

CONNECTIVITY ORACLES FOR GRAPHS SUBJECT TO VERTEX FAILURES*

RAN DUAN[†] AND SETH PETTIE[‡]

Abstract. We introduce new data structures for answering connectivity queries in graphs subject to batched *vertex failures*. A deterministic structure processes a batch of $d \leq d_*$ failed vertices in $\tilde{O}(d^3)$ time and thereafter answers connectivity queries in $O(d)$ time. It occupies space $O(d_* m \log n)$. We develop a randomized Monte Carlo version of our data structure with update time $\tilde{O}(d^2)$, query time $O(d)$, and space $\tilde{O}(m)$ for any failure bound $d \leq n$. This is the first connectivity oracle for general graphs that can efficiently deal with an unbounded number of vertex failures. We also develop a more efficient Monte Carlo *edge* failure connectivity oracle. Using space $O(n \log^2 n)$, d edge failures are processed in $O(d \log d \log \log n)$ time, and thereafter, connectivity queries are answered in $O(\log \log n)$ time, which are correct with high probability. Our data structures are based on a new decomposition theorem for an undirected graph $G = (V, E)$, which is of independent interest. It states that for any terminal set $U \subseteq V$ we can remove a set B of $|U|/(s-2)$ vertices such that the remaining graph contains a Steiner forest for $U - B$ with maximum degree s .

Key words. graph connectivity, dynamic graph, graph sketching, Steiner tree

AMS subject classifications. 05C40, 05C85, 68R10

DOI. 10.1137/17M1146610

1. Introduction. The *dynamic subgraph model* [20, 22, 38, 40, 47, 73] is a constrained dynamic graph model. Rather than allow the graph to evolve in completely arbitrary ways (via an unbounded sequence of edge insertions and deletions), there is assumed to be a *fixed* ideal graph $G = (V, E)$ that can be preprocessed in advance. The ideal graph is susceptible only to the *failure* of edges/vertices and their subsequent *recovery*, possibly with a bound d_* on the number of failures at one time. Queries naturally answer questions about the current failure-free subgraph. This model is useful because it more accurately represents the behavior of many real-world networks: changes to the underlying topology are relatively rare but transient failures very common. More importantly, this model offers the algorithm designer the freedom to explore exotic graph representations. Because *preprocessing time* is not the most critical measure of efficiency, it may be desirable to build a specialized graph representation that facilitates more efficient updates and queries.

Dynamic subgraph connectivity. The dynamic subgraph model was introduced by Frigioni and Italiano [47] who showed that when the ideal graph is planar, vertex failures/recoveries and connectivity queries could be handled in $O(\log^3 n)$ amortized time, after $\tilde{O}(n)$ preprocessing. Their algorithm even allowed the ideal graph to evolve via edge updates, also in $O(\log^3 n)$ amortized time, so long as it remained planar. Dynamic subgraph connectivity structures were later developed for general graphs [9, 20, 22, 38, 42]. Chan, Pătrașcu, and Roditty [22] gave an $O(m^{4/3})$ -space structure

*Received by the editors September 7, 2017; accepted for publication (in revised form) October 26, 2020; published electronically December 17, 2020. This paper includes material from two extended abstracts published in STOC 2010 and SODA 2017.

<https://doi.org/10.1137/17M1146610>

Funding: This research was supported by NSF CAREER grant CCF-0746673 and NSF grants CCF-1217338, CNS-1318294, CCF-1514383, CCF-1637546, and CCF-1815316. The first author is supported by a China Youth 1000-Talent grant.

[†]Tsinghua University, Beijing, China (duanran@mail.tsinghua.edu.cn).

[‡]University of Michigan, Ann Arbor, MI 48109 USA (pettie@umich.edu).

that handles vertex failures/recoveries in $\tilde{O}(m^{2/3})$ amortized time and connectivity queries in $O(m^{1/3})$ time. Duan [38] developed a different $O(m)$ -space structure with the same amortized update and query time as [22], and a new $\tilde{O}(m)$ -space structure with *worst-case* $\tilde{O}(m^{4/5})$ -time updates and $O(m^{1/5})$ time queries. More recently Duan and Zhang [42] presented a worst-case randomized (*Monte Carlo*) structure with update time $\tilde{O}(m^{3/4})$ and query time $\tilde{O}(m^{1/4})$. Each of [22, 38, 42] has an $\Omega(m)$ update time-query time product. Baswana et al. [9] (see also [26]) showed how to maintain a DFS tree in the dynamic subgraph model with $\tilde{O}(\sqrt{mn})$ update time, which supports $O(1)$ -time connectivity queries.

Pătrașcu and Thorup [73] considered a situation where a *batch* of d edges fail simultaneously. They showed that an $O(m)$ -space structure could be constructed that handles updates in $O(d \log^2 n \log \log n)$ time and subsequently answers connectivity queries in $O(\log \log n)$ time. Moreover, they observed that the query time could not be unilaterally improved, by a reduction to the predecessor problem [72, 76]. One downside of the Pătrașcu–Thorup structure is that it requires *exponential* time to compute: it involves solving *sparsest cut* $\tilde{O}(n)$ times on various subgraphs. Using a polynomial time $O(\sqrt{\log n})$ -approximate sparsest cut algorithm [5, 4, 78] instead increases the update time to $O(d \log^{5/2} n \log \log n)$. Pătrașcu and Thorup [73] were motivated by the absence of a fully dynamic connectivity data structure with $\text{poly}(\log n)$ worst-case update time.¹ Kapron, King, and Mountjoy [61] discovered a randomized dynamic connectivity structure with $O(c \cdot \text{poly}(\log n))$ update time that errs with probability n^{-c} . Gibb et al. [51] observed that this data structure can function correctly, with high probability (w.h.p.), without actually storing the graph. This leads to a d -edge failure connectivity oracle with update and query time similar to [73], but using just $\tilde{O}(n)$ space.

The analogous d -*vertex* failure connectivity problem is inherently more complex. Whereas removing d edges can only increase the number of connected components by d , removing d vertices can have an impact on the connectivity that is completely disproportionate to d . When $d = 1$ we can use the block tree representation of biconnected components to answer connectivity queries in constant time; see [16] for data structural details. When $d = 2$ we can use the SPQR tree [11, 16] of each biconnected component to answer queries in $O(1)$ time. A data structure of Kanevsky et al. [60] can answer queries in $O(1)$ time when $d = 3$. Similar ad hoc solutions can also be designed for d -edge failure connectivity oracles, for constant $d \leq 4$ [37, 50, 75, 82]. However, scaling these solutions up, even to an arbitrarily large constant d , becomes prohibitively complex, even in the simpler case of edge failures. In a λ -edge connected graph, encoding all λ -edge cuts is simple with the *cactus* [32] representation, but the simplicity is lost when encoding both λ - and $(\lambda+1)$ -edge cuts. See [34, 35, 36].

In previous work [40] we designed a d -edge failure oracle that reduces the problem to two-dimensional (2D) orthogonal range reporting. Using the range reporting structure of Chan, Larsen, and Pătrașcu [21] gives a d -edge failure structure with $O(d^2 \log \log n)$ update time, $O(\min\{\frac{\log d}{\log \log n}, \frac{\log \log n}{\log \log \log n}\})$ query time, and $O(m \log \log n)$ space, or a somewhat slower update time with $O(m)$ space. By itself, this structure compares favorably with the d -edge failure oracles of [73, 61] when $d = O(\log n)$. However, it has additional properties that make it attractive for use in d -*vertex*

¹There are dynamic connectivity structures with *amortized* $\text{poly}(\log n)$ update time [57, 56, 83]. However, the fastest deterministic worst-case update time is $O(\sqrt{\frac{n(\log \log n)^2}{\log n}})$ [62], a small improvement over the long-standing $O(\sqrt{n})$ bound of [45, 44]. See [68] for *Las Vegas* randomized dynamic connectivity structures with $n^{o(1)}$ worst-case bounds.

failure oracles. Specifically, if D is the set of failed vertices, the update time is actually $O((\sum_{v \in D} \deg_T(v))^2 \log \log n)$, where T is any spanning tree of the graph. In other words, the update time is quadratic in the sum of the T -degrees, independent of their degrees in G .

If G were guaranteed to have an $O(1)$ -degree spanning tree we would immediately have a satisfactory d -vertex failure connectivity oracle with update time $\tilde{O}(d^2)$ and query time $\tilde{O}(1)$. Of course, there is no such guarantee. Every *bridge* edge appears in every spanning tree T , so a vertex incident to many bridges must have high T -degree. Since bridges are easy to deal with this is not a very convincing counterexample. One might hope that if G had sufficient connectivity, a low degree spanning tree could be found. This is the approach taken by Borradaile, Pettie, and Wulff-Nilsen's [16] d -failure connectivity oracles for *planar* graphs. Barnette's theorem [7] states that every triconnected planar graph has a degree-3 spanning tree, which can be found in linear time [30, 79]. However, the analogues of Barnette's theorem for general graphs are too weak to be of any use. Czumaj and Strothmann [30, 79] proved that a k -connected graph with maximum degree $\Delta(G) \leq k(\Delta_T - 2) + 2$ has a degree- Δ_T spanning tree, which can be found in polynomial time. If, however, the maximum degree is at least $\Delta(G) \geq k(\Delta_T - 1)$ it is NP-hard to decide if there is a degree- Δ_T spanning tree. Thus, even if we could force G to be k -connected for some large constant k , it would not help to find a low degree spanning tree.

In [40] we developed a d_* -vertex failure connectivity oracle that offers a trade-off between update time and size. For any integer parameter $c \geq 1$, the space of the data structure is $O(d_*^{1-2/c} mn^{1/c-1/(c \log(2d_*))} \log^2 n)$ and the time to process $d \leq d_*$ vertex failures is $O(d_*^{2c+4} \log^2 n \log \log n)$. Thereafter connectivity queries can be answered in $O(d)$ time. The main drawbacks of [40] are its conceptual complexity and very poor trade-off between space and update time. Henzinger and Neumann [55] recently showed how any d -vertex failure connectivity oracle could be transformed to support fully dynamic updates in the dynamic subgraph model, where vertices fail and recover individually.

New results. In this paper we present dramatically better d -vertex failure connectivity oracles that match or improve on [40] in every measure of efficiency except construction time. Using space $O(d_* m \log n)$, a batch D of $d \leq d_*$ vertex failures is processed in $O(d^3 \log^3 n)$ time such that connectivity queries in $G - D$ can be answered in $O(d)$ time.² The construction time is $O(mn \log n)$. Note that there is now no trade-off between space and update time. Clearly any pair of $(d_* + 1)$ -connected vertices cannot be disconnected by d failures. By preprocessing the graph with the linear time Nagamochi–Ibaraki algorithm [67], we can replace $E(G)$ by an equivalent subgraph containing $\bar{m} = \min\{m, (d_* + 1)n\}$ edges. Thus, the factors of m in the space and construction time can be replaced with \bar{m} .

In the extended abstract of this work [41, section 7], we claimed a randomized Monte Carlo structure that occupies $\tilde{O}(m)$ space and has update and query times $\tilde{O}(d^2)$ and $O(d)$. This was an erroneous claim; we do not see any way to store this structure in less than $\Omega(d_* m)$ space. In this paper we present a different randomized Monte Carlo structure that uses space $O(m \log^6 n)$, and has update and query times $O(d^2 \log d \log^2 n \log \log n)$ and $O(d)$. This solution is more sophisticated than the one described in [41, section 7] and generalizes the Kapron, King, and Mountjoy [61] sketch technique in ways that may be of independent interest. We use vertex sampling rather than edge-sampling and show that sketches for certain subgraphs of a complete

²The notation $G - D$ is short for the subgraph of G induced by $V(G) - D$.

bipartite graph $A \times B$ can be generated “on the fly” using space $\tilde{O}(|A| + |B|)$ rather than a naive bound of $O(|A \times B|)$.

Some of the techniques used in our Monte Carlo d -vertex failure oracle can be repurposed to improve the state of the art in d -edge failure oracles [73, 61, 51]. We show that with $O(n \log^2 n)$ space, d edge failures can be processed in $O(d \log d \log \log n)$ time in expectation and thereafter support connectivity queries in $O(\log \log n)$ time, which are correct w.h.p.

Our data structures are based on a new graph decomposition theorem, which is obtained from a recursive version of the Fürer–Raghavachari [48] algorithm for approximating the minimum degree spanning tree. The theorem states that for any undirected graph $G = (V, E)$, terminal set $U \subseteq V$, and integer s , there exists a set of $n/(s-2)$ vertices B that can be removed, such that $U - B$ is spanned by a degree- s Steiner forest in the graph $G - B$. We believe this decomposition theorem is of independent interest.

Refer to Table 1 for a summary of d -edge failure and d -vertex failure connectivity oracles.

Lower bounds. One question raised by [40] is whether it is possible for a d -vertex failure oracle to match the $\tilde{O}(1)$ query time of existing d -edge failure oracles [73, 40, 61, 51]. There is now strong circumstantial evidence that no such data structure exists with reasonable update time. In particular, if the *integer 3SUM conjecture*³ holds, then any d -vertex failure connectivity oracle with subquadratic preprocessing and reasonable update time must have $\Omega(d^{1/2-o(1)})$ query time [63]. Henzinger et al. [54] showed that the *OMv conjecture*⁴ on the hardness of online matrix-vector multiplication implies an $\Omega(d^{1-o(1)})$ query lower bound, even if any polynomial preprocessing is allowed. Thus, beating $O(d)$ query time would require refuting a plausible conjecture. Of course, the plausibility of the 3SUM and OMv conjectures continues to be actively scrutinized. Stronger forms of the 3SUM and OMv conjectures have already been refuted; see [6, 53, 65]. Whereas d -edge failure connectivity oracles can be stored in sublinear $\tilde{O}(n)$ space [51], this is not possible for vertex failures. It is straightforward to see that any subgraph of the complete bipartite graph K_{n,d_*+1} can be reconstructed with a d_* -failure oracle, implying such an oracle occupies $\Omega(\min\{m, d_*n\})$ bits of space.

Related work. Much of the previous work in the d -failure model has focused on computing approximate shortest paths avoiding edge and vertex failures. Demetrescu et al. [31] gave an exact shortest path oracle for weighted directed graphs subject to $d = 1$ failure. It occupies $O(n^2 \log n)$ space and answers queries in constant time. The construction time for this oracle was later improved by Bernstein and Karger [14]. An analogous result for $d = 2$ failures was presented by Duan and Pettie [39], which uses space $O(n^2 \log^3 n)$ and query time $O(\log n)$. *Approximate* distance oracles for d edge failures were given for general graphs [25], with stretch that grows linearly in d .

These problems have also been studied on special graph classes. Borradaile, Pettie, and Wulff-Nilsen [16] described connectivity oracles for planar graphs subject to d -edge failures or d -vertex failures. See Baswana, Lath, and Mehta [8] for exact

³The 3SUM problem is, given a set A of n numbers, to determine if there exist $a, b, c \in A$ for which $a + b + c = 0$. There are now known to be $O(n^2/\text{poly}(\log n))$ algorithms for both integer inputs [6] and real inputs [53, 46, 52]. The integer 3SUM conjecture asserts that the problem requires $\Omega(n^{2-o(1)})$ time, even if $A \subset \{-n^3, \dots, n^3\}$.

⁴The OMv conjecture is that given a matrix $M \in \{0, 1\}^{n \times n}$ to be preprocessed and n vectors $v_1, \dots, v_n \in \{0, 1\}^n$ presented online, the total cost of preprocessing and computing the products $\{Mv_i\}_{1 \leq i \leq n}$ is $\Omega(n^{3-o(1)})$. Note that fast matrix multiplication is not obviously helpful in this context since Mv_i must be reported before receiving v_{i+1} .

TABLE 1
The lower bounds of Pătrașcu and Thorup are unconditional whereas the lower bounds of Kopelowitz et al. and Henzinger et al. rely on unproven conjectures. Whenever d_ is available at construction time we can replace m (the number of edges) with $\tilde{m} \stackrel{\text{def}}{=} \min\{m, (d_* + 1)n\}$ [67]. See [16, 11, 60] for data structure details when $d \in \{1, 2, 3\}$.*

d-EDGE FAILURE STRUCTURES AND LOWER BOUNDS

	Update	Query	Space	Preprocessing
Pătrașcu and Thorup [73]	$O(d \log^2 n \log \log n)$	$O(\log \log n)$	$O(m)$	exponential
	$O(d \log^{5/2} n \log \log n)$			polynomial
Duan and Pettie [40]	$O(d \text{poly}(\log n))$	$\Omega(\log \log n)$	any	any
	$O(d^2 \log \log n)$	$\min\left\{O\left(\frac{\log d}{\log \log n}\right), O\left(\frac{\log \log n}{\log \log \log n}\right)\right\}$	$O(m \log \log n)$	linear in space
Kapron, King, and Mountjoy [61] Gibb et al. [51]	$O(d^2 \log^{\epsilon} n)$	$O(\log \log \log n)$	$O(m)$	
	$O(d \log d \log^3 n)$	$O(\log \log n)$, connect w.h.p.	$O(n \log^2 n)$	near linear (Rand.)
New	$O(d \log d \log \log n)$ in expect. w.h.p.	$O(\log \log n)$ connect w.h.p.	$O(n \log^2 n)$	near linear (Rand.)
	$O(d \log n \log \log n)$			

d-VERTEX FAILURE STRUCTURES AND LOWER BOUNDS

	Update	Query	Space	Preprocessing
$d = 1$: Block tree $d = 2$: SPQR tree $d = 3$: Kanevsky et al. [60]	$O(1)$	$O(1)$	$O(n)$	linear
	$O(1)$	$O(1)$	$O(n)$	linear
	$O(1)$	$O(1)$	$O(n)$	near linear
Duan and Pettie [40]	$\tilde{O}(d^{2c+4})$	$O(d)$	$\tilde{O}(d_*^{1-2-\frac{1}{c \log 2d_*}})$	linear in space
	$O(\text{poly}(d, \log n))$ $O((dn)^{1-\epsilon})$	$\Omega(d^{1-o(1)})$	any	$O(\text{poly}(n))$
Henzinger et al. [54] (assuming OMv conjecture)	$O(\text{poly}(d, \log n))$ or $O((dn)^{1-\epsilon})$	$\Omega(d^{1-o(1)})$	any	$O(\text{poly}(n))$
	$O(\text{poly}(d, \log n))$ or $O((dn)^{1-\epsilon})$	$\Omega(d^{1/2-o(1)})$	any	$O(mn^{1-\epsilon} \text{poly}(d))$
Kopelowitz, Pettie, and Porat [63] (assuming 3SUM conjecture)	$O(d^3 \log^3 n)$	$O(d)$	$O(d_* \tilde{m} \log n)$	$O(\tilde{m} n \log n)$ (Det.)
	$O(d^2 \log d \log^2 n \log \log n)$ in exp. $O(d^2 \log^3 n \log \log n)$	$O(d)$, correct w.h.p.	$O(\tilde{m} \log^6 n)$	$O(\tilde{m} n \log n)$ (Rand.)
New				

distance oracles for planar graphs avoiding $d = 1$ failure, and see Abraham et al. [1, 2] for approximate distance oracles for planar graphs and graphs of bounded doubling dimension.

Parter and Peleg [70] considered the problem of computing a subgraph that preserves shortest paths from s sources after a single edge or vertex failure. They proved that $\Theta(s^{1/2}n^{3/2})$ edges are necessary and sufficient for every s . See also [15, 18, 19, 24, 33, 69, 71] for spanners (subgraphs) that preserve approximate distances subject to edge or vertex failures.

Recently researchers have considered reachability problems on directed graphs subject to vertex failures. Choudhary [29] gave an optimal $O(n)$ -space, $O(1)$ -query time reachability oracle for $d = 2$ failures. Baswana, Choudhary, and Roditty [10] considered the problem of finding a sparse subgraph that preserves reachability from a single source, subject to d vertex failures. They proved that $\Theta(2^d n)$ edges are necessary and sufficient.

1.1. Organization. In section 2 we review the Euler tour structure of [40] for handling d edge failures. We begin section 3 with a sketch of the Fürer–Raghavachari algorithm **FR-Tree**, then describe our decomposition algorithm **Decomp**. In section 4 we observe that by applying **Decomp** iteratively, we naturally obtain a representation of the graph as a *low degree hierarchy*. Section 4 describes how to build a d -failure connectivity oracle, by supplementing the low degree hierarchy with suitable data structures. The algorithms for deleting failed vertices and answering connectivity queries are presented in section 5. The basic algorithm for deleting failed vertices takes $\tilde{O}(d^4)$ time using standard 2D orthogonal range reporting data structures. In section 6 we give three distinct ways to reduce this to $\tilde{O}(d^3)$ using other orthogonal range searching structures. In section 7 we present a randomized Monte Carlo version of our data structure with update time $\tilde{O}(d^2)$ and space $\tilde{O}(m)$, and in section 7.4 we give a more efficient d -edge failure connectivity oracle. Several open problems are discussed in section 8.

2. The Euler tour structure. In this section we describe the *ET-structure* for handling connectivity queries avoiding multiple vertex and edge failures. When handling only d edge failures, the performance of the ET-structure is incomparable to that of Pătrașcu and Thorup [73] in nearly every respect.⁵ The strength of the ET-structure is that if the graph contains a low degree tree T , the time to delete a vertex is a function of its degree in T ; incident edges not in T are deleted implicitly. We prove Theorem 2.1 in the remainder of this section.

THEOREM 2.1. *Let $G = (V, E)$ be a graph, with $m = |E|$ and $n = |V|$, and let $\mathcal{F} = \{T_1, \dots, T_{|\mathcal{F}|}\}$ be a set of vertex disjoint trees in G . (\mathcal{F} does not necessarily span connected components of G .) There is a data structure **ET**(G, \mathcal{F}) that supports the following operations. Suppose D is a set of failed edges, of which d are tree edges in \mathcal{F} and d' are nontree edges. Deleting D splits some subset of the trees in \mathcal{F} into at most $2d$ trees $\mathcal{F}' = \{T'_1, \dots, T'_{2d}\}$. In $O(d^2q + d')$ time we can report which pairs of trees in \mathcal{F}' are connected by an edge in $E - D$. In $O(\min\{\frac{\log \log n}{\log \log \log n}, \frac{\log d}{\log \log n}\})$ time we can*

⁵The ET-structure is significantly faster in terms of construction time (near-linear versus a large polynomial or exponential time) though it may use slightly more space: $O(m \log \log n)$ versus $O(m)$. It handles d edge deletions exponentially faster for bounded d ($O(\log \log n)$ vs. $\Omega(\log^2 n \log \log n)$) but is slower as a function of d : $O(d^2 \log \log n)$ versus $O(d \log^2 n \log \log n)$ time. The query time is essentially the same for both structures, namely $O(\log \log n)$. Whereas the ET-structure naturally maintains a certificate of connectivity (a spanning tree), the Pătrașcu–Thorup structure requires modification and an additional logarithmic factor in the update time to maintain a spanning tree.

determine which tree in \mathcal{F}' contains a given vertex. Using space $O(m \log \log n)$ the value of q is $O(\log \log n)$; using space $O(m)$ the value of q is $O(\log^\epsilon n)$.

Our data structure uses Chan, Larsen, and Pătrașcu's [21] structure for orthogonal range reporting on the integer grid $[U] \times [U]$. They showed that given a set of N points, there is a data structure with size $O(N \log \log N)$ such that given $x, y, w, z \in [U]$, the set of points in $[x, y] \times [w, z]$ can be reported in $O(\log \log U + k)$ time, where k is the number of reported points. If the space is reduced to $O(N)$ the update time becomes $O(\log^\epsilon U + k)$ for any fixed $\epsilon > 0$.

For a tree T , let $\text{Euler}(T)$ be a list of its vertices encountered during an Euler tour of T (an undirected edge is treated as two directed edges), where we only keep the *first* occurrence of each vertex. One may easily verify that removing f edges from T partitions it into $f + 1$ connected subtrees and splits $\text{Euler}(T)$ into at most $2f + 1$ intervals, where the vertices of a connected subtree are the union of some subset of the intervals. To build $\mathbf{ET}(G = (V, E), \mathcal{F})$ we build the following structure for each pair of trees $(T_1, T_2) \in \mathcal{F} \times \mathcal{F}$; note that T_1 and T_2 may be the same. Let m' be the number of edges connecting T_1 and T_2 . Let $\text{Euler}(T_1) = (u_1, \dots, u_{|T_1|})$, $\text{Euler}(T_2) = (v_1, \dots, v_{|T_2|})$, and $U = \max\{|T_1|, |T_2|\}$. We define the point set $P \subseteq [U] \times [U]$ to be $P = \{(i, j) \mid \{u_i, v_j\} \in E\}$. Suppose D is a set of edge failures including d_1 edges in T_1 , d_2 in T_2 , and d' nontree edges. Removing D splits T_1 and T_2 into $d_1 + d_2 + 2$ connected subtrees and partitions $\text{Euler}(T_1)$ into a set $I_1 = \{[x_i, y_i]\}_i$ of $2d_1 + 1$ intervals and $\text{Euler}(T_2)$ into a set $I_2 = \{[w_i, z_i]\}_i$ of $2d_2 + 1$ intervals. For each pair i, j we query the 2D range reporting data structure for points in $P \cap ([x_i, y_i] \times [w_j, z_j])$. However, we stop the query the moment it reports some point corresponding to a nonfailed edge, i.e., one in $E - D$. Since there are $(2d_1 + 1) \times (2d_2 + 1)$ queries and each failed edge in D can only be reported in *one* such query, the total query time is $O(d_1 d_2 q + d')$, where q is either $\log \log n$ or $\log^\epsilon n$, depending on the space usage. See Figure 1 for an illustration.

Assuming that $m' \geq 1$, the space for the data structure restricted to T_1 and T_2 is $O(m' \log \log n)$ or $O(m')$. In order to avoid spending any space on pairs (T_1, T_2) with $m' = 0$, we maintain a hash table of tree pairs with at least one edge between them. Since each nontree edge contributes to the space of at most one tree pair (T_1, T_2) , the overall space for $\mathbf{ET}(G, \mathcal{F})$ is $O(m \log \log n)$ or $O(m)$. For the last claim of the theorem, observe that if a vertex u lies in an original tree $T_1 \in \mathcal{F}$, we can determine which tree in \mathcal{F}' contains it by performing a predecessor search over the left endpoints of intervals in I_1 . This can be accomplished in the minimum of $O(\frac{\log \log n}{\log \log \log n})$ time [72] or $O(\frac{\log d}{\log \log n})$ time [77] after $O(d^2)$ preprocessing on a $\Theta(\log n)$ -bit word-RAM.

Corollary 2.2 demonstrates how $\mathbf{ET}(G, \cdot)$ can be used to answer connectivity queries avoiding edge and vertex failures.

COROLLARY 2.2. *Let T be any spanning tree of $G = (V, E)$. The data structure $\mathbf{ET}(G, \{T\})$ occupies space $O(m \log \log n)$ (or $O(m)$) and supports the following operations. Given a set $D \subset E$ of edge failures, d of which are tree edges and d' are nontree edges, D can be processed in $O(d^2 \log \log n + d')$ time (or $O(d^2 \log^\epsilon n + d')$ time) so that connectivity queries in the graph $(V, E - D)$ can be answered in $O(\min\{\frac{\log \log n}{\log \log \log n}, \frac{\log d}{\log \log n}\})$ time. If $D \subset V$ is a set of vertex failures, let $d = \sum_{v \in D} \deg_T(v)$ be the sum of their T -degrees. The update time is $O(d^2 \log \log n)$ (or $O(d^2 \log^\epsilon n)$) and the query time is $O(\min\{\frac{\log \log n}{\log \log \log n}, \frac{\log d}{\log \log n}\})$.*

Proof. Using $\mathbf{ET}(G, \{T\})$ we split T into $d + 1$ subtrees and $\text{Euler}(T)$ into a set I of $2d + 1$ connected intervals, in which each connected subtree is made up of some

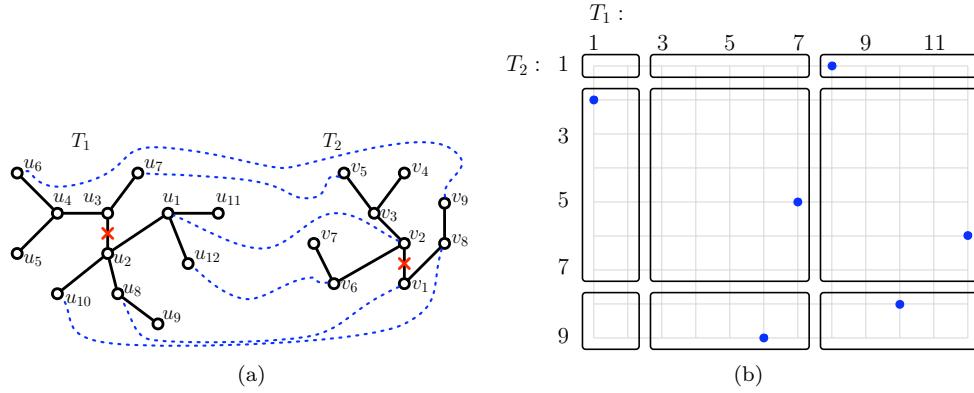


FIG. 1. (a) Here T_1 and T_2 are two trees where $\text{Euler}(T_1) = (u_1, \dots, u_{12})$ and $\text{Euler}(T_2) = (v_1, \dots, v_9)$ list their vertices according to first appearance in some Euler tours of T_1 and T_2 . (It does not matter which Euler tour we pick.) There are six nontree edges connecting T_1 and T_2 , marked by dashed curves. If the edges $\{u_2, u_3\}$ and $\{v_1, v_2\}$ are removed, T_1 and T_2 are split into four subtrees, say, T'_1, T'_2, T'_3, T'_4 , and both $\text{Euler}(T_1)$ and $\text{Euler}(T_2)$ are split into three intervals, namely $X_1 = (u_1, u_2), X_2 = (u_3, \dots, u_7), X_3 = (u_8, \dots, u_{12}), Y_1 = (v_1), Y_2 = (v_2, \dots, v_7)$, and $Y_3 = (v_8, v_9)$. Each tree T'_i is identified with some subset of the intervals: T'_1, \dots, T'_4 are identified with $\{X_1, X_3\}, \{X_2\}, \{Y_1, Y_3\}$, and $\{Y_2\}$. (b) The point (i, j) (marked by a blue dot) is in our point set if $\{v_i, u_j\}$ is a nontree edge. To determine if, for example, T'_1 and T'_4 are connected by an edge, we perform two 2D range queries, $X_1 \times Y_2$ and $X_3 \times Y_2$, and keep at most one point (i.e., a nontree edge) for each query. In general, removing d_1 edges from T_1 and d_2 edges from T_2 necessitates $(2d_1 + 1)(2d_2 + 1)$ 2D range queries to determine incidences between all pairs of subtrees. In this example we require nine 2D range queries, indicated by boxes in the point set diagram.

subset of the intervals. Using $O(d^2)$ 2D range queries, in $O(d^2 \log \log n + d')$ time we find at most one edge connecting each pair in $I \times I$. (In the case of vertex failures, no range queries are performed for the intervals containing singleton vertices in D .) In $O(d^2)$ time we find the connected components of $E - D$ or $V - D$ and store with each interval a representative vertex from its component. To answer a query (u, v) we only need to determine which subtree u and v are in, which involves two predecessor queries over the left endpoints of intervals in I . This takes $O(\min\{\frac{\log \log n}{\log \log \log n}, \frac{\log d}{\log \log n}\})$ time. \square

Corollary 2.2 motivates us to look for conditions under which G contains a low degree spanning forest, say, with degree at most s . In the next section we show that although G may not have a degree- s spanning forest, there are $O(n/s)$ critical nodes that, if they were removed, would let the remaining graph be spanned by a degree- s spanning forest.

3. A new graph decomposition theorem. Let $G = (V, E)$ be an undirected graph and $U \subseteq V$ be a set of *terminals*. We call a forest $T \subseteq E$ a *Steiner forest* for U if $u, v \in U$ are connected in T iff they are connected in G . Fürer and Raghavachari [48] proved that the minimum degree spanning forest (if $U = V$) and minimum degree Steiner forest could be approximated to within 1 of optimal in polynomial time.⁶

⁶Fürer and Raghavachari [48] claimed a running time of $O(|U| \alpha(m, n) \log |U|)$. The $\alpha(m, n)$ factor can be removed using the incremental-tree set-union structure of Gabow and Tarjan [49].

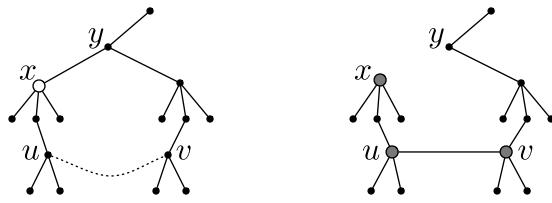


FIG. 2. A fragment of a larger tree is depicted. Swapping $\{u, v\}$ for $\{x, y\}$ yields a new tree with at least one fewer nodes with degree $\Delta(T_0)$.

THEOREM 3.1 (Fürer and Raghavachari [48]). *Suppose G contains a Steiner forest for U with maximum degree Δ^* . A Steiner forest T for U with maximum degree $\Delta^* + 1$ can be computed in $O(|U|m \log |U|)$ time.*

Let $\text{FR-Tree}(G, U)$ be the procedure that computes T . Our decomposition theorem is not concerned with Δ^* but with *other* properties of the forest T . In order to see how these properties arise, we sketch how the $\text{FR-Tree}(G, U)$ algorithm works in the simpler case in which $U = V$. Let $\Delta(G')$ denote the maximum degree in the graph G' .

The algorithm begins with any spanning forest T_0 and iteratively tries to *improve* T_0 , yielding $T_1, T_2, \dots, T_\omega$, such that (i) $\Delta(T_{i+1}) \leq \Delta(T_i)$ and (ii) the set of degree- $\Delta(T_i)$ nodes in T_{i+1} is a *strict* subset of the degree- $\Delta(T_i)$ nodes in T_i . The number of improvements is clearly finite. Since any tree contains fewer than $n/(k-1)$ nodes with degree at least k , the total number of improvements is at most $\sum_{k=\Delta(T_\omega)}^{\Delta(T_0)} n/(k-1) = O(n \log \frac{\Delta(T_0)}{\Delta(T_\omega)}) = O(n \log n)$.

The FR-Tree algorithm only searches for a particular class of improvements that can be found in linear time, leading to an $O(mn \log n)$ time bound. Let T_0 be the current spanning tree. All vertices with degree $\Delta(T_0)$ and $\Delta(T_0) - 1$ are initially marked *bad* and all others *good*. (In the diagrams in Figure 2 white nodes have degree $\Delta(T_0)$, gray nodes have degree $\Delta(T_0) - 1$, and black nodes have degree less than $\Delta(T_0) - 1$.) The simplest *single-swap* improvement arises if there is a non- T_0 edge $\{u, v\}$ such that u and v are good (black) and a bad vertex x with degree $\Delta(T_0)$ appears on the unique cycle of $T \cup \{\{u, v\}\}$. In this case we choose any edge $\{x, y\}$ incident to x on the cycle and set $T_1 \leftarrow T_0 - \{\{x, y\}\} \cup \{\{u, v\}\}$, thereby eliminating a degree- $\Delta(T_0)$ vertex (namely x , and perhaps even y) but possibly increasing the number of degree- $(\Delta(T_0) - 1)$ vertices (namely u and v).

In general the FR-Tree algorithm considers improvements composed of an arbitrarily large number of edge-swaps. While there exists an unscanned edge $\{u, v\}$ where both u and v are marked good, it marks all bad vertices *good* on the fundamental cycle of $T_0 \cup \{\{u, v\}\}$. Thus, a formerly bad good vertex is one whose degree can be reduced by 1 via a sequence of edge-swaps that does not introduce any degree- $\Delta(T_0)$ vertices. If a degree- $\Delta(T_0)$ vertex is ever marked good, an improvement has been detected and the sequence of swap edges that created it can easily be reconstructed. See Figure 3. Every time this procedure finds an improvement we obtain a new spanning tree and begin the search for another improvement from scratch. Let T_ω be the spanning tree for which this procedure fails to find an improvement. Let B be the set of vertices still marked *bad*. By definition B includes *all* vertices with degree $\Delta(T_\omega)$ and some subset of the vertices with degree $\Delta(T_\omega) - 1$. Consider what happens to G and T_ω if we removed all B -vertices from the graph. FR-Tree 's search for improvements guarantees

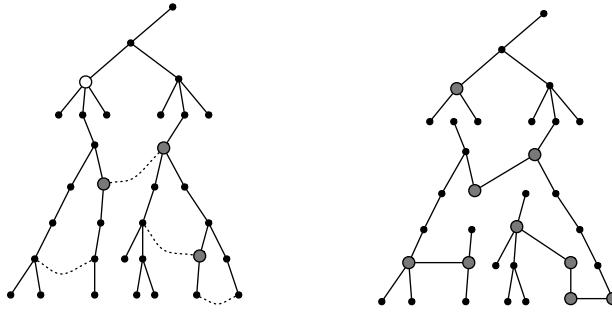


FIG. 3. A fragment of a larger tree is depicted. A sequence of edge-swaps reduces the number of degree- $\Delta(T_0)$ vertices but may increase the number of degree- $(\Delta(T_0) - 1)$ vertices.

that $T_\omega - B$ is a *spanning forest* of the graph $G - B$. Indeed, if there were an edge $\{u, v\}$ connecting two distinct trees of $T_\omega - B$, then all B -vertices on the fundamental cycle of $T_\omega \cup \{\{u, v\}\}$ would have been marked *good* and therefore u and v would not have been in distinct trees of $T_\omega - B$ after all. In general, the output of $\text{FR-Tree}(G, V)$ is the pair (T_ω, B) .

When the terminal set U is a strict subset of V , the execution of $\text{FR-Tree}(G, U)$ is similar, except that T_0, \dots, T_ω are *Steiner trees* (which might not span V). Each improvement to T_i substitutes for some edges in T_i an equal number of *paths*, whose intermediate vertices come from $V - V(T_i)$. See [48]. Theorem 3.2 summarizes the properties of the FR-Tree algorithm that we actually use.

THEOREM 3.2 (see [48]). *The $\text{FR-Tree}(G, U)$ algorithm returns a pair (T, B) , where T is a Steiner forest for U and $B \subset V$ comprises all vertices with T -degree $\Delta(T)$ and some subset of vertices with T -degree $\Delta(T) - 1$. If $u, v \in U$ are disconnected in $T - B$, then they are also disconnected in $G - B$.*

The degree $\Delta(T - B)$ is by definition at most $\Delta(T) - 1$, which may still be too large. Theorem 3.3 shows that by iteratively applying the FR-Tree algorithm to the components of $T - B$ we can reduce the maximum degree to any desired bound $s \geq 3$, at the cost of increasing the set B of “bad” vertices.

THEOREM 3.3 (the decomposition theorem). *Let $U \subseteq V$ be a terminal set in a graph $G = (V, E)$ and $s \geq 3$. There is an algorithm $\text{Decomp}(G, U, s)$ that returns a pair (T, B) such that the following hold.*

1. T is a Steiner forest for U and $T - B$ is a Steiner forest for $U - B$.
2. $\Delta(T - B) \leq s$.
3. $|B| < |U|/(s - 2)$ and $|B \cap U| < |U|/(s - 1)$.

The running time of Decomp is $O(|U|m \log |U|)$.

In the remainder of this section we give the $\text{Decomp}(G, U, s)$ algorithm and prove Theorem 3.3. An invocation of Decomp consists of the following three steps.

Step 1. Let (T', B') be the output of $\text{FR-Tree}(G, U)$. If $\Delta(T') \leq s$, then we are done, and return the pair (T', \emptyset) .

Step 2. Partition the edge set of T' into minimal trees $\{t_i\}$ such that the leaves of each t_i are either B' -nodes or leaves of T' , and hence U -nodes. Let $B'[t_i]$ be the B' nodes in t_i and $V[t_i]$ be the set of all vertices in $G - B'$ reachable from vertices in $V(t_i) - B'[t_i]$. (When $U = V$, $V[t_i]$ is exactly $V(t_i) - B'[t_i]$; in general $V[t_i]$ may contain vertices outside of $V(T')$. See Figure 4.) Let $G[t_i]$ be the graph whose vertex

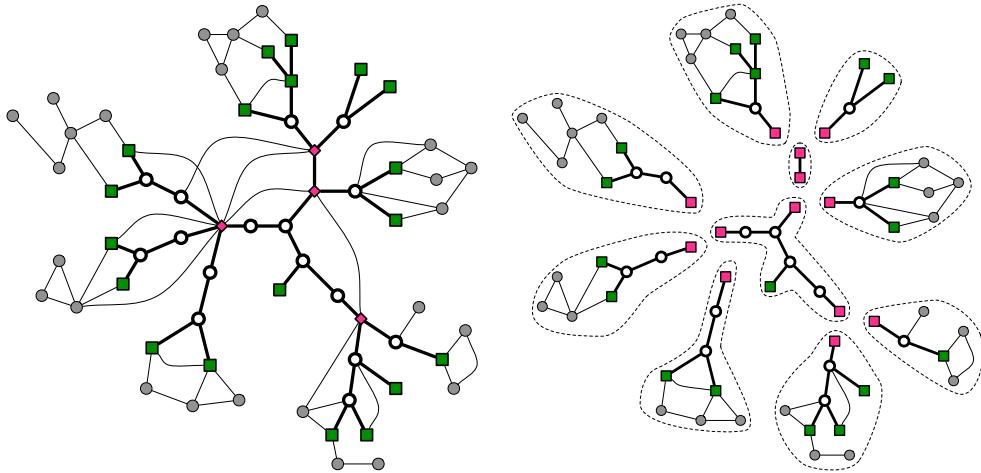


FIG. 4. Left: the output of FR-Tree. Square green nodes are terminals; pink diamonds are B' -nodes (and may be terminals); thick edges are part of T' ; gray vertices are outside $V(T')$. Right: detaching the edges adjacent to B' nodes creates 10 subtrees; non- $V(T')$ nodes are connected to at most one subtree; *Decomp* is called recursively on each subgraph; B' -nodes have degree 1 in these recursive calls and are designated terminals (square nodes).

set is $V[t_i] \cup B'[t_i]$ and whose edge set includes all edges induced by $V[t_i]$ and, for each $u \in B'[t_i]$, the unique T' -edge connecting u to $V(t_i)$. For each t_i , obtain a pair (T_i, B_i) by recursively calling *Decomp*($G[t_i], (V[t_i] \cap U) \cup B'[t_i], s$). Observe that $B'[t_i]$ are included as terminals in the recursive call, even if they are not members of U . See Figure 4 for an illustrative example.

Step 3. Return the pair (T, B) where

$$T = \bigcup_i T_i \quad \text{and} \quad B = B' \cup \bigcup_i B_i.$$

We need to establish all the claims: that $T - B$ is, in fact, a Steiner forest of $U - B$ with maximum degree s , that B has the right cardinality, and that the running time is $O(|U|m \log |U|)$.

If the algorithm halts at Step 1, then T' is, by Theorem 3.2, a Steiner forest for U in G . Suppose that the algorithm does not halt at Step 1 and let $P(u_0, u_k)$ be a path in T' between $u_0, u_k \in U$. Partition it into subpaths $P(u_0, u_1), \dots, P(u_{k-1}, u_k)$, where u_1, \dots, u_{k-1} are all the B' -nodes encountered on the path. By construction, each $P(u_i, u_{i+1})$ is completely contained in some tree t_i and the endpoints of this path are terminals in the recursive call to *Decomp*($G[t_i], (V[t_i] \cap U) \cup B'[t_i], s$), so, by the inductive hypothesis, the tree T_i returned contains a (possibly different) path between u_i and u_{i+1} . By Theorem 3.2 again, the graphs $\{G[t_i]\}$ intersect only at B' -nodes, which necessarily occur as leaves in the $\{T_i\}$ trees, so the edge set $T = \bigcup_i T_i$ returned is, in fact, a Steiner forest for U . By Theorem 3.2, all nodes in B have T -degree at least s and all nodes in $T - B$ have T -degree at most s . Moreover, if $u, v \in U$ are disconnected in $T - B$, then they are disconnected in $G - B$. This follows from Theorem 3.2 if u and v are in different trees t_i, t_j , and by induction on the output of *Decomp*($G[t_i], (V[t_i] \cap U) \cup B'[t_i], s$) if u, v are both in t_i .

We now prove that B has the claimed cardinality, using the property that all B -nodes have degree at least s in T .

LEMMA 3.4. *Let T be any minimal Steiner tree for U . The number of nodes in T with T -degree at least s is at most $g(|U|) = \lfloor \frac{|U|-2}{s-2} \rfloor$. The number of U -nodes in T with T -degree at least s is at most $h(|U|) = \lfloor \frac{|U|-2}{s-1} \rfloor$.*

Proof. Due to the minimality of T , all leaves are necessarily U -nodes. Moreover, we can assume without loss of generality that all internal nodes have degree at least 3, by splicing out paths of degree-2 vertices. When $|U| \leq s-1$ we have $g(|U|) = 0$ and when $|U| \leq s$ we have $h(|U|) = 0$. The claimed bounds on g and h hold when there is exactly one internal node. In general, choose an internal node u adjacent to exactly one internal (nonleaf) node. If u is adjacent to at least $s-1$ leaves, then it contributes 1 to the $g(|U|)$ tally; remove its incident leaves and designate u a U -node. We preserve the property that all leaves are U -nodes, and since the net loss in the number of U -nodes is at least $s-2$, we have $g(|U|) \leq g(|U| - (s-2)) + 1$. Observe that u only contributes to the $h(|U|)$ tally if it is already a U -node. In this case we have a loss of $s-1$ U -nodes, which implies that $h(|U|) \leq h(|U| - (s-1)) + 1$. The claimed bounds on g and h follow by induction on $|U|$. \square

To analyze the running time we imagine that a *single* global Steiner tree for U is being maintained, which is the union of the current Steiner trees in the deepest recursive calls. The initial tree provided to a call to **FR-Tree** is therefore just a fragment of the global Steiner tree, whose maximum degree is some $k \geq s+1$. Each iteration of this call to **FR-Tree**, except the last, finds an *improvement*, which reduces the number of maximum-degree nodes in its fragment by at least one. Say a k -improvement is one that reduces the number of degree- k nodes. If the current global Steiner tree has maximum degree k , the total number of k -improvements that can be found, in all recursive calls, is at most $|U|/(k-2)$. The initial value of k is certainly at most $|U|$. Since each improvement takes linear time, the total time for all improvements is at most $O(m) \cdot \sum_{k=s+1}^{|U|} |U|/(k-2) = O(|U|m \log(|U|/s))$.

4. The low degree hierarchy. We can apply Theorem 3.3 iteratively to create a *low degree hierarchy*. Fix $s = 4$ and generate a set of pairs $\{(T_i, B_i)\}$ as follows:

$$\begin{aligned} (T_0, B_0) &\leftarrow \text{Decomp}(G, V, 4), \\ (T_1, B_1) &\leftarrow \text{Decomp}(G, B_0, 4), \\ &\dots \\ (T_i, B_i) &\leftarrow \text{Decomp}(G, B_{i-1}, 4), \\ &\dots \\ (T_p, \emptyset) &\leftarrow \text{Decomp}(G, B_{p-1}, 4). \end{aligned}$$

In other words, the “bad” vertices for T_0 form the terminal set for T_1 , and in general, the bad vertices for T_{i-1} form the terminal set for T_i . We end, of course, at the first T_p with degree at most $s = 4$, so $B_p = \emptyset$. It follows from Theorem 3.3 that $|B_0| < n/3$ and in general that $|B_i| < |B_{i-1}|/2$, so $p < \log n - 1$ levels suffice.

Define \mathcal{T}_i to be the set of trees in $T_i - B_i$ and \mathcal{T} to be the set of all trees in $\mathcal{T}_0, \dots, \mathcal{T}_p$, as if each forest were on a disjoint vertex set. Theorem 3.3 implies that the forest \mathcal{T}_i has two useful properties: it has maximum degree 4, and it is a Steiner forest for $B_{i-1} - B_i$. Suppose $v \in V(\mathcal{T}_i) \cap B_{i-1}$ is a *terminal* for the first time in \mathcal{T}_i . We treat this copy of v as the “principal” copy in \mathcal{T} ; all other copies of v that may

appear in $\mathcal{T}_{i+1}, \dots, \mathcal{T}_p$ are dummies. For example, if $e = \{u, v\} \in E(G)$, we think of e joining the *terminal/principal* copies of u and v in \mathcal{T} .

DEFINITION 4.1. Suppose $\tau_i \in \mathcal{T}_i$ and $\tau_{i'} \in \mathcal{T}_{i'}, i \leq i'$. We say τ_i is a descendant of $\tau_{i'}$ if a connected component of $G - B_{i'}$ contains $V(\tau_{i'})$ and at least one vertex of $V(\tau_i)$.

Observe that if $V(\tau_i) \cap B_{i'} = \emptyset$, then τ_i can only have one ancestor at level i' ; if it had two distinct ancestors, then they would be connected by a path in $G - B_{i'}$, contradicting Theorem 3.3. Unfortunately, it seems that $V(\tau_i)$ can intersect $B_{i'}$, so in general the ancestry relation between trees in \mathcal{T} induces a $(p+1)$ -level *dag*, not a rooted tree. Algorithmically it is much easier to deal with trees rather than dags. For this reason we define a variant hierarchy \mathcal{C} that is more structured. Both \mathcal{C} and \mathcal{T} are used by our data structures.

DEFINITION 4.2. Define \mathcal{C}_i to be the set of connected components of $G - (B_i \cup B_{i+1} \cup \dots \cup B_{p-1})$ containing at least one B_{i-1} (terminal) vertex. Suppose $\gamma_i \in \mathcal{C}_i$ and $\gamma_{i'} \in \mathcal{C}_{i'}$, where $i \leq i'$. We say γ_i is a descendant of $\gamma_{i'}$, written $\gamma_i \preceq \gamma_{i'}$, if $V(\gamma_i) \cap V(\gamma_{i'}) \neq \emptyset$.

Lemma 4.3 identifies the critical properties of $\{\mathcal{C}_i\}$ used by our algorithm.

LEMMA 4.3. Consider the hierarchy of components $\{\mathcal{C}_i\}_{i \in [0, p]}$.

1. Each $\gamma \in \mathcal{C}_i$ has at most one ancestor in $\mathcal{C}_{i'}$ for each $i' \in [i, p]$.
2. $V(\gamma) \subseteq V(\gamma')$ for each $\gamma \preceq \gamma'$.
3. If $\{u, v\} \in E$ and $u \in V(\gamma), v \in V(\gamma')$, then $\gamma \preceq \gamma'$ or $\gamma' \prec \gamma$.
4. If $\gamma \in \mathcal{C}_i$, the terminals $V(\gamma) \cap B_{i-1}$ are contained in a single tree in \mathcal{T}_i , denoted $\tau(\gamma)$.

Proof. For part 1, note that any two distinct components $\gamma', \gamma'' \in \mathcal{C}_i$ have $V(\gamma') \cap V(\gamma'') = \emptyset$. Since, by construction, $V(\gamma) \cap (B_{i'} \cup \dots \cup B_{p-1}) = \emptyset$, γ cannot share vertices with both γ' and γ'' . We now turn to part 2. Suppose $\gamma \in \mathcal{C}_i, \gamma' \in \mathcal{C}_{i'}$ with $i < i'$. If γ and γ' share one vertex, then $V(\gamma) \subset V(\gamma')$ since γ is connected and $V(\gamma) \cap (B_{i'} \cup \dots \cup B_{p-1}) = \emptyset$. If part 3 were false, then γ and γ' would be unrelated. Let γ'' be the ancestor of γ at the same level as γ' , so γ', γ'' are two distinct components in some \mathcal{C}_i . Part 2 implies $u \in V(\gamma'')$, meaning γ' and γ'' are joined by an edge $\{u, v\}$ and are therefore not distinct components in \mathcal{C}_i . For part 4, consider a tree $\tau \in \mathcal{T}_i = T_i - B_i$. By Theorem 3.3, τ spans the terminals (B_{i-1} -nodes) in a connected component of $G - B_i$. A $\gamma \in \mathcal{C}_i$ represents a connected component in $G - (B_i \cup \dots \cup B_{p-1})$, so if $V(\gamma)$ intersects $V(\tau)$ at one terminal, every terminal of $V(\gamma)$ must be contained in $V(\tau)$. \square

Lemma 4.3.1 (unique ancestors) shows that the ancestry relationship on the components of $\mathcal{C}_0, \dots, \mathcal{C}_p$ can be succinctly encoded as a forest of rooted trees. Let \mathcal{C} be the *component hierarchy* defined by the \prec relation. The nodes of \mathcal{C} are in one-to-one correspondence with the components of $\mathcal{C}_0, \dots, \mathcal{C}_p$, where \mathcal{C}_0 form the leaves of \mathcal{C} . Slightly abusing notation, we shall say “ $\gamma \in \mathcal{C}$ ” to mean that γ is a node in \mathcal{C} or that γ is a component in some \mathcal{C}_i .

4.1. Stocking the low degree hierarchy. Our goal is to supplement \mathcal{C} and \mathcal{T} with useful data structures that allow us to reconnect the graph after a set of vertices fail. Recall that \mathcal{T} is composed of trees with maximum degree at most 4. If a single tree $\tau \in \mathcal{T}$ experiences the failure of some vertex set $D \subset V$, we can find *individual* edges that reconnect the subtrees of $\tau - D$ using $O(|D|^2)$ 2D range queries (Theorem 2.1). However, individual edges are, in general, insufficient to reconnect the

subtrees. There could be long paths that go through vertices that appear in ancestors or descendants of τ in \mathcal{T} . In order to quickly detect the existence of these paths we follow an idea from [22] and introduce *artificial* edges that capture connectivity via paths. We do not want to add too many artificial edges, for two reasons. First, they take up space, which we want to conserve, and second, after deleting vertices from the graph the *validity* of many artificial edges may be cast into doubt. Any invalid artificial edges must be ignored when reestablishing connectivity, so it is important that the algorithm not encounter too many of these edges. Before saying exactly how artificial edges are added we must introduce the concept of a d_\star -adjacency list. Recall that d_\star is the maximum number of vertex failures.

DEFINITION 4.4. *Let $L = (v_1, v_2, \dots, v_r)$ be a list of vertices and $d_\star \geq 1$ be an integer. The d_\star -adjacency edges $\Lambda_{d_\star}(L)$ connect all vertices at distance at most $d_\star + 1$ in the list L :*

$$\Lambda_{d_\star}(L) = \{\{v_i, v_j\} \mid 1 \leq i < j \leq r \text{ and } j - i \leq d_\star + 1\}.$$

LEMMA 4.5. *The following properties hold for any vertex list L :*

1. $\Lambda_{d_\star}(L)$ contains fewer than $(d_\star + 1)|L|$ edges.
2. If a set D of at most d_\star vertices are removed from L , then the subgraph of $\Lambda_{d_\star}(L)$ induced by $L - D$ remains connected.
3. Suppose L is partitioned into consecutive sublists L_1 and L_2 . Then at most $O(d_\star^2)$ edges from $\Lambda_{d_\star}(L)$ cross the partition (L_1, L_2) .

Proof. Part 1 is trivial, as is 2, since each pair of consecutive undeleted vertices is at distance at most $d_\star + 1$, and therefore adjacent. Part 3 is also trivial: the number of edges connecting any prefix and suffix of L is at most $(d_\star + 1)(d_\star + 2)/2$. \square

Fix a $\gamma_i \in \mathcal{C}_i$ and let $\gamma_{i+1}, \dots, \gamma_p$ be its ancestors in \mathcal{C} . Recall that the terminals of γ_i are contained in a single tree $\tau(\gamma_i) \in \mathcal{T}_i$. The mapping τ is not necessarily injective: one tree in \mathcal{T}_i could be the host for many components in \mathcal{C}_i . Define $A(\gamma_i, \gamma_j)$ to be a list of the terminals in $V(\gamma_j)$ that are adjacent to at least one vertex in $V(\gamma_i)$, listed according to an Euler tour $\text{Euler}(\tau(\gamma_j))$. (Recall that the terminals in $V(\gamma_j)$ are exactly those vertices in $V(\gamma_j) \cap (B_{j-1} - (B_j \cup \dots \cup B_{p-1}))$.) Let $A(\gamma_i)$ be the concatenation of $A(\gamma_i, \gamma_{i+1}), \dots, A(\gamma_i, \gamma_p)$. We interpret elements of $A(\gamma_i)$ as the terminal copies of vertices in \mathcal{T} .

DEFINITION 4.6. *The multigraph H is on the vertex set of \mathcal{T} . For each $\{u, v\} \in E$, H contains an original edge connecting the terminal copies of u and v . For each component $\gamma \in \mathcal{C}$, H includes $\Lambda(\gamma) \stackrel{\text{def}}{=} \Lambda_{d_\star}(A(\gamma))$. Each edge in H is labeled with its provenance: either original or the name of a γ if it appears in $\Lambda(\gamma)$. Note that H may contain multiple edges with the same endpoints but with different provenance.*

Lemma 4.7 exhibits the two salient properties of $\Lambda(\gamma)$: that it encodes useful connectivity information and that it is economical to effectively destroy $\Lambda(\gamma)$ when it is no longer valid, often in time sublinear in $|\Lambda(\gamma)|$.

LEMMA 4.7. *Consider a $\Lambda(\gamma_i) \subset E(H)$.*

1. *Suppose $d \leq d_\star$ vertices fail, none of which are in $V(\gamma_i)$, and suppose u and v are in components of ancestors of γ_i and are each adjacent to at least one vertex in $V(\gamma_i)$. Then u and v remain connected in the original graph and remain connected in H .*
2. *Suppose the proper ancestors of γ_i are $\gamma_{i+1}, \dots, \gamma_p$ and a total of f edges are removed from $\tau(\gamma_{i+1}), \dots, \tau(\gamma_p)$, breaking their Euler tours into intervals*

I_1, \dots, I_{p-i+2f} . Then at most $O(d_*^2(p+f))$ edges of $\Lambda(\gamma_i)$ connect distinct intervals $I_j, I_{j'}$.

Proof. For part 1, the vertices u and v are connected in the original graph because they are each adjacent to vertices in $V(\gamma_i)$ and, absent any failures, all vertices in $V(\gamma_i)$ remain connected. By Definition 4.6, u and v appear in $\Lambda(\gamma_i)$, and, by Lemma 4.5, $\Lambda(\gamma_i)$ remains connected after the removal of any d vertices. Turning to part 2, recall from Definition 4.6 that $A(\gamma_i)$ was the concatenation of $A(\gamma_i, \gamma_{i+1}), \dots, A(\gamma_i, \gamma_p)$ and each $A(\gamma_i, \gamma_{i'})$ was ordered according to an Euler tour of $\tau(\gamma_{i'}) \in \mathcal{T}_{i'}$. Removing f edges from $\tau(\gamma_{i+1}), \dots, \tau(\gamma_p)$ separates their Euler tours (and, hence, the lists $\{A(\gamma_i, \gamma_{i'})\}_{i'}$) into at most $2f+p-i$ intervals. By Lemma 4.5 at most $(2f+p-i) \cdot O(d_*^2)$ edges from $\Lambda(\gamma_i)$ connect distinct intervals. In other words, in order to “logically” delete $\Lambda(\gamma_i)$ it suffices to delete $O(d_*^2(p+f))$ edges from $\Lambda(\gamma_i)$ since all remaining edges do not add to the connectivity of the remaining graph. \square

We apply Theorem 2.1 and generate an ET-structure $\mathbf{ET}(H, \mathcal{T})$ for H . Lemma 4.8 bounds the space for the overall data structure.

LEMMA 4.8. *Given a graph G with m edges, n vertices, and a parameter $d_* \geq 1$, the d_* -failure connectivity oracle consists of $\mathcal{C}, \mathbf{ET}(H, \mathcal{T})$ and various linear space data structures supporting navigation around \mathcal{C} . The space required by the oracle is $O(d_* m \log n \log \log n)$ or $O(dm \log n)$, depending on the 2D range searching structure used in $\mathbf{ET}(H, \mathcal{T})$, and its construction time is $O(mn \log n)$.*

Proof. The number of vertices in H is at most $(p+1)n$, n per \mathcal{T}_i . (This is a pessimistic bound. We are unable to conceive of any graph G for which this is achieved.) The number of original edges in H is m . Each original edge contributes a vertex to at most p lists $A(\gamma)$, and each member of $A(\gamma)$ contributes at most $d_* + 1$ edges to $\Lambda(\gamma)$. The number of vertices and edges in H is therefore at most $m + (p+1)n + p(d_* + 1)m = O(d_* m \log n)$. By Theorem 2.1, each edge in H contributes $O(\log \log n)$ or $O(1)$ space to $\mathbf{ET}(H, \mathcal{T})$. Regarding construction time, by Theorem 3.3 the time to compute (T_0, B_0) is $O(mn \log n)$, and more generally, the time to compute (T_{i+1}, B_{i+1}) is $O(m|B_i| \log |B_i|)$ time, where $|B_i| < n/(s-2)^i = n/2^i$ decays geometrically with i . Thus, the total time to compute \mathcal{T} and \mathcal{C} is $O(mn \log n)$. \square

5. Recovery from failures. In this section we describe how, given a set of $d \leq d_*$ failed vertices, the data structure can be updated in time $\tilde{O}(d^2 d_*^2)$ such that connectivity queries can be answered in $O(d)$ time. Section 5.1 gives the algorithm to delete failed vertices and section 5.2 gives the query algorithm and proof of correctness. In section 6 we describe several ways to improve the update time to $\tilde{O}(d^3)$.

5.1. Deleting failed vertices.

Let $D \subset V$ be the set of d failed vertices.

Step 1. Begin by marking any $\gamma \in \mathcal{C}$ affected if $V(\gamma) \cap D \neq \emptyset$, and mark the corresponding tree $\tau(\gamma) \in \mathcal{T}$ affected as well. For each affected $\tau(\gamma)$, mark each D -node and its incident tree edges as deleted. This breaks up $\tau(\gamma)$ into *affected subtrees*, which must be reconnected, if possible.

LEMMA 5.1. *The number of affected trees is at most $d(p+1)$. The number of affected subtrees is at most $4d(p+1)$.*

Proof. By Lemma 4.3, any $u \in D$ appears in at most $p+1$ components of \mathcal{C} . Since all failed vertices have degree at most $s = 4$ in the \mathcal{T} trees in which they appear, there are at most $4d(p+1)$ affected subtrees. \square

Recall from the discussion above that if γ is affected, then $V(\gamma)$ contains failed

vertices and the connectivity provided by $\Lambda(\gamma)$ is presumed invalid. By Lemma 4.7 we can logically delete $\Lambda(\gamma)$ by ignoring $O(d_*^2)$ edges for *each* of $O(pd)$ breaks in the list $A(\gamma)$. Since there are at most $O(pd)$ affected (sub)trees, the number of edges that need to be ignored is $O((pd)^2 d_*^2)$. Let H' denote the graph H with these $O((pd)^2 d_*^2)$ edges removed.

Step 2. We now attempt to reconnect all affected subtrees using valid edges, i.e., those in H' . Let R be a graph whose vertex set $V(R)$ represents the $O(pd)$ affected subtrees such that $\{t_1, t_2\} \in E(R)$ if t_1 and t_2 are connected by an edge in H' . Using the structure $\mathbf{ET}(H, \mathcal{T})$ (see Theorem 2.1) we populate the edge set of R in time $O(|V(R)|^2 q + (pd)^2 d_*^2)$, where $q = \log \log n$ or $\log^\epsilon n$, depending on the space of the 2D range structure [21]. For each 2D range query, we halt the enumeration of points/edges as soon as an H' -edge is reported. Recall that a point/edge is tagged with its provenance, so we can check in $O(1)$ time whether it came from an affected $\Lambda(\gamma)$ and must be discarded. Since $|V(R)| = O(pd)$ and $p < \log n$, the time to perform these queries is $O(d^2(q + d_*^2) \log^2 n)$. In $O(|E(R)|) = O((pd)^2)$ time we determine the connected components of R .

This concludes the deletion algorithm. The running time is dominated by Step 2.

5.2. Answering a connectivity query. To answer a connectivity query between u and v we first check to see if there is a path between them that avoids affected trees, then consider paths that intersect one or more affected trees.

Step 1. We first find the components in \mathcal{C} containing u and v as terminals; let them be $\gamma(u)$ and $\gamma(v)$. If $\gamma(u)$ is unaffected, let $\hat{\gamma}(u)$ be the most ancestral unaffected ancestor of $\gamma(u)$, and let $\hat{\gamma}(v)$ be defined analogously. If $\hat{\gamma}(u), \hat{\gamma}(v)$ exist and are equal, then $V(\hat{\gamma}(u))$ contains u and v but no failed vertices. If this is the case we declare u and v *connected* and stop.

We can find $\hat{\gamma}(u)$ and $\hat{\gamma}(v)$ in $O(\log p) = O(\log \log n)$ time using a binary search over the ancestors of $\gamma(u)$ and $\gamma(v)$. Alternatively, we can find them in time $O(\log d)$, independent of n , using relatively simple data structures. Fix any postordering of the nodes of \mathcal{C} . Find the predecessor γ_{pred} and successor γ_{succ} of $\gamma(u)$ among all components whose *terminal set* contains a D -vertex. There are at most d such nodes, so the cost to find them is $O(\log d)$ via binary search. Let $\gamma_{\text{pred}}^{\text{lca}}, \gamma_{\text{succ}}^{\text{lca}}$ be the least common ancestors of $\gamma(u)$ and $\gamma_{\text{pred}}, \gamma_{\text{succ}}$, respectively. Without loss of generality suppose $\gamma_{\text{pred}}^{\text{lca}}$ is closer to $\gamma(u)$. Since $V(\gamma_{\text{pred}}^{\text{lca}}) \cap D \neq \emptyset$, $\gamma_{\text{pred}}^{\text{lca}}$ is affected. If $\gamma_{\text{pred}}^{\text{lca}}$ is at depth k from its root in \mathcal{C} , the node $\hat{\gamma}(u)$ that we are looking for is the ancestor of $\gamma(u)$ at depth $k+1$. Refer to [12, 13] for linear space data structures for least common ancestor and level ancestors.

Step 2. We now try to find vertices u' and v' in affected subtrees that are connected to u and v , respectively. If $\gamma(u)$ is affected, then $u' = u$ clearly suffices, so we only need to consider the case when $\gamma(u)$ is unaffected and $\hat{\gamma}(u)$ exists. Recall from Definition 4.6 that $A(\hat{\gamma}(u))$ is the list of terminals in proper ancestors of $\hat{\gamma}(u)$ that are adjacent to some vertex in $V(\hat{\gamma}(u))$. We scan $A(\hat{\gamma}(u))$ looking for *any* nonfailed vertex u' adjacent to $V(\hat{\gamma}(u))$. Since $V(\hat{\gamma}(u))$ is unaffected, u is connected to u' , and since all of $\hat{\gamma}(u)$'s proper ancestors are affected, u' must appear in an affected subtree in \mathcal{T} . Since there are at most d failed vertices we must inspect at most $d+1$ elements of $A(\hat{\gamma}(u))$. This takes $O(d)$ time to find u' and v' , if they exist. If one or both of u' and v' does not exist we declare u and v *disconnected* and stop.

Step 3. We have the terminal copies of u' and v' in \mathcal{T} . We find the affected subtrees t'_1 and t'_2 containing u' and v' in $O(\min\{\frac{\log \log n}{\log \log \log n}, \frac{\log d}{\log \log n}\})$ time via predecessor search over the left endpoints of the Euler-tour intervals that remain after deleting D

and their incident tree edges. Note that t'_1 and t'_2 are vertices in R , from Step 2 of the deletion algorithm. We declare u and v to be connected iff t'_1 and t'_2 are in the same connected component of R . This takes $O(1)$ time.

LEMMA 5.2. *The query algorithm correctly determines whether u and v are connected in $G - D$, in $O(d)$ time.*

Proof. If the query algorithm halts in Step 1 it is because both u and v are in the unaffected component $\hat{\gamma}(u)$, and since $V(\hat{\gamma}(u)) \cap D = \emptyset$, all vertices in $\hat{\gamma}(u)$ are still connected. If the query algorithm halts in Step 2 it is because $u \in V(\hat{\gamma}(u))$, $v \notin V(\hat{\gamma}(u))$, and $A(\hat{\gamma}(u)) - D = \emptyset$. Since $A(\hat{\gamma}(u))$ contains all vertices adjacent to $\hat{\gamma}(u)$ there can be no path from u to v in $G - D$.

At Step 3 we have discovered u', v' such that u is connected to u' , which appears as a terminal in some affected subtree t'_1 and similarly for v, v' , and t'_2 . Since t'_1, t'_2 are vertices in R , the correctness of the query algorithm hinges on whether the graph R correctly represents the connectivity between affected subtrees.

We first argue that if t'_1 and t'_2 are connected by a path in R , then they are connected in $G - D$. Each edge on this path is either an original edge or a $\Lambda(\gamma)$ -edge for some unaffected γ . All original edges not incident to D are still valid and each $\Lambda(\gamma)$ edge can, when γ is unaffected, be replaced by a path in $G - D$ using intermediate nodes in $V(\gamma)$.

We now argue that if $P = (u' = u_0, u_1, \dots, u_{|P|} = v')$ is a $u'-v'$ path in $G - D$, that there exists a $t'_1-t'_2$ path in R . Partition $P = P_1 P_2 \dots P_\omega$ into maximal subpaths ($P_i = (u_{a(i)}, \dots, u_{b(i)})$) such that $V(P_i)$ is either (i) contained in a single affected subtree or (ii) contained in $V(\gamma)$ for some unaffected $\gamma \in \mathcal{C}$. Observe that because of the maximality criterion, no two type (ii) subpaths can be adjacent. Since P_1 and P_ω contain u' and v' , they must be type (i) subpaths. We want to show that all type (i) subpaths are connected in R by considering how consecutive type (i) subpaths could be connected by valid edges in H' . (Recall that H' is H after deleting all $\Lambda(\gamma)$ edges for affected $\gamma \in \mathcal{C}$.) There are two cases to consider. \square

Case 1. Suppose P_i and P_{i+1} are type (i) subpaths. Then $\{u_{b(i)}, u_{a(i+1)}\}$ is an original edge in H' , so it or some other edge will be discovered that puts the affected subtrees of P_i and P_{i+1} in the same connected component in R .

Case 2. Suppose P_i and P_{i+2} are type (i) subpaths, but P_{i+1} is a type (ii) subpath. Let $\gamma \in \mathcal{C}$ be the component for which $V(P_{i+1}) \subset V(\gamma)$, so $u_{b(i)}, u_{a(i+2)} \notin V(\gamma)$. It must be that $u_{b(i)}, u_{a(i+2)} \in A(\gamma)$, and since $\Lambda(\gamma)$ remains connected after any d vertex deletions, $u_{b(i)}$ and $u_{a(i+2)}$ are connected by a path in $\Lambda(\gamma) - D$. All the $\Lambda(\gamma) - D$ edges straddling two affected subtrees are eligible to be discovered when populating the edge set of R , so the affected subtrees of P_i and P_{i+2} must be in the same connected component in R .

6. Improving the update time. In this section we present not one, not two, but three different methods to reduce the update time from $\tilde{O}(d^2 d_\star^2)$ to $\tilde{O}(d^3)$. Each of the three methods uses a different, more sophisticated orthogonal range searching structure. In section 6.1 we show how $\tilde{O}(d^3)$ time can be achieved with a 2D *colored* (a.k.a. *categorical*) range searching structure [64]. Section 6.2 uses a 2D range *counting* [23] data structure, and section 6.3 uses a 3D range *emptiness* data structure [21]. The method of section 6.3 was suggested to us by Shiri Chechik.

6.1. Method 1: Colored range searching. We use the following theorem from Larsen and van Walderveen [64].

THEOREM 6.1 (see [64]). *Given a multiset $P \subset [U] \times [U]$ of n points and coloring $\phi : P \rightarrow \mathbb{N}$, there is a data structure occupying space $O(n \log n)$ that answers the following type of query. Given x, x', y, y' , report the color set $\Phi = \{\phi(p) \mid p \in P \cap [x, y] \times [x', y']\}$. The query time is $O(\log \log U + |\Phi|)$.*

Assign each component $\gamma \in \mathcal{C}$ a distinct color $\phi(\gamma) \in \{1, \dots, |\mathcal{C}|\}$. Recall that each edge in H is tagged with its provenance. All *original* edges receive color zero and all $\Lambda(\gamma)$ edges receive color $\phi(\gamma)$. Each 2D range query now returns a list of colors in the query rectangle. We halt the search the moment it returns color 0 (an original edge), or the color of any unaffected component. Since there are at most $d(p+1)$ affected components, each of the $O((pd)^2)$ 2D range queries is halted after time $O(\log \log n + pd)$.

Using Method 1, the space of the resulting d_* -failure connectivity oracle becomes $O(d_* m \log^2 n)$ and the update time $O((pd)^3) = O(d^3 \log^3 n)$.

6.2. Method 2: 2D range counting. We use the following theorem of JaJa, Mortensen, and Shi [58].

THEOREM 6.2 (see [58]). *Given a multiset $P \subset [U] \times [U]$ of n points there is an $O(n)$ -space data structure answering the following type of query in $O(\log n / \log \log n)$ time. Given x, x', y, y' , report the number $k = |P \cap [x, y] \times [x', y']|$.*

Consider an affected component γ_i and recall that its adjacency list $A(\gamma_i)$ is the concatenation of $A(\gamma_i, \gamma_{i+1}), \dots, A(\gamma_i, \gamma_p)$, where $\gamma_{i+1}, \dots, \gamma_p$ are its ancestors in \mathcal{C} . The 2D range queries that are influenced by $\Lambda(\gamma_i)$ involve two trees, say, $\tau = \tau(\gamma_j)$ and $\tau' = \tau(\gamma_{j'})$, where $i < j \leq j' \leq p$. Each query is the product $Q = I \times I'$ of an interval $I \subset \text{Euler}(\tau)$ and another $I' \subset \text{Euler}(\tau')$. Given the indices of the first and last elements of $A(\gamma_i, \gamma_j) \cap I$ and $A(\gamma_i, \gamma_{j'}) \cap I'$, we can determine in $O(1)$ time how many $\Lambda(\gamma_i)$ edges (points) appear in Q . Call these *affected points*. For each affected component γ and each query Q to be performed by the update algorithm, we calculate the number of affected $\Lambda(\gamma)$ points in Q . This takes time $O(pd \cdot (pd)^2) = O(d^3 \log^3 n)$.

Let k_Q be the *total* number of affected points in Q , over all affected γ . In $O(\log n / \log \log n)$ time we compute the number k of points in Q . If $k = k_Q$, then there are no unaffected points in Q , and if $k > k_Q$ we deduce that there is an unaffected point (a valid edge connecting the two intervals). The total time for all $O((pd)^2)$ queries is therefore $O(d^2 \log^3 n / \log \log n)$ time. The bottleneck in this approach is computing the set $\{k_Q\}$ of critical thresholds.

Using Method 2 the space of our d_* -failure connectivity oracle is $O(d_* m \log n)$ and the update time is $O(d^3 \log^3 n)$.

6.3. Method 3: 3D range emptiness. We use the following theorem of Chan, Larsen, and Pătrașcu [21].

THEOREM 6.3 (see [21]). *Given a set $P \subset [U] \times [U] \times [U]$ of n points there is an $O(n \log^{1+\epsilon} n)$ -space data structure answering queries of the following type, in $O(\log \log U)$ time. Given x, x', x'', y, y', y'' , determine if $P \cap [x, y] \times [x', y'] \times [x'', y''] = \emptyset$.*

List the nodes in \mathcal{C} as $\{\gamma_1, \dots, \gamma_{|\mathcal{C}|}\}$. Suppose that $\{u, v\}$ is an original edge in H and τ_u, τ_v are the trees in \mathcal{T} containing the terminal copies of u and v , where u appears at position i in $\text{Euler}(\tau_u)$ and v appears at position j of $\text{Euler}(\tau_v)$. Rather than map $\{u, v\}$ to the point (i, j) in the 2D structure of $\text{ET}(H, \mathcal{T})$ we map it to the 3D point $(i, j, 0)$. If $\{u, v\}$ is an edge of $\Lambda(\gamma_k)$ we map it to the point (i, j, k) .

Let $(\gamma_{k_1}, \gamma_{k_2}, \dots, \gamma_{k_{d(p+1)}})$ be the affected components and Q be a 2D query performed by the update algorithm. We are interested in knowing whether there is a point whose first two coordinates are in Q and whose third coordinate is *not* a member of $\{k_1, \dots, k_{d(p+1)}\}$. Thus the 2D query Q can be reduced to $d(p+1)$ 3D emptiness queries $Q \times [0, k_1]$, $Q \times (k_1, k_2)$, and so on. Each 3D query is answered in $O(\log \log n)$ time, so the total update time is $O(d^3 \log^3 n \log \log n)$.

With the current state-of-the-art range searching data structures [21, 23, 58, 64], Method 2 is always strictly superior to Methods 1 and 3 in update time or space or both. Method 2 also leaves the most room for improvement since the bottleneck is not range counting queries per se but computing the critical thresholds $\{k_Q\}$ for the queries.

7. A Monte Carlo connectivity oracle. In the extended abstract [41] of this work, we claimed a Monte Carlo d_* -failure connectivity oracle with near optimum space $\tilde{O}(m)$, update time $\tilde{O}(d^2)$, and query time $O(d)$. The data structure described in [41, section 7] functions correctly but occupies space $\Omega(d_* m)$, not $\tilde{O}(m)$.⁷ In this section we present the first Monte Carlo connectivity oracle that achieves the claimed specifications of [41, section 7]. Our data structure is inspired by the graph sketching techniques of Ahn, Guha, and McGregor [3] and Kapron, King, and Mountjoy [61] but applies the ideas differently. In particular, by using vertex sampling rather than edge sampling, we show that it is possible to form sketches of complete bipartite subgraphs “on the fly” using minimal storage.

Let us first take one step back and discuss why achieving near-linear space is difficult. Recall from section 4 that $A(\gamma)$ is a list of all vertices adjacent to the component γ , and D is the set of failed vertices. So long as γ suffers no vertex failures, we want the subgraph induced by $A(\gamma) - D$ to remain connected. On the other hand, if γ does suffer a vertex failure, we want to be able to efficiently dispose of any suspect edges induced by $A(\gamma)$. Adding the d_* -adjacency edges $\Lambda_{d_*}(A(\gamma))$ solved both problems, but with some significant losses in efficiency. The space required to store $\Lambda_{d_*}(A(\gamma))$ is $\Omega(d_* \cdot |A(\gamma)|)$, and in order to ignore suspect edges, the update times for our deterministic solutions are $\Omega(d^3)$. It seems very difficult to avoid an $\Omega(d_*)$ factor overhead in space. Indeed, if $|A(\gamma)| = d + 2$ and all but two random elements of $A(\gamma)$ fail, we want to be able to quickly determine that those last two elements are still connected. In this situation, is it possible to avoid storing a clique on $A(\gamma)$?

By introducing Monte Carlo randomness, we are able to save both space and time simultaneously. The high-level ideas are as follows.

- Rather than use a d_* -adjacency list $\Lambda_{d_*}(A(\gamma))$ to maintain connectivity information within $A(\gamma)$, we pick a random subset $B(\gamma) \subseteq A(\gamma)$ and represent the complete bipartite graph $A(\gamma) \times B(\gamma)$.⁸ The total number of edges in the multigraph, over all γ (i.e., $\sum_{\gamma} |A(\gamma) \times B(\gamma)|$), could be quite large. One property of our graph sketch is that the space is actually proportional to the number of *distinct* edges in $\bigcup_{\gamma} A(\gamma) \times B(\gamma)$, not counting multiplicity, which is just $\tilde{O}(m)$.
- Observe that a complete bipartite graph $A(\gamma) \times B(\gamma)$ preserves the connectivity on $A(\gamma)$ iff $B(\gamma) - D \neq \emptyset$, i.e., if at least one nonfailed vertex is (randomly)

⁷After showing that a certain data structure $\mathcal{S}[\tau, \tau']$ occupies $O(m \log^4 n)$ space, we stated [41, p. 505], “By a similar analysis, the space for $\mathcal{S}[v, \tau]$ and $\mathcal{S}[\gamma, \tau]$ are also upper bounded by $O(m \log^4 n)$.” Unfortunately, we see no way to store $\mathcal{S}[v, \tau]$ in less than $\tilde{O}(d_* m)$ space.

⁸Throughout this section we abusively write $A \times B$ to be the set of *undirected* edges $\{u, v\}$ with $u \in A, v \in B$.

selected for inclusion in $B(\gamma)$. In some situations we can guarantee that this property holds, w.h.p. For example, if $|A(\gamma)| \geq 2|D|$ and $\Omega(\log n)$ vertices are included in $B(\gamma)$, then w.h.p. one vertex in $A(\gamma) - D$ is included in $B(\gamma)$. However, in general it is impossible to guarantee this property w.h.p., short of setting $B(\gamma) = A(\gamma)$. Our solution depends on a particular accounting scheme used in choosing the $B(\cdot)$ sets. We process the components $\gamma_1, \dots, \gamma_{|\mathcal{C}|}$ in an arbitrary sequential order. When it is γ_j 's turn we examine the subgraph induced by $A(\gamma_j)$ and choose $|B(\gamma_j)|$ such that the expected number of *new* edges contributed by $A(\gamma_j) \times B(\gamma_j)$ is $\tilde{O}(|A(\gamma_j)|)$. Every time a *new* edge is added we label it with its *owner* “ γ_j .” These labels are not simply used for accounting. We prove that for any failed set D , at least one of the following two events occurs, w.h.p.: (i) either $B(\gamma_j) - D \neq \emptyset$ (and connectivity information via γ is maintained) or (ii) at least $|A(\gamma_j)|$ pairs in $\binom{D}{2}$ are owned by γ_j . Event (ii) is a happy outcome because it *reveals* a small number of components whose connectivity information was not maintained as in (i), and those components can be processed separately in $O(|D|^2)$ total time.

- When γ suffers a vertex failure, the entire bipartite graph contributed by γ , namely $A(\gamma) \times B(\gamma)$, is suspect. Our sketch has the property that complete bipartite subgraphs of $A(\gamma) \times B(\gamma)$ can be efficiently generated by a data structure occupying space $\tilde{O}(|A(\gamma)| + |B(\gamma)|)$ rather than $O(|A(\gamma)| \cdot |B(\gamma)|)$. Thus, it is efficient to subtract from all relevant graph sketches the contribution of edges from affected components.

Organization of section 7. In section 7.1 we show how the B -sets are chosen and analyze their properties. In section 7.2 we introduce two sketches. Original graph edges are sketched exactly as in Kapron, King, and Mountjoy [61], but “artificial” edges in $A(\gamma) \times B(\gamma)$ are sketched in a new way. The total size of all sketches and their attendant data structures is $O(m \log^6 n)$. In section 7.3 we show how to handle a batch of d vertex failures in $O(d^2 \log^6 n)$ time and subsequently answer connectivity queries in $O(d)$ time. In section 7.4 we observe that it is often unnecessary to explicitly form complete graph sketches. This allows us to reduce the update time of the best d -edge failure oracles [73, 61, 51, 40] to $O(d \log d \log \log n)$ expected time and reduce the update time of section 7.3 to $O(d^2 \log d \log^2 n \log \log n)$ expected time.

7.1. The B -sets and their properties. Recall that the “artificial” edges associated with $\gamma \in \mathcal{C}$ will be a complete bipartite graph $A(\gamma) \times B(\gamma)$. The algorithm for generating $B(\gamma) \subseteq A(\gamma)$ is as follows. Choose an arbitrary order $\gamma_1, \dots, \gamma_{|\mathcal{C}|}$ of the components. Each pair $\{u, v\} \in \binom{V}{2}$ (regardless of whether it is in E or not) is initially *unlabeled* and may become labeled as we proceed. After $B(\gamma_1), \dots, B(\gamma_{j-1})$ have been selected we consider the pairs on elements of $A(\gamma_j)$. Let $n' = |A(\gamma_j)|$ and let $m' \in [0, \binom{n'}{2}]$ be the number of *unlabeled* pairs in $\binom{A(\gamma_j)}{2}$. The set $B(\gamma_j)$ is selected by sampling each vertex in $A(\gamma_j)$ independently with probability $\min\{1, \frac{n' \cdot c \ln n}{m'}\}$, where c is a sufficiently large constant that controls the error probability $n^{-\Omega(c)}$. Every unlabeled pair in the set $\{u, v\} \mid u \in A(\gamma_j), v \in B(\gamma_j)\}$ is now *owned* by γ_j and labeled “ γ_j .”

LEMMA 7.1. *For each $\gamma \in \mathcal{C}$, $O(|A(\gamma)| \log n)$ pairs are labeled “ γ ” in expectation. The total number of pairs in $\bigcup_{\gamma} A(\gamma) \times B(\gamma)$ is $O(m \log^2 n)$ in expectation.*

Proof. The probability that an unlabeled pair $\{u, v\}$ with $u, v \in A(\gamma)$ is labeled “ γ ” is exactly the probability that either u or v (or both) is selected for inclusion

in $B(\gamma)$. Recalling the definitions of n' and m' , the number of edges labeled γ is, by linearity of expectation, at most $m' \cdot (2n' \cdot c \ln n)/m' = O(|A(\gamma)| \log n)$.

Every pair $\{u, v\} \in \bigcup_{\gamma} A(\gamma) \times B(\gamma)$ must be owned by *some* component. By the first part of the lemma,

$$\left| \bigcup_{\gamma} A(\gamma) \times B(\gamma) \right| \leq \sum_{\gamma} O(|A(\gamma)| \log n) = O(m \log^2 n).$$

The last equality holds because each edge contributes one element to at most $p < \log n$ $A(\cdot)$ -lists. \square

By design, the B -sets are chosen to keep the total number of owned pairs $\tilde{O}(m)$. Lemma 7.2 indicates why this method of choosing B -sets is useful when vertices fail.

LEMMA 7.2. *Fix any γ and any set D of (failed) vertices such that $A(\gamma) - D \neq \emptyset$. With probability $1 - n^{-\Omega(c)}$, one of the following two events occurs.*

1. $B(\gamma) - D \neq \emptyset$.
2. *The number of pairs in $\binom{D}{2}$ owned by γ is at least $|A(\gamma)|$.*

Proof. Consider the moment in the algorithm just before $B(\gamma)$ is selected, and let n', m' be defined as usual. We consider two possible scenarios, depending on how many of the m' pairs are completely contained in D or straddle/lie outside of D .

Case I. At least $m'/2$ of the unlabeled pairs contain at least one vertex in $A(\gamma) - D$. There must be at least $(m'/2)/n'$ vertices in $A(\gamma) - D$, and each one is sampled with probability $\min\{1, cn' \ln n/m'\}$. The probability that some vertex in $A(\gamma) - D$ is sampled into $B(\gamma)$ is

$$1 - \left(1 - \min\left\{1, \frac{cn' \ln n}{m'}\right\}\right)^{\frac{m'}{2n'}} > 1 - n^{-c/2},$$

in which case part 1 of the lemma holds.

Case II. At least $m'/2$ of the unlabeled pairs are contained in $\binom{D}{2}$. We can assume without loss of generality that $m' \geq cn' \ln n$ for otherwise $B(\gamma) = A(\gamma)$ and part 1 of the lemma is already satisfied. Assign each unlabeled pair in $\binom{D}{2}$ to one of its endpoints, and let $\deg'(v)$ be the number of pairs assigned to v , so $\sum_v \deg'(v) = m'/2$. Partition the vertices into $\lfloor \log n' \rfloor + 1$ classes where class i contains those vertices for which $\deg'(v) \in [2^i, 2^{i+1})$. Let the sum of degrees in class i be $\epsilon_i(m'/2)$, i.e., $\sum_i \epsilon_i = 1$. The number of vertices in class i is at least $\epsilon_i m'/2^{i+2}$ since each accounts for at most 2^{i+1} distinct edges. The expected number of vertices in class i included in $B(\gamma)$ is therefore at least $\epsilon_i cn' \ln n/2^{i+2}$, and by a Chernoff bound, the probability that at least half the expected number are sampled is $1 - \exp(-\epsilon_i cn' \ln n/2^{i+5})$. If so, this contributes at least $\epsilon_i cn' \ln n/2^5$ pairs owned by γ . Call a class i *good* if $\epsilon_i cn' \ln n/2^{i+5} \geq (c/2^7) \ln n$, or equivalently, if $\epsilon_i \geq 2^{i-2}/n'$. The fraction of pairs contributed by *bad* classes is at most $\sum_{i=0}^{\lfloor \log n' \rfloor} 2^{i-2}/n' < 1/2$. Thus, with probability $1 - n^{-\Omega(c)}$, the number of unlabeled pairs in $\binom{D}{2}$ that are covered by $B(\gamma)$ -vertices in good classes (which become owned by γ) is at least $(1/2) \cdot cn' \ln n/2^5 > n'$. This satisfies part 2 of the lemma. \square

Remark 7.3. The proof of Case II of Lemma 7.2 is necessarily ad hoc. We are trying to lower bound a sum $X = X_1 + \dots + X_k$ of independent random variables, which seems to be well suited to some variant of the Azuma–Hoeffding inequality [43]. However, in our case $\mathbb{E}[X]$ is small but the variances $V[X_i]$ large. In this regime the standard concentration bounds do not offer strong enough guarantees.

7.2. Graph sketches. We use the graph sketch of Kapron et al. [61, 51, 81] to store original edges but develop a new sketch for artificial edges of the form $A(\gamma) \times B(\gamma)$. It is convenient to rename the vertex ids in $\{1, \dots, n\}$. For each $\tau \in \mathcal{T}$, the ids of the terminals in $V(\tau)$ occupy a contiguous interval of $[1, n]$, and moreover, their ids are consistent with the ordering of $\text{Euler}(\tau)$.

7.2.1. Sketching original edges. An edge $e = \{u, v\}$ is represented by the bit string $\langle e \rangle = \langle \min\{u, v\}, \max\{u, v\} \rangle$. For $i \in [0, \log m], j \in [1, c \log n]$, the edge sets $E = E_{0,j} \supseteq E_{1,j} \supseteq \dots \supseteq E_{\log m-1,j}$ are generated such that all edges are sampled for inclusion in $E_{i,j}$ independently with probability 2^{-i} . The sketch $\Upsilon^{E'}$ for an edge set $E' \subseteq E$ is a $\log m \times c \log n$ matrix in which

$$\Upsilon^{E'}(i, j) = \bigoplus_{e \in E' \cap E_{i,j}} \langle e \rangle.$$

That is, the (i, j) th entry contains the bitwise XOR of all edge names in $E' \cap E_{i,j}$. Clearly sketches are additive: for any E', E'' , $\Upsilon^{E' \oplus E''} = \Upsilon^{E'} \oplus \Upsilon^{E''}$. Lemma 7.4 illustrates why this sketch is useful for quickly finding edges crossing cuts.

LEMMA 7.4. Define E_u to be the edges incident to u . For any subset $S \subset V$, define $\Upsilon = \bigoplus_{u \in S} \Upsilon^{E_u}$ to be the componentwise XOR of all Υ^{E_u} sketches. For each j , there exists some i such that with constant probability $\Upsilon(i, j)$ is the name of some edge crossing the cut $(S, V - S)$.

Proof. Edges with two endpoints in S contribute nothing to Υ since $\langle e \rangle \oplus \langle e \rangle = \langle 0 \rangle$. Let i be such that the number of edges crossing the cut is between 2^i and $2^{i+1} - 1$. Then with constant probability, exactly one such edge is sampled for inclusion in $E_{i,j}$. \square

When a batch D of vertices fail we get a set $\{t_l\}$ of $O(|D| \log n)$ affected subtrees. For each t_l , we need to be able to obtain a sketch of all edges $\{u, v\}$ where $u \in t_l$, $v \in t_{l'}, l' \neq l$. The data structures \mathcal{V} and \mathcal{C} report sketches of edges incident to one vertex and one component, respectively.

$\mathcal{C}(\gamma, I)$: The input is a component $\gamma \in \mathcal{C}$ and an interval I of some $\text{Euler}(\tau)$, where τ could be equal to $\tau(\gamma)$. Define $E_{u,\gamma}$ to be the original edges joining u to the terminals of $V(\gamma)$. Report the sketch $\Upsilon = \bigoplus_{u \in I} \Upsilon^{E_{u,\gamma}}$.

$\mathcal{V}(v, I)$: The input is a vertex v and an interval I of some $\text{Euler}(\tau)$. Let $E_{v,I}$ be the original edges joining v to the terminals in I . Report the sketch $\Upsilon^{E_{v,I}}$.

LEMMA 7.5. The structures \mathcal{V}, \mathcal{C} occupy $O(m \log^2 n)$ space and answer queries in $O(\log^2 n)$ time.

Proof. First consider a fixed $v \in V(\gamma)$. Let $L_v = (v_1, \dots, v_{\deg(v)})$ be a list of v 's neighbors, in increasing order of vertex id. By how we chose the vertex id assignment, any interval I of some $\text{Euler}(\tau)$ corresponds to an interval of L_v . Let Υ^r be the sketch for the single edge $\{v, v_r\}$. In $O(\deg(v) \log^2 n)$ space we store all prefix sums $(\alpha_1, \dots, \alpha_{\deg(v)})$, where $\alpha_k = \bigoplus_{r \leq k} \Upsilon^r$. To answer a query $\mathcal{V}(v, I)$, we simply need to identify the sublist of $(v_1, \dots, v_{\deg(v)})$ covered by interval I , say it is (v_k, \dots, v_l) , and report $\alpha_l \oplus \alpha_{k-1}$ in $O(\log^2 n)$ time.

We now turn to \mathcal{C} . As before, let $L_\gamma = (v_1, v_2, \dots)$ be a list of all neighbors of terminals in $V(\gamma)$, listed in increasing order of vertex id, let Υ^r be the sample matrix for $E_{v_r, \gamma}$, and let $\beta_k = \bigoplus_{r \leq k} \Upsilon^r$. Suppose the query is $\mathcal{C}(\gamma, I)$. We do a binary search to find the sublist of (v_1, v_2, \dots) covered by I , then report the interval-sum in $O(\log^2 n)$ time by XORing two β -sketches.

Each original edge $\{u, v\}$ may contribute two $O(\log^2 n)$ -size sketches to \mathcal{V} and \mathcal{C} . The total space is therefore $O(m \log^2 n)$. \square

7.2.2. Sketching artificial edges. Artificial edges are encoded differently than original edges. Let $e = \{u, v\}$ be an artificial edge in $A(\gamma) \times B(\gamma)$. The encoding $\langle e \rangle = \langle u, v, \gamma \rangle$ puts $u \in A(\gamma)$ before $v \in B(\gamma)$, and includes the *provenance* identifier γ .⁹ Given a bit string $\langle u, v, \gamma \rangle$, we can easily verify whether it corresponds to a legitimate edge by checking whether $u \in A(\gamma), v \in B(\gamma)$.

The sketches for artificial edges are obtained via *vertex*-sampling rather than edge-sampling. For $i \in [0, \log n)$, and $j \in [1, c \log n]$, we choose sets $A_{i,j}, B_{i,j}, \mathcal{C}_{i,j}$ such that

$$\begin{aligned} V &= A_{0,j} \supseteq A_{1,j} \supseteq \cdots \supseteq A_{\log n-1,j}, \\ V &= B_{0,j} \supseteq B_{1,j} \supseteq \cdots \supseteq B_{\log n-1,j}, \\ \mathcal{C} &= \mathcal{C}_{0,j} \supseteq \mathcal{C}_{1,j} \supseteq \cdots \supseteq \mathcal{C}_{\log n-1,j}. \end{aligned}$$

Each $\gamma \in \mathcal{C}$ is included in $\mathcal{C}_{i,j}$ independently with probability 2^{-i} . Similarly, each $u \in V$ is included in $A_{i,j}$ and $B_{i,j}$ independently with probability 2^{-i} . Define $E_{i_a, i_b, i_c, j}$ to be the edge set

$$E_{i_a, i_b, i_c, j} = \{\langle u, v, \gamma \rangle \mid u \in A_{i_a, j}, v \in B_{i_b, j}, \gamma \in \mathcal{C}_{i_c, j}\}.$$

Let $\widehat{E} = E_{0,0,0,0}$ be the union of all edges contained in $A(\gamma) \times B(\gamma)$ over all $\gamma \in \mathcal{C}$.¹⁰ The sketch of $E' \subset \widehat{E}$ is a 4D matrix $\widehat{\Upsilon}^{E'}$, where

$$\widehat{\Upsilon}^{E'}(i_a, i_b, i_c, j) = \bigoplus_{e \in E' \cap E_{i_a, i_b, i_c, j}} \langle e \rangle.$$

Lemma 7.6 is the analogue of Lemma 7.4 for vertex-sampled sketches.

LEMMA 7.6. *Let E_u be the edges adjacent to u in \widehat{E} , and let $\widehat{\Upsilon}^{E_u}$ be the sketch for E_u . Suppose that for $S \subset V$, the cut $(S, V - S)$ is nonempty, and let $\widehat{\Upsilon} = \bigoplus_{u \in S} \widehat{\Upsilon}^{E_u}$ be the componentwise XOR of the sketches of S -vertices. For each j , with constant probability there exists i_a, i_b, i_c such that $\widehat{\Upsilon}(i_a, i_b, i_c, j)$ is the name of some edge crossing the cut $(S, V - S)$.*

Proof. In contrast to the proof of Lemma 7.4, there is not necessarily a *specific* triple (i_a, i_b, i_c) that satisfies the lemma; we only claim that one of the $O(\log^3 n)$ triples will work, with constant probability. Let $C \subseteq \mathcal{C}$ be the subset of components such that for each $\gamma \in C$, some edge of $A(\gamma) \times B(\gamma)$ crosses the cut. With constant probability, $|\mathcal{C}_{i_c, j} \cap C| = 1$, where $i_c = \lfloor \log |C| \rfloor$. Suppose that $\gamma \in C$ is the component isolated by $\mathcal{C}_{i_c, j}$. Let $A' \subseteq A(\gamma)$ be the subset of vertices adjacent to edges with provenance γ crossing the cut. Note that A' may include vertices on both sides of the cut. With constant probability $|A_{i_a, j} \cap A'| = 1$, where $i_a = \lfloor \log |A'| \rfloor$. Let $v_a \in A'$ be the vertex isolated by $A_{i_a, j}$, and let $B' \subseteq B(\gamma)$ be the neighbors of v_a on the other side of the cut. With constant probability $|B_{i_b, j} \cap B'| = 1$, where $i_b = \lfloor \log |B'| \rfloor$, isolating some vertex $v_b \in B'$. Thus, in this case $\widehat{\Upsilon}(i_a, i_b, i_c, j) = \langle v_a, v_b, \gamma \rangle$ is the name of an edge crossing the cut. \square

⁹Here “ γ ” refers to a $\log n$ -bit identifier for the component γ .

¹⁰Because each edge in $A(\gamma) \times B(\gamma)$ is tagged with its provenance γ , edges with the same endpoints but different provenances are distinguishable edges. Thus, we usually think of \widehat{E} as a set rather than a multiset.

The structures $\widehat{\mathcal{C}}$ and $\widehat{\mathcal{V}}$ are analogues of \mathcal{C} and \mathcal{V} but report sketches of edges in \widehat{E} . The structure $\widehat{\mathcal{B}}$ is new and is used to efficiently generate sketches of complete bipartite subgraphs of $A(\gamma) \times B(\gamma)$ on the fly.

$\widehat{\mathcal{C}}(\gamma, I)$: The input is a component $\gamma \in \mathcal{C}$ and interval I of some Euler(τ). Define $E_{u,\gamma}$ to be the \widehat{E} -edges joining u to the terminals of $V(\gamma)$. Report the sketch $\widehat{\Upsilon} = \bigoplus_{u \in I} \widehat{\Upsilon}^{E_{u,\gamma}}$.

$\widehat{\mathcal{V}}(v, I)$: The input is a vertex v and interval I of some Euler(τ). Let $E_{v,I}$ be the \widehat{E} edges joining v to terminals in I . Report the sketch $\widehat{\Upsilon}^{E_{v,I}}$.

$\widehat{\mathcal{B}}(\gamma, I, D)$: The input is a component γ , an interval $I \subseteq A(\gamma)$, and a set D of failed vertices such that $I \cap D = \emptyset$. Let $E_{I,D} = I \times (B(\gamma) - D) \oplus (A(\gamma) - D) \times (I \cap B(\gamma))$ be the subset of provenance- γ edges in $(A(\gamma) - D) \times (B(\gamma) - D)$ crossing the cut $(I, A(\gamma) - I)$; see Figure 5. Report the sketch matrix $\widehat{\Upsilon}^{E_{I,D}}$.

LEMMA 7.7. *The structures $\widehat{\mathcal{V}}$, $\widehat{\mathcal{C}}$, and $\widehat{\mathcal{B}}$ occupy $O(m \log^6 n)$ space. The query time for $\widehat{\mathcal{V}}$ and $\widehat{\mathcal{C}}$ is $O(\log^4 n)$, whereas the query time of $\widehat{\mathcal{B}}$ is $O(|D| \log^2 n + \log^4 n)$.*

Proof. The implementation of $\widehat{\mathcal{V}}(v, I)$ is exactly like $\mathcal{V}(v, I)$, except that $\widehat{\Upsilon}^r$ occupies $O(\log^4 n)$ space and is the sketch for *all* edges joining v and v_r (with different provenances). According to Lemma 7.1, the number of edges in \widehat{E} (ignoring multiplicity) is $O(m \log^2 n)$. Thus, the space for $\widehat{\mathcal{V}}$ is $O(m \log^6 n)$. The query time is still linear in the sketch size: $O(\log^4 n)$.

The implementation of $\widehat{\mathcal{C}}(\gamma, I)$ is also similar to $\mathcal{C}(\gamma, I)$, with a $O(\log^4 n)$ query time. We now analyze its space. Let γ be a component and v be a neighbor of γ that is a terminal in $V(\gamma')$. Each such pair (v, γ) contributes $O(\log^4 n)$ space to $\widehat{\mathcal{C}}$. We consider the pairs when $\gamma' \preceq \gamma$ and $\gamma' \succ \gamma$ separately. There are at most $O(pn)$ pairs (v, γ) when $\gamma' \preceq \gamma$ since v has at most p ancestral components, so the contribution of these is $O(pn \log^4 n) = O(n \log^5 n)$. Now suppose $\gamma' \succ \gamma$. Let u be some vertex in $V(\gamma)$ adjacent to v , and let γ'' be the provenance of the edge $\{u, v\}$. It must be that $\gamma'' \prec \gamma$ is a strict descendant of γ and that both $u, v \in A(\gamma'')$. This also implies that $v \in A(\gamma)$, hence the contribution of all pairs (v, γ) when $\gamma' \succ \gamma$ is $O(|A(\gamma)| \log^4 n)$, which is $O(m \log^5 n)$ over all γ .

We now turn to the new structure that answers the query $\widehat{\mathcal{B}}(\gamma, I, D)$. Let $\widehat{\Upsilon}_0$ be the sketch for $I \times (B(\gamma) - D)$ and $\widehat{\Upsilon}_1$ be the sketch for $(A(\gamma) - D) \times (I \cap B(\gamma))$. The output sketch is exactly $\widehat{\Upsilon}_0 \oplus \widehat{\Upsilon}_1$. We focus on the computation of $\widehat{\Upsilon}_0$; computing $\widehat{\Upsilon}_1$ is symmetric. Figures 5(a), (b) illustrate $\widehat{\Upsilon}_0$ and $\widehat{\Upsilon}_1$, respectively.

If $\gamma \notin \mathcal{C}_{i_c,j}$ then $\widehat{\Upsilon}_0(i_a, i_b, i_c, j) = \langle 0 \rangle$. Otherwise, define $A' = I \cap A_{i_a,j}$ and $B' = B(\gamma) \cap B_{i_b,j} - D$. Then

$$\widehat{\Upsilon}_0(i_a, i_b, i_c, j) = \left\langle \left(\bigoplus_{u \in A'} \langle u \rangle \right)^{|B'|}, \left(\bigoplus_{v \in B'} \langle v \rangle \right)^{|A'|}, \langle \gamma \rangle^{|A'| \cdot |B'|} \right\rangle,$$

where x^k is short for $\overbrace{x \oplus \cdots \oplus x}^k$. Note that to compute this entry of $\widehat{\Upsilon}_0$, we only need to be able to compute the parities of $|A'|$ and $|B'|$, and the sums $\bigoplus_{u \in A'} \langle u \rangle$ and $\bigoplus_{v \in B'} \langle v \rangle$.

Let $A(\gamma) = (u_1, \dots, u_{|A(\gamma)|})$ and $B(\gamma) = (v_1, \dots, v_{|B(\gamma)|})$. We store parity prefix sum matrices $(\sigma_k)_{k \in [1, |A(\gamma)|]}$ and $(\sigma'_k)_{k \in [1, |B(\gamma)|]}$ and name prefix sum matrices $(\rho_k)_{k \in [1, |A(\gamma)|]}$, and $(\rho'_k)_{k \in [1, |B(\gamma)|]}$, where

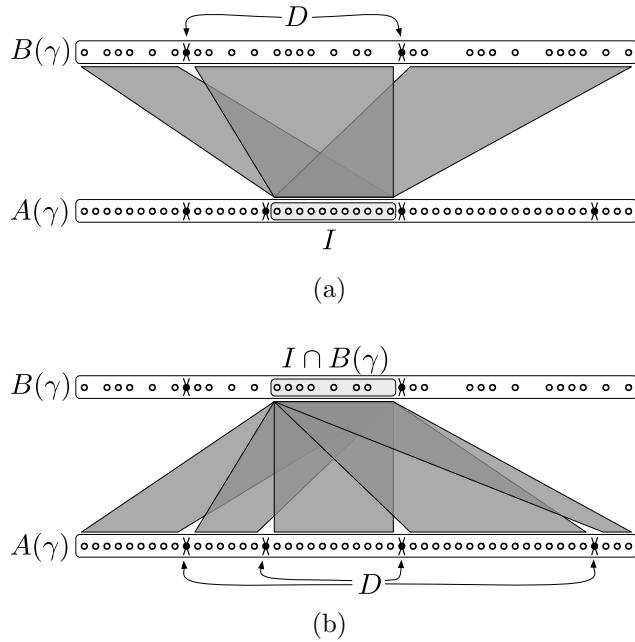


FIG. 5. (a) The complete bipartite graph $I \times (B(\gamma) - D)$ sketched by $\widehat{\Upsilon}_0$. (b) The complete bipartite graph $(I \cap B(\gamma)) \times (A(\gamma) - D)$ sketched by $\widehat{\Upsilon}_1$. Note that in $\widehat{\Upsilon} = \widehat{\Upsilon}_0 \oplus \widehat{\Upsilon}_1$, edges with both endpoints in I are included twice and cancel each other out.

$$\begin{aligned}\sigma_k(i_a, j) &= |\{u_1, \dots, u_k\} \cap A_{i_a, j}| \pmod 2, \\ \sigma'_k(i_b, j) &= |\{v_1, \dots, v_k\} \cap B_{i_b, j}| \pmod 2, \\ \rho_k(i_a, j) &= \bigoplus_{k' \in [1, k] : u_{k'} \in A_{i_a, j}} \langle u_{k'} \rangle, \\ \rho'_k(i_b, j) &= \bigoplus_{k' \in [1, k] : v_{k'} \in B_{i_b, j}} \langle v_{k'} \rangle.\end{aligned}$$

Suppose $I = (u_k, \dots, u_l)$ is the query interval. First compute the parity matrices $\sigma = \sigma_l \oplus \sigma_{k-1}$ and $\sigma' = \sigma'_{|B(\gamma)|}$ in $O(\log^2 n)$ time, and then compute σ'' in $O(|D| \log^2 n)$ time, where $\sigma''(i_b, j) = |D \cap B_{i_b, j}| \bmod 2$. Next compute name matrices $\rho = \rho_l \oplus \rho_{k-1}$ and $\rho' = \rho'_{|B(\gamma)|}$ in $O(\log^2 n)$ time and ρ'' in $O(|D| \log^2 n)$ time, where $\rho''(i_b, j) = \bigoplus_{v \in D \cap B_{i_b, j}} \langle v \rangle$. One may easily verify that σ, ρ are the correct parity and name matrices for I and that $\sigma' \oplus \sigma'', \rho' \oplus \rho''$ are the correct parity and name matrices for $B(\gamma) - D$. Each entry of the output matrix \widehat{Y}_0 is then computed in $O(1)$ time as follows:

$$\begin{aligned} & \widehat{\Upsilon}_0(i_a, i_b, i_c, j) \\ &= \begin{cases} \langle 0 \rangle & \text{if } \gamma \notin \mathcal{C}_{i_c, j}, \\ \left\langle (\rho(i_a, j))^{\sigma'(i_b, j) + \sigma''(i_b, j)}, \right. \\ \quad (\rho'(i_b, j) \oplus \rho''(i_b, j))^{\sigma(i_a, j)}, \\ \quad \left. \langle \gamma \rangle^{\sigma(i_a, j) \cdot (\sigma'(i_b, j) + \sigma''(i_b, j))} \right\rangle & \text{otherwise.} \end{cases} \end{aligned}$$

The overall time to compute $\widehat{\Upsilon}_0$ is therefore $O(|D| \log^2 n + \log^4 n)$.

1

Remark 7.8. Observe that in the proof of Lemma 7.7, the matrices σ'', ρ'' depended only on D , not I . Thus, once they are computed we can answer a query for a different triple (γ, I', D) in just $O(\log^4 n)$ time. This fact will be used in sections 7.3 and 7.4.

7.3. Update and query algorithms. At a high level, the deletion algorithm has four major steps.

1. The first task is to mark up to $(p+1)d \leq d \log n$ components $\gamma_1, \dots, \gamma_{(p+1)d}$ as *affected*, as well as the corresponding trees $\tau_1, \dots, \tau_{(p+1)d}$. (Because the component-to-tree mapping is not injective, the number of distinct trees may be smaller.) We mark all tree edges incident to D as deleted, which breaks up $\text{Euler}(\tau_1), \dots, \text{Euler}(\tau_{pd})$ into $O(pd)$ intervals, call them $I_1, \dots, I_{O(pd)}$, with the property that each affected subtree (i.e., those in $\tau_1 - D, \dots, \tau_{(p+1)d} - D$) is the union of some subset of the intervals.
2. The next task is to generate two sketches $\Upsilon[I_q], \widehat{\Upsilon}[I_q]$ for each interval representing valid edges in E and \widehat{E} , respectively, joining I_q to another interval. In other words, we do not want to consider original or artificial edges adjacent to D , nor invalid artificial edges with provenance γ for some affected γ , nor valid artificial edges joining I_q to an unaffected tree in \mathcal{T} . The structures $\mathcal{V}, \mathcal{C}, \widehat{\mathcal{V}}, \widehat{\mathcal{C}}, \widehat{\mathcal{B}}$ are used to build these sketches.
3. Let $t_1, \dots, t_{O(pd)}$ be the affected subtrees. We form the sketches $\Upsilon[t_q]$ and $\widehat{\Upsilon}[t_q]$ for each tree, by XORing the sketches of the constituent intervals of t_q . According to Lemmas 7.4 and 7.6, these sketches reveal one edge crossing the cut defined by $V(t_q)$, with constant probability. We can implement a probabilistic version of Borůvka's algorithm in order to compute the connected components among the affected subtrees. The j th Borůvka step only examines parts of the sketch with matching j -coordinate. Using “fresh” randomness for each Borůvka step is essential for showing the procedure succeeds w.h.p.
4. Lastly, we must account for any unaffected components $\gamma \in \mathcal{C}$ that were unlucky enough to see all vertices in $B(\gamma)$ fail. According to Lemma 7.2, at least $|A(\gamma)|$ of the pairs in $\binom{D}{2}$ are owned by γ , w.h.p. We scan all $\binom{|D|}{2}$ labels, tallying up how many times each owner label occurs. Any owner label γ that appears $|A(\gamma)|$ times might provide additional connectivity not captured by the components discovered at the end of step 3. We merge any connected components from step 3 that contain at least one $A(\gamma)$ vertex. This takes $O(|A(\gamma)|)$ time to process γ and hence $O(d^2)$ time overall.

7.3.1. Generating sketches. We show how to generate $\widehat{\Upsilon}[I_q]$. The process for $\Upsilon[I_q]$ is analogous but simpler and faster. For each $z \leq (p+1)d$, consult with $\widehat{\mathcal{C}}$ to get a sketch $\widehat{\Upsilon}(\gamma_z, I_q)$ covering edges in \widehat{E} joining I_q to terminals in $V(\gamma_z)$. These sketches include two types of edges we must subtract off: (i) those incident to D and (ii) those with provenance γ for some affected γ .¹¹ For each $v \in D$, consult with $\widehat{\mathcal{V}}$ to get a sketch $\widehat{\Upsilon}(v, I_q)$ covering edges in \widehat{E} joining I_q to v . These sketches cover type (i) bad edges. Suppose I is an interval containing terminals of $V(\gamma_y)$. For each $z \leq (p+1)d$, if $\gamma_z \prec \gamma_y$ is a strict descendant of γ_y , consult $\widehat{\mathcal{B}}$ to get a sketch

¹¹Note that the intersection of (i) and (ii) is generally nonempty, so it is not sufficient to subtract off (i) and (ii) separately as this will inadvertently add back edges in (i) \cap (ii).

$\widehat{\Upsilon}(\gamma_z, I_q, D)$. This covers all remaining edges with provenance γ_z not already covered by $\{\widehat{\Upsilon}(v, I_q)\}_{v \in D}$. Finally, we compute $\widehat{\Upsilon}[I_q]$ by combining these sketches.

$$\widehat{\Upsilon}[I_q] = \left(\bigoplus_{z \leq (p+1)d} \widehat{\Upsilon}(\gamma_z, I_q) \right) \oplus \left(\bigoplus_{v \in D} \widehat{\Upsilon}(v, I_q) \right) \oplus \left(\bigoplus_{z \leq (p+1)d} \widehat{\Upsilon}(\gamma_z, I_q, D) \right).$$

For the time analysis, recall that there are $O(pd)$ affected components, $O(pd)$ affected subtrees, and $O(pd)$ relevant Euler tour intervals. By Lemma 7.7, the time to compute all $\widehat{\Upsilon}(\gamma_z, I_q)$ sketches is $O((pd)^2 \log^4 n) = O(d^2 \log^6 n)$, and the time to compute $\widehat{\Upsilon}(v, I_q)$ sketches $O(pd^2 \log^4 n) = O(d^2 \log^5 n)$. By Lemma 7.7 and Remark 7.8, the time to compute all $\widehat{\Upsilon}(\gamma_z, I_q, D)$ sketches is $O((pd)d \log^2 n + (pd)^2 \log^4 n) = O(d^2 \log^6 n)$.

7.3.2. Executing Boruvka's algorithm. Once the sketches for each interval are generated we can combine them to form sketches $\Upsilon[t_l], \widehat{\Upsilon}[t_l]$ for each affected subtree t_l .

We proceed as in Boruvka's MST algorithm [17] and many parallel connectivity algorithms that use the “hook and contract” technique [27, 28, 59, 74]. In each round, each affected subtree will pick an arbitrary edge joining it to a different affected subtree. The affected subtrees will be merged into larger affected subtrees, which participate in the next round. Under error-free conditions—which we do not have—this process will halt after $\log_2(O(dp))$ rounds since each round reduces the number of nonisolated affected subtrees by at least half.

The formal procedure is as follows. Let $C_{j-1} = \{t_{j-1,1}, t_{j-1,2}, \dots, t_{j-1,|C_{j-1}|}\}$ be the affected trees after $j-1$ rounds, where $C_0 = \{t_{0,1}, \dots, t_{0,O(dp)}\}$. We maintain the invariant that we have, for each $t_{j-1,l}$, sketches $\Upsilon[t_{j-1,l}], \widehat{\Upsilon}[t_{j-1,l}]$ covering original and artificial edges joining $t_{j-1,l}$ to a different tree. In the j th round, loop over each $t_{j-1,l} \in C_{j-1}$ and look for the name of any valid original/artificial edge in the $\log n$ entries of $\Upsilon[t_{j-1,l}](\star, j)$ and the $\log^3 n$ entries of $\widehat{\Upsilon}[t_{j-1,l}](\star, \star, \star, j)$. Such an edge $e_{j-1,l}$, if it exists, has one endpoint in $V(t_{j-1,l})$. Let C_j be the components induced by the C_{j-1} trees and the intertree edges $\{e_{j-1,l}\}$ just selected. Suppose the constituent trees of some $t_{j,r} \in C_j$ are $S \subseteq C_{j-1}$. The sketches for $t_{j,r}$ are computed as $\Upsilon[t_{j,r}] = \bigoplus_{t \in S} \Upsilon[t]$ and $\widehat{\Upsilon}[t_{j,r}] = \bigoplus_{t \in S} \widehat{\Upsilon}[t]$. The total time to compute sketches for C_j is just $O((|C_{j-1}| - |C_j|) \log^4 n)$.

Observe that just before executing the j th round we have only examined sketch entries whose final coordinate is in $\{1, \dots, j-1\}$. Hence, the contents of the sketches with final coordinate j reflect “fresh” randomness, and we can apply Lemmas 7.4 and 7.6. If there exists at least one edge crossing the cut defined by $t_{j-1,l}$, then with constant probability, either $\Upsilon[t_{j-1,l}]$ or $\widehat{\Upsilon}[t_{j-1,l}]$ will reveal the name of one such edge. Letting $\|C_k\|$ denote the number of nonisolated components in C_k , we have $\mathbb{E}[\|C_j\|] \leq (1 - \epsilon)\|C_{j-1}\|$ for some absolute constant $\epsilon > 0$. Thus, after $c \log n$ rounds $\mathbb{E}[\|C_{c \log n}\|] \leq (1 - \epsilon)^c \log n \|C_0\| < n^{-\Omega(c)}$ and by Markov's inequality, the probability that $C_{c \log n}$ has nonisolated components (an error) is $n^{-\Omega(c)}$.

7.3.3. Recapitulation. The high level update algorithm in section 7.3 was divided into four major steps. Step 1 (marking affected components and subtrees, enumerating relevant intervals) takes $O(d \log n)$ time. Step 2 (generating sketches) takes $O(d^2 \log^6 n)$ time. Step 3 (Boruvka's algorithm) takes time linear in the sum of

the sketches: $O(d \log^4 n)$. Finally, Step 4 (processing γ with $B(\gamma) \subseteq D$) takes $O(d^2)$ time. Observe that due to the probabilistic nature of Lemmas 7.2, 7.4, and 7.6, Steps 3 and 4 can have both detected and undetected errors, with probability $n^{-\Omega(c)}$.¹²

The final output of this algorithm (a partition of the affected subtrees into connected components) is exactly the same as in the deterministic algorithms of sections 4–6. Thus, the same deterministic query algorithm works in $O(d)$ time. In the next section we shall see some general methods to shave $\text{poly}(\log n)$ -factors off some algorithms that use graph sketches.

7.4. Improving update times with on-demand sketching. Recall that existing d -edge failure connectivity oracles have update times that are linear in d but have $\text{poly}(\log n)$ factors ($O(d \log^2 n \log \log n)$ [73] or $O(d \log d \log^3 n)$ [61]) or have a quadratic dependence on d , but better dependence on n , namely $O(d^2 \log \log n)$ [40]. In this section we show how to improve all of these bounds and use sublinear space, as in [51].

THEOREM 7.9. *A connectivity oracle for $G = (V, E)$ with size $O(n \log^2 n)$ can be constructed in $O(m \log n + n \log^2 n)$ time. Any set $D \subseteq E(G)$ of d edges can be processed in $O(d \log d \log \log n)$ time in expectation (and $O(d \log n \log \log n)$ time w.h.p.) such that connectivity queries in the graph $(V, E - D)$ can be answered in time $O(\min\{\log \log n, \frac{\log d}{\log \log n}\})$. With high probability, the query is answered correctly.*

Proof. Because the space is sublinear in m we cannot afford to store the graph, nor can we explicitly record for each edge which samples it appears in. Assume the initial vertex ids are $\{1, \dots, n\}$. We assign u the bit string $\phi(u)$, where $\phi : \{1, \dots, n\} \rightarrow \{0, 1\}^{c \log n}$ is a uniformly random injective function. The encoding of an edge $e = \{u, v\}$ is $\langle e \rangle = \langle \min\{\phi(u), \phi(v)\}, \max\{\phi(u), \phi(v)\} \rangle$.

Sketching. We use hash functions to decide whether to include edges in sampled sets. Choose pairwise independent hash functions $h_1, \dots, h_{c \log n} : \{0, 1\}^{2 \log n} \rightarrow \{0, \dots, 2^w - 1\}$, and for each $i \in [0, \log m]$ and $j \in [1, c \log n]$, let $E_{i,j}$ be the edge set

$$E_{i,j} = \{e \in E \mid h_j(e) \in [0, 2^{w-i})\}.$$

The sketch $\Upsilon^{E'}$ is a $\log m \times c \log n$ matrix defined exactly as before. Pairwise independence suffices to guarantee the claim of Lemma 7.4, that for any set $E' \subset E$ and any j , there exists an i such that with constant probability, $\Upsilon^{E'}(i, j)$ is the name of one edge in E' . (See [51, Appendix A] for a short proof.) Moreover, since $E = E_{0,j} \supseteq \dots \supseteq E_{\log m-1,j}$, the right value of i is, w.h.p., the unique value for which $\Upsilon^{E'}(i, j) \neq \langle 0 \rangle$ and $\Upsilon^{E'}(i+1, j) = \langle 0 \rangle$. We also need to be able to tell that a bit string $\Upsilon^{E'}(i, j)$ encodes an edge rather than garbage. Since ϕ assigns random $c \log n$ -bit strings, the XOR of multiple edge names is a random $2c \log n$ -bit string. Thus, the probability that a garbage string looks like a legitimate edge name is $n^{-2(c-1)}$.

The construction. At preprocessing time, choose an arbitrary spanning tree $T \subseteq E(G)$ and an arbitrary tour $\text{Euler}(T) = (v_1, \dots, v_n)$. Initialize sketch matrices $\Upsilon^{v_1}, \dots, \Upsilon^{v_n}$ to be all zero. For each $e = (v_k, v_l) \in E(G)$, evaluate $h_1(e), \dots, h_{c \log n}(e)$ to determine which sets $E_{i,j}$ contain e . If $e \in E_{i,j}$, update $\Upsilon^{v_k}(i, j) \leftarrow \Upsilon^{v_k}(i, j) \oplus \langle e \rangle$ and likewise with $\Upsilon^{v_l}(i, j)$. Finally, compute all prefix sum sketches (μ_1, \dots, μ_n) ,

¹²One undetected error that has nothing to do with sketching is if $B(\gamma) \subseteq D$, but γ is not processed in Step 4. An undetected sketch failure occurs if $t_{j,l}$ is not an isolated tree, but nonetheless $\Upsilon[t_{j,l}]$ and $\widehat{\Upsilon}[t_{j,l}]$ are the all-zero matrices. A detected error would be if $\Upsilon[t_{c \log n+1,l}]$ or $\Upsilon[t_{c \log n+1,l}]$ were not the all-zero matrices, indicating that $c \log n$ Borůvka steps failed to detect all connected components.

where $\mu_k = \bigoplus_{k' \leq k} \Upsilon^{v_{k'}}$. The data structure stores T , $\text{Euler}(T)$, and (μ_k) . The space is dominated by (μ_k) , which takes $O(n \log^2 n)$ words. The construction time is $O(m \log n + n \log^2 n)$ in expectation. Observe that each edge causes just $O(c \log n)$ entries of the sketches to be updated, in expectation, and that computing (μ_k) takes $O(n \log^2 n)$ time once the (Υ^{v_k}) are computed.

Handling edge failures. Suppose a subset $D \subseteq E(G)$ of edges are deleted.¹³ Removing D partitions $\text{Euler}(T)$ into a set of $2|D \cap T| + 1$ intervals, call them \mathcal{I} . For each interval $I \in \mathcal{I}$, suppose it is $\{v_p, \dots, v_q\}$, we compute its initial sketch $\Upsilon[I] \leftarrow \mu_q \oplus \mu_{p-1}$, then proceed to delete D from the sketches. For each $e = (v_k, v_l) \in D$, find the intervals $I, I' \in \mathcal{I}$ containing v_k, v_l , respectively, and for each $E_{i,j} \ni e$, update $\Upsilon[I](i, j) \leftarrow \Upsilon[I](i, j) \oplus \langle e \rangle$ and update $\Upsilon[I']$ likewise. Once we have sketches for all intervals, we execute Borůvka's algorithm as in section 7.3.2. The time to generate the sketches and execute Borůvka's algorithm takes time linear in the size of all sketches, namely $O(d \log^2 n)$.

To improve the update time we calculate entries in sketch matrices in an on-demand fashion. Suppose t is a tree encountered during Borůvka's algorithm. We maintain a linked list $\mathcal{L}[t]$ of sketches satisfying the invariant $\Upsilon[t] = \bigoplus_{\sigma \in \mathcal{L}[t]} \sigma$. (For example, before the first Borůvka step, t is an interval in \mathcal{I} and $\mathcal{L}[t]$ consists of two μ sketches and possibly several single-edge sketches, one for each edge in D with an endpoint in t .) Thus, any entry $\Upsilon[t](i, j)$ can be looked up in $|\mathcal{L}[t]|$ time. In the j th Borůvka step, for each current tree t we do a binary search for the maximum i such that $\Upsilon[t](i, j) \neq \langle 0 \rangle$ and check whether it is a legitimate encoding of an edge. In this Borůvka step, if trees t_1, \dots, t_r are merged into one tree t' , we simply set $\mathcal{L}[t']$ to be the concatenation of $\mathcal{L}[t_1], \dots, \mathcal{L}[t_r]$.

The number of basic sketches appearing in any list $\mathcal{L}[\cdot]$ is $O(d)$: there are at most $2|\mathcal{I}| = O(d)$ μ -sketches of interest and at most d single-edge sketches for edges in D . If Borůvka's algorithm terminates after b steps, then we have probed $O(b \log \log n)$ locations in each of the basic sketches, for a total time of $O(d b \log \log n)$. The claimed update time follows from the fact that b is $O(\log d)$ in expectation and $O(\log n)$ w.h.p.

Queries. A query (u_k, u_l) simply needs to find the intervals I, I' containing u_k, u_l , respectively, and check whether I, I' are in the same connected component discovered by Borůvka's algorithm. Finding I, I' can be done with predecessor search, in $O(\log \log n)$ time [80] or $O(\log d / \log \log n)$ time [77]. \square

The same technique allows us to shave four log factors off the update time from section 7.3.

THEOREM 7.10. *A connectivity oracle for $G = (V, E)$ with size $O(m \log^6 n)$ can be constructed in $O(mn \log n)$ time. Any set $D \subseteq V(G)$ of d vertices can be processed in $O(d^2 \log d \log^2 n \log \log n)$ time in expectation (and $O(d^2 \log^3 n \log \log n)$ time w.h.p.) such that connectivity queries in $G - D$ can be answered in $O(d)$ time. With high probability, the query is answered correctly.*

Proof. Consider how we construct the sketch matrix $\hat{\Upsilon}[I]$ for an interval I . For each affected component γ_z , $\hat{\Upsilon}(\gamma_z, I)$ is the sum of two β sketches, and for each $v \in D$, $\hat{\Upsilon}(v, I)$ is the sum of two α sketches. Recall that entries of $\hat{\Upsilon}(\gamma_z, I, D)$ are computable in $O(1)$ time, given matrices $\sigma, \sigma', \rho, \rho'$, and σ'', ρ'' . The first four matrices depend only on γ_z and entries in them can be computed in $O(1)$ time. The last two matrices

¹³We are promised that $D \subseteq E(G)$, which cannot be verified with only $\tilde{O}(n)$ space. Strictly speaking, we will be preparing a data structure that answers connectivity queries in $G' = (V, E \oplus D)$, i.e., any edge $e \in E(G) - D$ is treated as an insertion, not a deletion.

depend on both γ_z and D , and each of their entries takes $O(d)$ time to compute. Thus, if b Borůvka steps suffice, it takes $O(bd^2 \log^2 n)$ time to compute the relevant entries of the σ'', ρ'' matrices, over all $O(d \log n)$ affected γ_z .

Now consider a tree t in the j th Borůvka step. We look for an edge with one endpoint in t via three binary searches over $\widehat{\Upsilon}[t]$. We find the maximum i_c for which $\widehat{\Upsilon}[t](0, 0, i_c, j) \neq \langle 0 \rangle$, then find the maximum i_a for which $\widehat{\Upsilon}[t](i_a, 0, i_c, j) \neq \langle 0 \rangle$, then find the maximum i_b for which $\widehat{\Upsilon}[t](i_a, i_b, i_c, j) \neq \langle 0 \rangle$. With constant probability, this entry contains the name of an edge with one endpoint in t . Thus, each of the $O((d \log n)^2)$ basic sketches is probed in $O(b \log \log n)$ locations, for a total time of $O(d^2 b \log^2 n \log \log n)$. Once again, b is $O(\log d)$ in expectation and $O(\log n)$ w.h.p. \square

8. Conclusions. In this paper we illustrated the power of a new graph decomposition theorem by giving time- and space-efficient connectivity oracles for graphs subject to vertex failures. Our data structures perform well in all the major measures of efficiency (space, update time, query time, and preprocessing time) but leave many opportunities for improvement. The following open problems are quite challenging.

- The Fürer–Raghavachari [48] algorithm FR-Tree for computing the near-minimum degree spanning tree takes $O(mn \log n)$ time, which is the main bottleneck in our construction. Is it possible to reduce the running time of FR-Tree to $\tilde{O}(m)$, or compute spanning trees with similar decomposition properties in $\tilde{O}(m)$ time? Would such a result contradict a popular hardness conjecture?¹⁴
- The conditional lower bounds of [63, 54] show that any connectivity oracle with reasonable update time cannot have $\tilde{O}(1)$ query time, independent of d , but they do not preclude a data structure having both query and update time $\tilde{O}(d)$. Is it possible to reduce the update time below $O(d^2)$ without disturbing the space or query time?
- Is it possible to reduce the space of our deterministic d_\star -failure connectivity oracle to $\tilde{O}(m)$ (independent of d_\star) or perhaps $\tilde{O}(d_\star n)$?

A more accessible problem is to eliminate log factors, especially in our Monte Carlo structure, which still has an extra $\log^6 n$ factor in space and $\log^2 n$ factor in update time.

Acknowledgment. We would like to thank Kasper Green Larsen and Peyman Afshani for help with the navigating the range searching literature, Shiri Chechik for suggesting the reduction to 3D range searching in section 6.3, and Veronika Loitzenbauer for bringing [55] to our attention and pointing out the $\Omega(\min\{m, d_\star n\})$ lower bound on d_\star -failure connectivity oracles.

REFERENCES

- [1] I. ABRAHAM, S. CHECHIK, AND C. GAVOILLE, *Fully dynamic approximate distance oracles for planar graphs via forbidden-set distance labels*, in Proceedings of the 44th ACM Symposium on Theory of Computing, 2012, pp. 1199–1218, <https://doi.org/10.1145/2213977.2214084>.
- [2] I. ABRAHAM, S. CHECHIK, C. GAVOILLE, AND D. PELEG, *Forbidden-set distance labels for graphs of bounded doubling dimension*, ACM Trans. Algorithms, 12 (2016), 22, <https://doi.org/10.1145/2818694>.
- [3] K. J. AHN, S. GUHA, AND A. MCGREGOR, *Analyzing graph structure via linear measurements*, in Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms, 2012, pp. 459–467.

¹⁴See Open Problem 24 from the *Structure and Hardness in P* open problems list [66].

- [4] S. ARORA, E. HAZAN, AND S. KALE, *$O(\sqrt{\log n})$ approximation to SPARSEST CUT in $\tilde{O}(n^2)$ time*, SIAM J. Comput., 39 (2010), pp. 1748–1771, <https://doi.org/10.1137/080731049>.
- [5] S. ARORA, S. RAO, AND U. V. VAZIRANI, *Expander flows, geometric embeddings and graph partitioning*, J. ACM, 56 (2009).
- [6] I. BARAN, E. D. DEMAINE, AND M. PĂTRAȘCU, *Subquadratic algorithms for 3SUM*, Algorithmica, 50 (2008), pp. 584–596.
- [7] D. BARNETTE, *Trees in polyhedral graphs*, Canad. J. Math., 18 (1966), pp. 731–736.
- [8] S. BASWANA, U. LATH, AND A. S. MEHTA, *Single source distance oracle for planar digraphs avoiding any failed node or link*, in Proceedings of the 23rd ACM-SIAM Symposium on Discrete Algorithms, 2012, pp. 223–232.
- [9] S. BASWANA, S. R. CHAUDHURY, K. CHAUDHARY, AND S. KHAN, *Dynamic DFS in undirected graphs: Breaking the $o(m)$ barrier*, in Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms, 2016, pp. 730–739, <https://doi.org/10.1137/1.9781611974331.ch52>.
- [10] S. BASWANA, K. CHAUDHARY, AND L. RODITTY, *Fault tolerant subgraph for single source reachability: generic and optimal*, in Proceedings of the 48th Annual ACM Symposium on Theory of Computing, 2016, pp. 509–518, <https://doi.org/10.1145/2897518.2897648>.
- [11] G. D. BATTISTA AND R. TAMASSIA, *On-line maintenance of triconnected components with SPQR-trees*, Algorithmica, 15 (1996), pp. 302–318.
- [12] M. A. BENDER AND M. FARACH-COLTON, *The LCA problem revisited*, in Proceedings of the 4th Latin American Symposium on Theoretical Informatics, Lecture Notes in Comput. Sci. 1776, Springer, New York, 2000, pp. 88–94.
- [13] M. A. BENDER AND M. FARACH-COLTON, *The level ancestor problem simplified*, Theoret. Comput. Sci., 321 (2004), pp. 5–12.
- [14] A. BERNSTEIN AND D. KARGER, *A nearly optimal oracle for avoiding failed vertices and edges*, in Proceedings of the 41st Annual ACM Symposium on Theory of Computing, 2009, pp. 101–110.
- [15] D. BILÒ, F. GRANDONI, L. GUALÀ, S. LEUCCI, AND G. PROIETTI, *Improved purely additive fault-tolerant spanners*, in Proceedings of the 23rd Annual European Symposium on Algorithms, 2015, pp. 167–178, https://doi.org/10.1007/978-3-662-48350-3_15.
- [16] G. BORRADAILE, S. PETTIE, AND C. WULFF-NILSEN, *Connectivity oracles for planar graphs*, in Proceedings of the 13th Scandinavian Symposium and Workshops on Algorithm Theory, 2012, pp. 316–327, https://doi.org/10.1007/978-3-642-31155-0_28.
- [17] O. BORŮVKA, *O jistém problémě minimálním*, Práce Moravské Přírodovědecké Společnosti, 3 (1926), pp. 37–58.
- [18] G. BRAUNSCHVIG, S. CHECHIK, AND D. PELEG, *Fault tolerant additive spanners*, in Proceedings of the 38th International Workshop on Graph-Theoretic Concepts in Computer Science, 2012, pp. 206–214, https://doi.org/10.1007/978-3-642-34611-8_22.
- [19] G. BRAUNSCHVIG, S. CHECHIK, D. PELEG, AND A. SEALFON, *Fault tolerant additive and (μ, α) -spanners*, Theoret. Comput. Sci., 580 (2015), pp. 94–100, <https://doi.org/10.1016/j.tcs.2015.02.036>.
- [20] T. CHAN, *Dynamic subgraph connectivity with geometric applications*, SIAM J. Comput., 36 (2006), pp. 681–694.
- [21] T. M. CHAN, K. G. LARSEN, AND M. PĂTRAȘCU, *Orthogonal range searching on the RAM, revisited*, in Proceedings of the 27th ACM Symposium on Computational Geometry, 2011, pp. 1–10, <https://doi.org/10.1145/1998196.1998198>.
- [22] T. M. CHAN, M. PĂTRAȘCU, AND L. RODITTY, *Dynamic connectivity: Connecting to networks and geometry*, SIAM J. Comput., 40 (2011), pp. 333–349.
- [23] T. M. CHAN AND B. T. WILKINSON, *Adaptive and approximate orthogonal range counting*, in Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms, 2013, pp. 241–251, <https://doi.org/10.1137/1.9781611973105.18>.
- [24] S. CHECHIK, M. LANGBERG, D. PELEG, AND L. RODITTY, *Fault tolerant spanners for general graphs*, SIAM J. Comput., 39 (2010), pp. 3403–3423.
- [25] S. CHECHIK, M. LANGBERG, D. PELEG, AND L. RODITTY, *f -sensitivity distance oracles and routing schemes*, Algorithmica, 63 (2012), pp. 861–882, <https://doi.org/10.1007/s00453-011-9543-0>.
- [26] L. CHEN, R. DUAN, R. WANG, H. ZHANG, AND T. ZHANG, *An improved algorithm for incremental DFS tree in undirected graphs*, in Proceedings of the 16th Scandinavian Symposium and Workshops on Algorithm Theory, LIPIcs. Leibniz Int. Proc. Inform. 101, Schloss Dagstuhl, 2018, pp. 16:1–16:12, <https://doi.org/10.4230/LIPIcs.SWAT.2018.16>.
- [27] K. W. CHONG, Y. HAN, AND T. W. LAM, *Concurrent threads and optimal parallel minimum spanning trees algorithm*, J. ACM, 48 (2001), pp. 297–323.

- [28] K. W. CHONG AND T. W. LAM, *Finding connected components in $O(\log n \log \log n)$ time on the EREW PRAM*, J. Algorithms, 18 (1995), pp. 378–402.
- [29] K. CHOUDHARY, *An optimal dual fault tolerant reachability oracle*, in Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming, 2016.
- [30] A. CZUMAJ AND W.-B. STROTHMANN, *Bounded degree spanning trees*, in Proceedings of the 5th Annual European Symposium on Algorithms, 1997, pp. 104–117, https://doi.org/10.1007/3-540-63397-9_9.
- [31] C. DEMETRESCU, M. THORUP, R. A. CHOWDHURY, AND V. RAMACHANDRAN, *Oracles for distances avoiding a failed node or link*, SIAM J. Comput., 37 (2008), pp. 1299–1318.
- [32] E. A. DINIC, A. V. KARZANOV, AND M. V. LOMONOSOV, *On the structure of the system of minimum edge cuts in a graph*, Stud. Discrete Optim., (1976), pp. 290–306 (in Russian).
- [33] M. DINITZ AND R. KRAUTHGAMER, *Fault-tolerant spanners: Better and simpler*, in Proceedings of the 30th ACM Symposium on Principles of Distributed Computing, 2011, pp. 169–178.
- [34] Y. DINITZ AND Z. NUTOV, *A 2-level cactus model for the system of minimum and minimum+1 edge-cuts in a graph and its incremental maintenance*, in Proceedings of the 27th ACM Symposium on Theory of Computing, 1995, pp. 509–518.
- [35] Y. DINITZ AND Z. NUTOV, *A 2-Level Cactus Tree Model for the System of Minimum and Minimum+1 Edge Cuts of a Graph and Its Incremental Maintenance. Part I: The Odd Case*, unpublished manuscript, 1999.
- [36] Y. DINITZ AND Z. NUTOV, *A 2-Level Cactus Tree Model for the System of Minimum and Minimum+1 Edge Cuts of a Graph and Its Incremental Maintenance. Part II: The Even Case*, unpublished manuscript, 1999.
- [37] Y. DINITZ AND J. WESTBROOK, *Maintaining the classes of 4-edge-connectivity in a graph online*, Algorithmica, 20 (1998), pp. 242–276.
- [38] R. DUAN, *New data structures for subgraph connectivity*, in Proceedings of the 37th International Colloquium on Automata, Languages and Programming, 2010, pp. 201–212.
- [39] R. DUAN AND S. PETTIE, *Dual-failure distance and connectivity oracles*, in Proceedings of the 20th ACM-SIAM Symposium on Discrete Algorithms, 2009, pp. 506–515.
- [40] R. DUAN AND S. PETTIE, *Connectivity oracles for failure prone graphs*, in Proceedings of the 42nd ACM Symposium on Theory of Computing, 2010, pp. 465–474.
- [41] R. DUAN AND S. PETTIE, *Connectivity oracles for graphs subject to vertex failures*, in Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms, 2017, pp. 490–509.
- [42] R. DUAN AND L. ZHANG, *Faster randomized worst-case update time for dynamic subgraph connectivity*, in Proceedings of the 15th International Symposium on Algorithms and Data Structures, 2017, pp. 337–348, https://doi.org/10.1007/978-3-319-62127-2_29.
- [43] D. P. DUBHASHI AND A. PANCONESI, *Concentration of Measure for the Analysis of Randomized Algorithms*, Cambridge University Press, Cambridge, MA, 2009.
- [44] D. EPPSTEIN, Z. GALIL, G. ITALIANO, AND A. NISSENZWEIG, *Sparsification—a technique for speeding up dynamic graph algorithms*, J. ACM, 44 (1997), pp. 669–696.
- [45] G. FREDERICKSON, *Data structures for on-line updating of minimum spanning trees, with applications*, SIAM J. Comput., 14 (1985), pp. 781–798.
- [46] A. FREUND, *Improved subquadratic 3SUM*, Algorithmica, 77 (2017), pp. 440–458, <https://doi.org/10.1007/s00453-015-0079-6>.
- [47] D. FRIGIONI AND G. F. ITALIANO, *Dynamically switching vertices in planar graphs*, Algorithmica, 28 (2000), pp. 76–103.
- [48] M. FÜRER AND B. RAGHAVACHARI, *Approximating the minimum-degree Steiner tree to within one of optimal*, J. Algorithms, 17 (1994), pp. 409–423, <https://doi.org/10.1006/jagm.1994.1042>.
- [49] H. N. GABOW AND R. E. TARJAN, *A linear-time algorithm for a special case of disjoint set union*, J. Comput. System Sci., 30 (1985), pp. 209–221.
- [50] Z. GALIL AND G. ITALIANO, *Maintaining the 3-edge-connected components of a graph on-line*, SIAM J. Comput., 22 (1993), pp. 11–28.
- [51] D. GIBB, B. M. KAPRON, V. KING, AND N. THORN, *Dynamic Graph Connectivity with Improved Worst Case Update Time and Sublinear Space*, CoRR, abs/1509.06464, 2015.
- [52] O. GOLD AND M. SHARIR, *Improved Bounds for 3SUM, k-SUM, and Linear Degeneracy*, in Proceedings of the 25th Annual European Symposium on Algorithms, K. Pruhs and C. Sohler, eds., LIPIcs. Leibniz Int. Proc. Inform. 87, Schloss Dagstuhl, 2017, pp. 42:1–42:13, <https://doi.org/10.4230/LIPIcs.ESA.2017.42>.
- [53] A. GRØNLUND AND S. PETTIE, *Threesomes, degenerates, and love triangles*, J. ACM, 65 (2018), pp. 22:1–22:25, <https://doi.org/10.1145/3185378>.
- [54] M. HENZINGER, S. KRINNINGER, D. NANONGKAI, AND T. SARANURAK, *Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture*,

in Proceedings of the 47th Annual ACM Symposium on Theory of Computing, 2015, pp. 21–30.

[55] M. HENZINGER AND S. NEUMANN, *Incremental and fully dynamic subgraph connectivity for emergency planning*, in Proceedings of the 24th Annual European Symposium on Algorithms, 2016, pp. 48:1–48:11.

[56] J. HOLM, K. DE LICHTENBERG, AND M. THORUP, *Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity*, J. ACM, 48 (2001), pp. 723–760.

[57] S.-E. HUANG, D. HUANG, T. KOPELOWITZ, AND S. PETTIE, *Fully dynamic connectivity in $O(\log n(\log \log n)^2)$ amortized expected time*, in Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms, 2017, pp. 510–520, <https://doi.org/10.1137/1.9781611974782.32>.

[58] J. JAJA, C. W. MORTENSEN, AND Q. SHI, *Space-efficient and fast algorithms for multidimensional dominance reporting and counting*, in Proceedings of the 15th International Symposium on Algorithms and Computation, 2004, pp. 558–568, <https://doi.org/10.1007/978-3-540-30551-4-49>.

[59] D. B. JOHNSON AND P. METAXAS, *Connected components in $O(\log^{3/2} n)$ parallel time for the CREW PRAM*, J. Comput. System Sci., 54 (1997), pp. 227–242.

[60] A. KANEVSKY, R. TAMASSIA, G. D. BATTISTA, AND J. CHEN, *On-line maintenance of the four-connected components of a graph*, in Proceedings of the 32nd IEEE Symposium on Foundations of Computer Science, 1991, pp. 793–801.

[61] B. M. KAPRON, V. KING, AND B. MOUNTJOY, *Dynamic graph connectivity in polylogarithmic worst case time*, in Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms, 2013, pp. 1131–1142.

[62] C. KEJLBERG-RASMUSSEN, T. KOPELOWITZ, S. PETTIE, AND M. THORUP, *Faster worst case deterministic dynamic connectivity*, in Proceedings of the 24th European Symposium on Algorithms, 2016, pp. 53:1–53:15.

[63] T. KOPELOWITZ, S. PETTIE, AND E. PORAT, *Higher lower bounds from the 3SUM conjecture*, in Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms, 2016, pp. 1272–1287, <https://doi.org/10.1137/1.9781611974331.ch89>.

[64] K. G. LARSEN AND F. VAN WALDERVEEN, *Near-optimal range reporting structures for categorical data*, in Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms, 2013, pp. 265–276, <https://doi.org/10.1137/1.9781611973105.20>.

[65] K. G. LARSEN AND R. R. WILLIAMS, *Faster online matrix-vector multiplication*, in Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms, 2017, pp. 2182–2189, <https://doi.org/10.1137/1.9781611974782.142>.

[66] M. LEWENSTEIN, S. PETTIE, AND V. V. WILLIAMS, *Structure and hardness in P (Dagstuhl Seminar 16451)*, Dagstuhl Reports, 6 (2016), pp. 1–34, <https://doi.org/10.4230/DagRep.6.11.1>.

[67] H. NAGAMOCHI AND T. IBARAKI, *A linear-time algorithm for finding a sparse k -connected spanning subgraph of a k -connected graph*, Algorithmica, 7 (1992), pp. 583–596.

[68] D. NANONGKAI, T. SARANURAK, AND C. WULFF-NILSEN, *Dynamic minimum spanning forest with subpolynomial worst-case update time*, in Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science, 2017, pp. 950–961.

[69] M. PARTER, *Fault-tolerant logical network structures*, Bull. Eur. Assoc. Theor. Comput. Sci. EATCS, 118 (2016).

[70] M. PARTER AND D. PELEG, *Sparse fault-tolerant BFS trees*, in Proceedings of the 21st Annual European Symposium on Algorithms, 2013, pp. 779–790, <https://doi.org/10.1007/978-3-642-40450-4-66>.

[71] M. PARTER AND D. PELEG, *Fault tolerant approximate BFS structures*, in Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms, 2014, pp. 1073–1092, <https://doi.org/10.1137/1.9781611973402.80>.

[72] M. PĂTRAȘCU AND M. THORUP, *Time-space trade-offs for predecessor search*, in Proceedings of the 38th ACM Symposium on Theory of Computing, 2006, pp. 232–240.

[73] M. PĂTRAȘCU AND M. THORUP, *Planning for fast connectivity updates*, in Proceedings of the 48th IEEE Symposium on Foundations of Computer Science, 2007, pp. 263–271.

[74] S. PETTIE AND V. RAMACHANDRAN, *A randomized time-work optimal parallel algorithm for finding a minimum spanning forest*, SIAM J. Comput., 31 (2002), pp. 1879–1895.

[75] J. A. L. POUTRÉ, J. VAN LEEUWEN, AND M. H. OVERMARS, *Maintenance of 2- and 3-edge-connected components of graphs I*, Discrete Math., 114 (1993), pp. 329–359, [https://doi.org/10.1016/0012-365X\(93\)90376-5](https://doi.org/10.1016/0012-365X(93)90376-5).

- [76] M. PĂTRAŞCU AND M. THORUP, *Randomization does not help searching predecessors*, in Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms, 2007, pp. 555–564.
- [77] M. PĂTRAŞCU AND M. THORUP, *Dynamic integer sets with optimal rank, select, and predecessor search*, in Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science, 2014, pp. 166–175, <https://doi.org/10.1109/FOCS.2014.26>.
- [78] J. SHERMAN, *Breaking the multicommodity flow barrier for $O(\sqrt{\log n})$ -approximations to sparsest cut*, in Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science, 2009, pp. 363–372, <https://doi.org/10.1109/FOCS.2009.66>.
- [79] W.-B. STROTHMANN, *Bounded Degree Spanning Trees*, dissertation, Universität Paderborn, Heinz Nixdorf Institut, Theoretische Informatik, 1997.
- [80] P. VAN EMDE BOAS, R. KAAS, AND E. ZIJLSTRA, *Design and implementation of an efficient priority queue*, Math. Syst. Theory, 10 (1977), pp. 99–127.
- [81] Z. WANG, *An Improved Randomized Data Structure for Dynamic Graph Connectivity*, CoRR, abs/1510.04590, 2015, <http://arxiv.org/abs/1510.04590>.
- [82] J. WESTBROOK AND R. E. TARJAN, *Maintaining bridge-connected and biconnected components on-line*, Algorithmica, 7 (1992), pp. 433–464, <https://doi.org/10.1007/BF01758773>.
- [83] C. WULFF-NILSEN, *Faster deterministic fully-dynamic graph connectivity*, in Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms, 2013, pp. 1757–1769.