge $(y, y') \in \pi_x(s, C)$ such that the component $C_{y,y'} \in \mathcal{C}_y$ is x-sensitive. Note that non-y-sensitive components are clearly connected to s in $G \setminus \{x, y\}$. We later on show that this is true also for pseudo y-sensitive components, therefore their sensitivity to y is superficial. Let $\mathcal{S}(x, y), \mathcal{PS}(x, y), \mathcal{FS}(x, y)$ denote the components in $\mathcal{C}_x$ that are y-sensitive, pseudo y-sensitive and fully y-sensitive, respectively[7]. We next

[7] Notice that these notations are not symmetric in x, y, e.g. $\mathcal{S}(x, y)$ is different than $\mathcal{S}(y, x)$.

show that each vertex x can determine, for every $C \in \mathcal{C}_x$, certain y vertices for which C is fully y-sensitive by running the procedure described of the following lemma. Note that by having the compressed-path $\pi^*_y(s, C_{y,y'})$ and the $\pi^*(s,x)$, it is possible to determine if $x \in \pi_y(s, C_{y,y'})$, hence determining if $C_{y,y'}$ is x-sensitive.

▶ **Lemma 18.** *There is an $\widetilde{O}(D)$-round algorithm that computes the following for every $x \in V$ (in parallel):*

- *$\pi^*_y(s, C_{y,y'})$ for every edge $(y,y') \in \pi_x(s,C)$ and every $C \in \mathcal{C}_x \setminus \{H_x\}$.*
- *$\pi^*_y(s, C_{y,y'})$ for every* light *edge $(y,y') \in \pi_x(s, H_x)$.*

3.3 xy-Connectivity Algorithms Under a Promise

Throughout, we assume that all vertices applied the pre-processing steps of computing the x-connectivity trees $\widehat{T}_x$, as well as, applied the $\widetilde{O}(D)$-round procedures of Lemma 16 and 18. From this point on, we explain how to determine the connectivity in $G \setminus \{x,y\}$, first for a single pair xy, and then for all pairs that satisfy a given promise.

Recall that $\mathsf{LD}(x)$ is the collection of light descendants of x in T. For a vertex y, let $\mathsf{LDS}(x,y)$ be the collection of light descendants of x that are sensitive to y. Formally, the **l**ight x-**d**escendants y-**s**ensitive vertices are defined by:

$$\mathsf{LDS}(x,y) = \{v \in \mathsf{LD}(x) \;\mid\; y \in \pi_x(s, C_{x,v}) \setminus V(T_x)\} \;. \tag{3}$$

▶ **Observation 19.** *Every vertex v belongs to a total of $O(D \log n)$ sets $\mathsf{LDS}(x,y)$ for $x,y \in V$.*

Proof. A vertex $v \in V$ has $O(\log n)$ light ancestors (i.e., belongs to $O(\log n)$ sets of $\mathsf{LD}(x)$). In addition, for each light ancestor $x \in \pi(s,v)$, there are $O(D)$ vertices $y \in \pi_x(s, C_{x,v})$. Therefore, it belongs to $O(D \log n)$ sets as required. ◀

▶ **Theorem 20** (xy-Connectivity Given an x-y Path)**.** *Fix $x,y \in V$ and assume that there is an x-y path $\Pi_{x,y} \subseteq G$ (known in a distributed manner) of length $O(D)$. Then, there is an xy-connectivity algorithm $\mathcal{A}^P_{x,y}$ (i.e., that determines the connectivity in $G \setminus \{x,y\}$) in $\widetilde{O}(D)$ and $\widetilde{O}(1)$-congestion, by sending messages only along on the edges of $\Pi_{x,y}$ or edges incident to $\mathsf{LDS}(x,y) \cup \mathsf{LDS}(y,x)$. At the end of the computation, both x and y know whether $G \setminus \{x,y\}$ is connected or not.*

Before proving Theorem 20, we show that given a set of pairs $Q \subseteq V \times V$, then all algorithms $\{\mathcal{A}^P_{x,y} \mid (x,y) \in Q\}$ can be scheduled simultaneously when provided a path collection $\mathcal{P}_Q = \{\Pi_{x,y} \mid (x,y) \in Q\}$ that satisfies the following promise:

[**Promise**:] *$\mathcal{P}_Q$-paths have length $O(D)$, and each edge appears on $\widetilde{O}(D)$ paths in $\mathcal{P}_Q$.*

By a straightforward application of the random delay approach, we obtain:

▶ **Corollary 21.** *[All Pairs xy-Connectivity Under a Promise] Let $Q \subseteq V \times V$ be a collection of independent pairs and let $\mathcal{P}_Q = \{P_{x,y} \mid (x,y) \in Q\}$ be a collection of x-y paths that satisfy the promise. Then, the collection of algorithms $\{\mathcal{A}^P_{x,y} \mid (x,y) \in Q\}$, where each $\mathcal{A}_{x,y}$ uses the corresponding path $\Pi_{x,y} \in \mathcal{P}_Q$, can be scheduled simultaneously within $\widetilde{O}(D)$ rounds, w.h.p.*

Description of the Connectivity Algorithm $\mathcal{A}^P_{x,y}$. The algorithm is based on simulating the Borůvka algorithm using the sketch information of connected subsets in $G \setminus \{x,y\}$, held jointly by x and y. Throughout, we refer to the given x-y path $\Pi_{x,y}$ as *the xy channel*. Recall that the algorithm can send only $\widetilde{O}(1)$ bits on that channel. The input for the $i \geq 1$

phase of Borůvka is the following. There is a partitioning $\mathcal{P}_{i-1} = \{P_{i-1,1}, \ldots, P_{i-1,k_{i-1}}\}$ of the vertices in $V \setminus \{x, y\}$ into connected subsets (in $G \setminus \{x, y\}$). We call each $P \in \mathcal{P}_{i-1}$ a *part* (to avoid confusion with the term "component" reserved for sets in $\mathcal{C}_x$ and $\mathcal{C}_y$). We mark a special vertex in each $P_{i,j} \in \mathcal{P}_i$, called the *leader* of the part. The source vertex s is the leader of its own part (called the s-part), and the leaders of the other parts are some chosen T-children of x or y in these parts. The part-ID is the ID of its leader. The part containing x_h (resp., y_h) is called *x-heavy* (resp., *y-heavy*)[8]. The parts that are free of s, x_h, y_h are called *light*. Hence every light part is contained in $\mathsf{LD}(x) \cup \mathsf{LD}(y)$. A part P is denoted as *growable* if there is an outgoing G-edge connecting P to $V \setminus (P \cup \{x, y\})$. The Borůvka algorithm has $K = O(\log n)$ forest growing phases in $G \setminus \{x, y\}$, each phase reduces the number of growable parts by a constant factor, in expectation. We maintain the following invariant for the beginning of each phase $i \in \{1, \ldots, K\}$:

(I1) x, y know $\mathsf{Sketch}_{G\setminus\{x,y\}}(P)$ of the part $P \in \mathcal{P}_{i-1}$ containing s.
(I2) $z \in \{x, y\}$ knows $\mathsf{Sketch}_{G\setminus\{x,y\}}(P)$ for every *light* part $P \in \mathcal{P}_{i-1}$ whose leader is in T_z.
(I3) x, y know $\mathsf{Sketch}_{G\setminus\{x,y\}}(P)$ as well as the part-IDs of the heavy parts P in $\mathcal{P}_{i-1}$.
(I4) $z \in \{x, y\}$ knows, for each T-child z' of z, the part-ID of the part containing z' in $\mathcal{P}_{i-1}$.

Satisfying the Invariant for the First Borůvka Phase. We start by defining the partitioning $\mathcal{P}_0$ and in particular, focus first on the definition of the part containing s. Recall Def. 17 and that $\mathcal{S}(x, y), \mathcal{PS}(x, y), \mathcal{FS}(x, y) \subseteq \mathcal{C}_x$ are the y-sensitive, pseudo y-sensitive and fully y-sensitive components, respectively. Let $\mathsf{NS}(x, y) = \bigcup_{C \in \mathcal{C}_x \setminus \mathcal{FS}(x,y)} C$. The set $\mathsf{NS}(y, x)$ is defined in an analogous manner. Then the s-part in $\mathcal{P}_0$ is given by $U(x, y) = (V \setminus (V(T_x) \cup (T_y))) \cup \mathsf{NS}(x, y) \cup \mathsf{NS}(y, x)$. The next observation exploits the fact that the pseudo y-sensitive components in $\mathcal{C}_x$ and the pseudo x-sensitive components in $\mathcal{C}_y$ are all connected to s in $G \setminus \{x, y\}$.

▶ **Observation 22.** *$G[U(x, y)]$ is connected.*

We partition the responsibilities on the parts in $\mathcal{P}_0$ between x and y, as follows. Let $\mathcal{P}_{0,x} = \mathcal{FS}(x, y)$ be the components in $\mathcal{C}_x$ that are fully-sensitive to y. Similarly, $\mathcal{P}_{0,y} = \mathcal{FS}(y, x)$. The 0^{th} partitioning of $V \setminus \{x, y\}$ is given by $\mathcal{P}_0 = \{U(x, y)\} \cup \mathcal{P}_{0,x} \cup \mathcal{P}_{0,y}$. For every $z \in \{x, y\}$, the leader of each $C \in \mathcal{P}_{0,z}$ is chosen as the vertex of largest ID among all the T-children of x, y in C. The leader of $U(x, y)$ is the root s. To satisfy the invariants for the beginning of phase $i \geq 1$, it is sufficient to show the following claims for x (as they apply in a symmetric manner also for y):

▷ Claim 23. Within $\widetilde{O}(D)$ rounds, the vertex x can compute $\mathsf{Sketch}_{G\setminus\{x,y\}}(C)$ for every component $C \in \mathcal{S}(x, y)$. In addition, the vertex y can determine its neighbors in $\{v \in V_x \mid y \notin \pi_x(s, C_{x,v})\}$. The communication is restricted to the edges of $\mathsf{LDS}(x, y) \cup \mathsf{LDS}(y, x)$ and using the xy channel.

▷ Claim 24. By exchanging $\widetilde{O}(1)$ bits of information (using the promised channel), invariants (I1-I4) hold w.r.t $\mathcal{P}_0$.

Simulation of the i^{th} Borůvka Phase. We now describe the execution of phase $i \geq 1$ assuming that at the beginning of the phase the invariant holds w.r.t $\mathcal{P}_{i-1}$. The output of the execution will be the partitioning $\mathcal{P}_i$, for which we later show that the invariant holds

[8] A part can be both x-heavy and y-heavy.

as well. Our goal is to let x, y simulate a Borůvka phase in which parts of $\mathcal{P}_{i-1}$ are merged along their outgoing edges. See Figure 2 in Appendix B for an illustration of this process. The main objective of this phase is to reduce the number of *growable* parts by a constant factor, in expectation. Throughout, we use the following auxiliary claim which allows the vertices in every light part to exchange $\widetilde{O}(1)$ bits, in parallel.

▷ Claim 25. Let P be a light part in $\mathcal{P}_{i-1}$ such that each vertex $v \in P$ holds a $\widetilde{O}(1)$-bit value $val(v)$. Then, there is an $\widetilde{O}(D)$-round algorithm that allows all vertices in P to compute any aggregate function of the $val(v)$ values for $v \in P$, by sending messages *only* along edges incident to P. Consequently, all light parts in $\mathcal{P}_{i-1,x} \cup \mathcal{P}_{i-1,y}$ can compute their respective aggregate functions, in parallel.

For efficiency of computation, we restrict the merge shapes to be star shapes by using random coins (see e.g., [15]). Such star merges are obtained by letting each part $\mathcal{P}_{i-1}$ toss a random coin, and allowing only merges centered on head-parts, each accepting incoming suggested merge-edges from tail-parts. The leader of this head-part becomes the leader of the merged part. We show that under the promise and the $(i-1)^{th}$ invariant, this merging phase can be implemented in $\widetilde{O}(D)$ rounds as follows. W.l.o.g., we make x be responsible for the s-part $P_s \in \mathcal{P}_{i-1}$.

Implementing Merges. Each vertex $z \in \{x, y\}$ tosses a (fresh) random coin for each of its parts in $\mathcal{P}_{i-1,z}$. In addition, x tosses a coin for the s-part P_s. Next, for each of the tail part $P \in \mathcal{P}_{i-1,z}$, z locally computes an outgoing edge for each of its *tail* parts in $\mathcal{P}_{i-1,z}$. In addition, x computes an outgoing edge for the s-part (in case that the coin toss of that part is tail). For each growable part $P \in \mathcal{P}_{i-1,z}$, such an edge can be detected from $\mathsf{Sketch}_{G\setminus\{x,y\}}(P)$ with constant probability. The parts of $\mathcal{P}_i$ are formed by merging every head part $P^* \in \mathcal{P}_{i-1}$ with all the tail parts in $\mathcal{P}_{i-1}$ whose outgoing edges point at P^*. The leader of the merged part is the leader of the head part P^*. For every tail part $P \in \mathcal{P}_{i-1,x}$, let $e_P = (u_P, v_P)$ be the detected outgoing edge obtained by x from $\mathsf{Sketch}_{G\setminus\{x,y\}}(P)$.

▷ Claim 26. Using $\widetilde{O}(D)$ rounds of communication over edges incident to $\mathsf{LDS}(x, y)$ and the given xy channel, z can determine for all its tail parts $P \in \mathcal{P}_{i-1,z}$ with an outgoing edge $e_P = (u_P, v_P)$, the following information: (i) the part-ID of the second endpoint $v_P \notin P$ and, (ii) the coin-toss of the part of v_P.

To implement the merges and satisfy the invariant, it is required for $z \in \{x, y\}$ to learn the updated sketch information of their head parts in $\mathcal{P}_{i-1,z}$. We next explain how y can compute the sketch information of each of its head parts $P^* \in \mathcal{P}_{i-1,y}$. (A similar procedure would work for x).

By Claim 26, x knows for every head part $P^* \in \mathcal{P}_{i-1,y}$, the collection of tail parts in $\mathcal{P}_{i-1,x}$ that should be merged with P^*.

Any → Non-Light Merges. There are (at most three) non-light parts in $\mathcal{P}_{i-1}$, corresponding to at most two heavy parts and the s-part[9]. For each of these non-light part P^*, x aggregates that sketch information of the corresponding tail parts $P \in \mathcal{P}_{i-1,x}$, and send it to y over the xy channel.

From the point on, x considers the transfer of information concerning the light head parts P^* in $\mathcal{P}_{i-1,y}$.

[9] The latter is held by x, so when revering the roles of x, y, y might be required to send x the sum of sketch information of the tail parts in $\mathcal{P}_{i-1,y}$ that got merged with the s-part.

Non-Light → Light Merges. It uses the xy channel to send y the sketch information of its non-light tail parts P, along with the part-ID of their head parts (to which they should be merged).

Light → Light Merges. The sketch of all other (light) parts in $\mathcal{P}_{i-1,x}$ are communicated to y over the edges incident to the light sensitive xy descendants, $\mathsf{LDS}(x,y) \cup \mathsf{LDS}(y,x)$, as follows. Using Claim 25, each light and tail part $P \in \mathcal{P}_{i-1,x}$ can learn $\mathsf{Sketch}_{G\setminus\{x,y\}}(P)$ (as x holds this information, by the invariant). Note that by definition $P, P^* \subseteq \mathsf{LDS}(x,y) \cup \mathsf{LDS}(y,x)$. The vertices of P then send this received information to all their neighbors. At this point, for every light head part P^* in $\mathcal{P}_{i-1,y}$, and for every tail *light* part P in $\mathcal{P}_{i-1,x}$, there is a vertex $v_P \in P^*$ that holds $\mathsf{Sketch}_{G\setminus\{x,y\}}(P)$. By applying Claim 25, all vertices in P^* can learn the sum of all these sketches. This provides y with all the required information from x to compute $\mathsf{Sketch}_{G\setminus\{x,y\}}(P^*)$ for each head part $P^* \in \mathcal{P}_{i-1,y}$. In a symmetric manner, x can compute the sketch of the merged parts for all its head parts in $\mathcal{P}_{i-1,x}$. Using the xy channel, x and y can exchange the part-ID and sketch information of the heavy parts and the s-part in $\mathcal{P}_i$. This satisfies (I1,I2,I3) for the partitioning $\mathcal{P}_i$.

To satisfy (I4), note that the part-ID has changed only for tail parts in $\mathcal{P}_{i-1}$. For the tail-parts in $\mathcal{P}_{i-1,z}$, z holds their new part-ID using Claim 26 (i.e., this is the part-ID of the detected outgoing edges). This completes the description for phase i.

We are now ready to complete the proof of Theorem 20.

Proof of Theorem 20. By the description of the i^{th} phase, the invariant holds w.r.t $\mathcal{P}_i$. We next show that the i^{th} phase sends $\widetilde{O}(1)$ messages along edges incident to $\mathsf{LDS}(x,y) \cup \mathsf{LDS}(y,x)$, as well as over the xy channel. It is also easy to see that given the promised channel that running time is $\widetilde{O}(D)$ using Claim 25. Finally, we show that within $k = O(\log n)$ phases it holds that there are no growable components in $\mathcal{P}_k$.

Recall that a part P in $\mathcal{P}_j$ is denoted as *growable* if there is a G-edge $(u,v) \in P \times (V \setminus (\{x,y\} \cup P))$. We claim that the number of growable part reduces by a constant factor in each Borůvka phase. Given a sketch information $\mathsf{Sketch}_{G\setminus\{x,y\}}(P)$ for a growable part P, one can infer an outgoing edge (u,v) from P with constant probability. In addition, with probability $1/4$ this edge is valid (i.e., P is a tail part and v is in a head part). Therefore, overall the number of growable parts reduces by a constant factor, in expectation. By the Markov inequality, w.h.p. there is no growable part after $O(\log n)$ phases. Since x, y jointly hold the sketch information of all parts in $\mathcal{P}_k$ they can determine if there is more than one part in $\mathcal{P}_k$ by exchanging information along their channel (i.e., if $G \setminus \{x,y\}$ is not connected, then w.h.p. either x or y holds a part whose leader is not s). The theorem follows. ◀

References

1 Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Analyzing graph structure via linear measurements. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 459–467. SIAM, 2012.

2 Giuseppe Di Battista and Roberto Tamassia. Incremental planarity testing (extended abstract). In *30th Annual Symposium on Foundations of Computer Science, Research Triangle Park, North Carolina, USA, 30 October – 1 November 1989*, pages 436–441. IEEE Computer Society, 1989.

3 Giuseppe Di Battista and Roberto Tamassia. On-line maintenance of triconnected components with spqr-trees. *Algorithmica*, 15(4):302–318, 1996.

4 Keren Censor-Hillel, Mohsen Ghaffari, and Fabian Kuhn. Distributed connectivity decomposition. In *Proceedings of the 2014 ACM symposium on Principles of distributed computing*, pages 156–165. ACM, 2014.

5 Li Chen, Rasmus Kyng, Yang P. Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. *CoRR*, abs/2203.00671, 2022. doi:10.48550/arXiv.2203.00671.

6 Keerti Choudhary. An optimal dual fault tolerant reachability oracle. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, volume 55 of *LIPIcs*, pages 130:1–130:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016.

7 Michal Dory, Yuval Efron, Sagnik Mukhopadhyay, and Danupon Nanongkai. Distributed weighted min-cut in nearly-optimal time. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 1144–1153. ACM, 2021.

8 Michal Dory and Merav Parter. Fault-tolerant labeling and compact routing schemes. In Avery Miller, Keren Censor-Hillel, and Janne H. Korhonen, editors, *PODC '21: ACM Symposium on Principles of Distributed Computing, Virtual Event, Italy, July 26-30, 2021*, pages 445–455. ACM, 2021.

9 Ran Duan and Seth Pettie. Dual-failure distance and connectivity oracles. In *Proceedings of the twentieth annual ACM-SIAM symposium on Discrete algorithms*, pages 506–515. SIAM, 2009.

10 Ran Duan and Seth Pettie. Connectivity oracles for graphs subject to vertex failures. *CoRR*, abs/1607.06865, 2016. arXiv:1607.06865.

11 Ran Duan and Seth Pettie. Connectivity oracles for graphs subject to vertex failures. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 490–509, 2017.

12 Pawel Gawrychowski, Shay Mozes, and Oren Weimann. Minimum cut in o(m $\log^2$ n) time. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPIcs*, pages 57:1–57:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020.

13 Loukas Georgiadis, Giuseppe F. Italiano, Luigi Laura, and Nikos Parotsidis. 2-vertex connectivity in directed graphs. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *Automata, Languages, and Programming – 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, volume 9134 of *Lecture Notes in Computer Science*, pages 605–616. Springer, 2015.

14 Mohsen Ghaffari. Near-optimal scheduling of distributed algorithms. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC*, pages 3–12, 2015.

15 Mohsen Ghaffari and Bernhard Haeupler. Distributed algorithms for planar networks II: low-congestion shortcuts, MST, and min-cut. In Robert Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 202–219. SIAM, 2016.

16 Mohsen Ghaffari, Krzysztof Nowicki, and Mikkel Thorup. Faster algorithms for edge connectivity via random 2-out contractions. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 1260–1279. SIAM, 2020.

17 Mohsen Ghaffari and Merav Parter. MST in log-star rounds of congested clique. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing, PODC 2016, Chicago, IL, USA, July 25-28, 2016*, pages 19–28, 2016.

18 David Gibb, Bruce M. Kapron, Valerie King, and Nolan Thorn. Dynamic graph connectivity with improved worst case update time and sublinear space. *CoRR*, abs/1509.06464, 2015. arXiv:1509.06464.

19 Manoj Gupta and Shahbaz Khan. Multiple source dual fault tolerant BFS trees. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl, editors, *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland*, volume 80 of *LIPIcs*, pages 127:1–127:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017.

20 Zhiyang He, Jason Li, and Magnus Wahlström. Near-linear-time, optimal vertex cut sparsifiers in directed acyclic graphs. In Petra Mutzel, Rasmus Pagh, and Grzegorz Herman, editors, *29th Annual European Symposium on Algorithms, ESA 2021, September 6-8, 2021, Lisbon, Portugal (Virtual Conference)*, volume 204 of *LIPIcs*, pages 52:1–52:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.

21 John E. Hopcroft and Robert Endre Tarjan. Dividing a graph into triconnected components. *SIAM J. Comput.*, 2(3):135–158, 1973.

22 Michael Kapralov and David Woodruff. Spanners and sparsifiers in dynamic streams. In *Proceedings of the 2014 ACM symposium on Principles of distributed computing*, pages 272–281, 2014.

23 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*, pages 1131–1142. SIAM, 2013.

24 David R Karger. Random sampling in cut, flow, and network design problems. *Mathematics of Operations Research*, 24(2):383–413, 1999.

25 Karthik C. S. and Merav Parter. Deterministic replacement path covering. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10–13, 2021*, pages 704–723. SIAM, 2021.

26 Valerie King, Shay Kutten, and Mikkel Thorup. Construction and impromptu repair of an MST in a distributed network with o(m) communication. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC 2015, Donostia-San Sebastián, Spain, July 21–23, 2015*, pages 71–80, 2015.

27 Frank Thomson Leighton, Bruce M Maggs, and Satish B Rao. Packet routing and job-shop scheduling ino (congestion+ dilation) steps. *Combinatorica*, 14(2):167–186, 1994.

28 Jason Li, Danupon Nanongkai, Debmalya Panigrahi, Thatchaphol Saranurak, and Sorrachai Yingchareonthawornchai. Vertex connectivity in poly-logarithmic max-flows. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 317–329. ACM, 2021.

29 Ali Mashreghi and Valerie King. Broadcast and minimum spanning tree with o(m) messages in the asynchronous CONGEST model. In *32nd International Symposium on Distributed Computing, DISC 2018, New Orleans, LA, USA, October 15-19, 2018*, pages 37:1–37:17, 2018.

30 Danupon Nanongkai, Thatchaphol Saranurak, and Sorrachai Yingchareonthawornchai. Breaking quadratic time for small vertex connectivity and an approximation scheme. In Moses Charikar and Edith Cohen, editors, *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 241–252. ACM, 2019.

31 Jaroslav Nešetřil, Eva Milková, and Helena Nešetřilová. Otakar Borůvka on minimum spanning tree problem translation of both the 1926 papers, comments, history. *Discrete Mathematics*, 233(1):3–36, 2001.

32 Merav Parter. Dual failure resilient BFS structure. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing*, pages 481–490, 2015.

33 Merav Parter. Small cuts and connectivity certificates: A fault tolerant approach. In *33rd International Symposium on Distributed Computing*, 2019.

34 Merav Parter. Distributed constructions of dual-failure fault-tolerant distance preservers. In Hagit Attiya, editor, *34th International Symposium on Distributed Computing, DISC 2020, October 12-16, 2020, Virtual Conference*, volume 179 of *LIPIcs*, pages 21:1–21:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020.

35 David Peleg. *Distributed Computing: A Locality-sensitive Approach*. SIAM, 2000.

36 Seth Pettie and Longhui Yin. The structure of minimum vertex cuts. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPIcs*, pages 105:1–105:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.

37 David Pritchard and Ramakrishna Thurimella. Fast computation of small cuts via cycle space sampling. *ACM Transactions on Algorithms (TALG)*, 7(4):46, 2011.

38 Daniel Dominic Sleator and Robert Endre Tarjan. A data structure for dynamic trees. *J. Comput. Syst. Sci.*, 26(3):362–391, 1983.

39 Robert Endre Tarjan. Depth-first search and linear graph algorithms. *SIAM J. Comput.*, 1(2):146–160, 1972.

40 Ramakrishna Thurimella. Sub-linear distributed algorithms for sparse certificates and biconnected components. *Journal of Algorithms*, 23(1):160–179, 1997.

41 Salil P. Vadhan. Pseudorandomness. *Foundations and Trends® in Theoretical Computer Science*, 7(1–3):1–336, 2012. doi:10.1561/0400000010.

42 Oren Weimann and Raphael Yuster. Replacement paths and distance sensitivity oracles via fast matrix multiplication. *ACM Transactions on Algorithms (TALG)*, 9(2):14, 2013.

A Single Cut Vertices

In this section we describe the distributed algorithm for detecting single vertex cuts of Theorem 1. This serves both as a warm-up to our approach in the subsequent sections devoted to dual vertex cuts detection, as well as for a detailed presentation of basic tools used in these next sections. We assume each vertex v is equipped with its heavy/light classification in T and with its ancestry label which is its compressed path, $\mathsf{ANC}_T(v) = \pi^*(s, v, T)$. This can be achieved in $\widetilde{O}(D)$ rounds by Lemma 5.

Step 0: Computing Extended Edge IDs. The source s samples a random seed $\mathcal{S}_{ID}$ of $\widetilde{O}(1)$ bits and shares it with all vertices. Then, using Lemma 6, each vertex v can then locally compute the unique edge-ID $\mathrm{UID}(e)$ for each of its incident edges. By letting all neighbors in G exchange their ANC_T-labels, each $\mathrm{UID}(e)$ can be concatenated with the required information to create $\mathrm{EID}(e)$.

Step 1: Computing Subtree Sketches. The source s locally samples the random seed $\mathcal{S}_h$ of $\widetilde{O}(1)$ bits and sends it to all the vertices. Along with the extended edge IDs, this provides all the required information for the computation of $\mathsf{Sketch}_G(v)$ locally in each vertex v. By XOR-aggregation of the individual sketches from the leaves of T up to the root s, each vertex v obtains its subtree sketch, given by $\mathsf{Sketch}_G(V(T_v)) = \oplus_{u \in T_v} \mathsf{Sketch}_G(v)$. Next, within $\widetilde{O}(1)$ rounds, each vertex passes its subtree sketch to its parent, so that each vertex now holds the subtree sketch for each of its children. Finally, the source s also broadcasts its subtree sketch, which is $\mathsf{Sketch}_G(V)$, to all the other vertices.

Step 2: Local Connectivity Computation. This step is locally applied at every vertex x, and requires no additional communication. We show that each vertex x, given the received sketch information in Step 1, can locally simulate the Borůvka's algorithm [31] in the graph $G \setminus \{x\}$, and consequently determine if $G \setminus \{x\}$ is connected. Let $x_1, \ldots, x_k$ be the children of x in T. We assume that $x \neq s$; the case $x = s$ is easier and requires only slight modifications. The connected components in $T \setminus \{x\}$ are denoted by $\mathcal{C}_x = \{V(T_{x_j}) \mid j = 1, \ldots, k\} \cup \{V \setminus V(T_x)\}$. By Step 1, x holds the G-sketch of each component in $\mathcal{C}_x$: It

has explicitly received $\mathsf{Sketch}_G(V(T_{x_j}))$ from each child x_j. In addition, it can locally infer $\mathsf{Sketch}(V \setminus V(T_x)) = \mathsf{Sketch}(V) \oplus \mathsf{Sketch}(V(T_x))$. To implement Borůvka's algorithm on these components, we first need to update these G-sketches into $(G \setminus \{x\})$-sketches.

2.1: Obtaining Sketch Information in $G \setminus \{x\}$. Recall x knows the random seed $\mathcal{S}_h$ as well as the extended identifiers of its incident edges (from Step 0). For each such edge (x, u), it first uses the ancestry label of u and of its T-children (found in the EID_T's) to determine the component C of u in $\mathcal{C}_x$. It then cancel this edges from the sketch of the component C using Lemma 8. This allows x to obtain $\mathsf{Sketch}_{G\setminus\{x\}}(C)$ for every $C \in \mathcal{C}_x$.

2.2: Simulating Borůvka in $G \setminus \{x\}$. The input to this step is the $(G \setminus \{x\})$-sketch information of the components in $\mathcal{C}_{x,0} = \mathcal{C}_x$. The desired output is determining the connectivity of $G \setminus \{x\}$. The algorithm consists of $O(\log n)$ phases of the Borůvka algorithm, and is very similar to the (centralized) decoding algorithm of [8]. Each phase i will be given as input a partitioning $\mathcal{C}_{x,i} = \{C_{i,1}, \ldots, C_{i,k_i}\}$ of (not necessarily maximal) connected components in $G \setminus \{x\}$ along with their sketch information $\mathsf{Sketch}_{G\setminus\{x\}}(C_{i,j})$. The output of the phase is a coarser partitioning $\mathcal{C}_{x,i+1}$, along with the sketch information of the new parts. A component $C_{i,j} \in \mathcal{C}_{x,i}$ is said to be *growable* if it has at least one outgoing edge to a vertex in $V \setminus (C_{i,j} \cup \{x\})$. To obtain outgoings edges from the growable components in $\mathcal{C}_{x,i}$, the algorithm uses the i^{th} basic-unit sketch $\mathsf{Sketch}_{G\setminus\{x\},i}(C_{i,j})$ of each $C_{i,j} \in \mathcal{C}_{x,i}$. By Lemma 7, from every growable component $C_{i,j} \in \mathcal{C}_{x,i}$, we get one outgoing edge $e = (u, v)$ with constant probability. To find the component $C_{i,j'}$ containing the other endpoint of e (to be merged with $C_{i,j}$), we use the T-ancestry labels found in $\mathrm{EID}_T(e)$. Say this endpoint is v. We determine the component of v in $T \setminus \{x\}$, i.e. the component $C_{0,q}$ containing v in $\mathcal{C}_{x,0}$, by querying the ancestry relation between v and each child of x using $\mathsf{ANC}_T(v)$ and the labels of x's children. Then v belongs to the unique component $C_{i,j'} \in \mathcal{C}_{x,i}$ containing $C_{0,q}$. The sketch information for the next phase $i + 1$ is given by XORing over the sketches of the components in $\mathcal{C}_{x,i}$ that got merged into a single component in $\mathcal{C}_{x,i+1}$. Note that it is important to use fresh randomness (i.e., independent sketch information) in each of the Borůvka phases [1, 23, 10]. Since each growable component gets merged with constant probability, the expected number of growable components is reduced by a constant factor in each phase. Thus after $O(\log n)$ phases, the expected number of growable components is at most $1/n^5$, and by Markov's inequality we conclude that w.h.p. there are no growable components. The partitioning at this point corresponds to the maximal connected components in $G \setminus \{x\}$, so its connectivity can be inferred. This concludes the proof of Theorem 1.

Finally, we note that by tracking the merges throughout the Borůvka simulation, x can also find a subset $\widetilde{E}$ of the outgoing edges received throughout the simulation such $(T \setminus \{x\}) \cup \widetilde{E}$ is a maximal spanning forest of $G \setminus \{x\}$.

B Figures

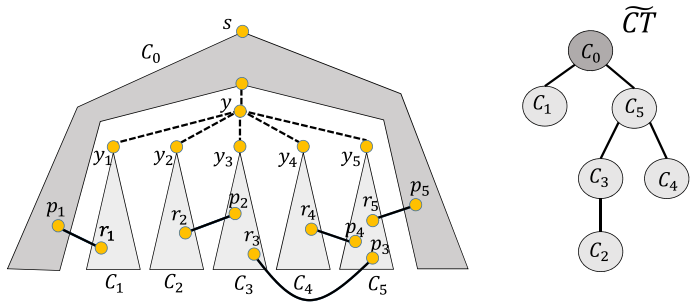

Figure 1 Left: Illustration of the trees T and $\widetilde{T}$. The dashed edges are T-edges adjacent to y, and the solid edges are $\widetilde{E}$-edges. The components $C_0, C_1, \ldots, C_5$ are each internally connected via original T-edges. The tree $\widetilde{T}$ is obtained by removing y and its incident edges from the T and adding the $\widetilde{E}$ edges. Right: The component tree $\widetilde{CT}$.

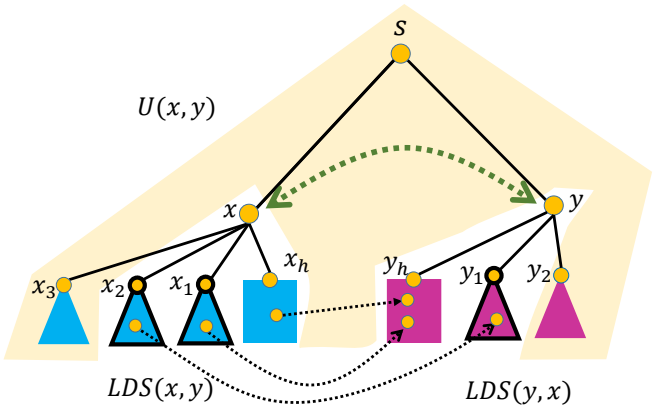

Figure 2 Simulating the first Borůvka phase in algorithm $\mathcal{A}^P_{x,y}$. Each triangle corresponds to a light component in $\mathcal{C}_x, \mathcal{C}_y$. The square boxes correspond to the heavy components H_x, H_y. The framed triangles correspond the subtrees of x, y that belong to the set $\mathsf{LDS}(x,y) \cup \mathsf{LDS}(y,x)$. The dashed green bidirectional arrow represents the xy channel given by the promise. The dashed black arrows correspond to the outgoing edges obtained by x, y from the sketch information of their components. In the example, the light subtrees T_{x_2} and T_{y_1} exchange information over their outgoing edge, which allows y to compute the sketch of the merged component $V(T_{x_2}) \cup V(T_{y_1})$. The sketch of the merged component $V(T_{y_h}) \cup V(T_{x_1}) \cup V(T_{x_h})$ is computed by y by letting x send $\mathsf{Sketch}_{G\setminus\{x,y\}}(V(T_{x_1})) \oplus \mathsf{Sketch}_{G\setminus\{x,y\}}(V(T_{x_h}))$.