

Multiple-Edge-Fault-Tolerant Approximate Shortest-Path Trees*

Davide Bilò¹, Luciano Gualà², Stefano Leucci³, and Guido Proietti⁴

- 1 DUMAS, Università di Sassari, Italy
davide.bilo@uniss.it
- 2 DII, Università di Roma “Tor Vergata”, Italy
guala@mat.uniroma2.it
- 3 DISIM, Università degli Studi dell’Aquila, Italy; and
DI, “Sapienza” Università di Roma, Italy
Stefano.leucci@univaq.it
- 4 DISIM, Università degli Studi dell’Aquila, Italy; and
IASI, CNR, Roma, Italy
guido.proietti@univaq.it

Abstract

Let G be an n -node and m -edge positively real-weighted undirected graph. For any given integer $f \geq 1$, we study the problem of designing a sparse f -edge-fault-tolerant (f -EFT) σ -approximate single-source shortest-path tree (σ -ASPT), namely a subgraph of G having as few edges as possible and which, following the failure of a set F of at most f edges in G , contains paths from a fixed source that are stretched at most by a factor of σ . To this respect, we provide an algorithm that efficiently computes an f -EFT $(2|F| + 1)$ -ASPT of size $O(fn)$. Our structure improves on a previous related construction designed for *unweighted* graphs, having the same size but guaranteeing a larger stretch factor of $3(f + 1)$, plus an additive term of $(f + 1) \log n$.

Then, we show how to convert our structure into an efficient f -EFT *single-source distance oracle* (SSDO), that can be built in $\tilde{O}(fm)$ time, has size $O(fn \log^2 n)$, and is able to report, after the failure of the edge set F , in $O(|F|^2 \log^2 n)$ time a $(2|F| + 1)$ -approximate distance from the source to any node, and a corresponding approximate path in the same amount of time plus the path’s size. Such an oracle is obtained by handling another fundamental problem, namely that of updating a *minimum spanning forest* (MSF) of G after that a *batch* of k simultaneous edge modifications (i.e., edge insertions, deletions and weight changes) is performed. For this problem, we build in $O(m \log^3 n)$ time a *sensitivity oracle* of size $O(m \log^2 n)$, that reports in $O(k^2 \log^2 n)$ time the (at most $2k$) edges either exiting from or entering into the MSF. As a result of independent interest, it is worth noticing that our MSF oracle can be employed to handle arbitrary sequences of $o(\sqrt[4]{n}/\log n)$ (non-simultaneous) updates with a worst-case time per update of $o(\sqrt{n})$. Thus, for relatively short sequences of updates, our oracle should be preferred w.r.t. the best-known (in a worst-case sense) MSF *fully-dynamic* algorithm, requiring $O(\sqrt{n})$ time per update.

1998 ACM Subject Classification G.2.2 [Graph Theory] Graph algorithms, Trees

Keywords and phrases fault-tolerant shortest-path tree, distance oracle, minimum spanning tree

Digital Object Identifier 10.4230/LIPIcs.STACS.2016.18

* This work was partially supported by the Research Grant PRIN 2010 “ARS TechnoMedia”, funded by the Italian Ministry of Education, University, and Research.

1 Introduction

Let $G = (V(G), E(G), w)$ be a positively real-weighted undirected graph of n nodes and m edges. A *shortest-path tree* (SPT) of G rooted at a distinguished source vertex, say s , is one of the most popular structures in communication networks. For example, it can be used for implementing the fundamental *broadcasting* operation. However, the SPT, as any tree-based topology, is highly sensitive to edge/vertex failures, which cause the undesired effect of disconnecting sets of vertices from the source.

Therefore, a general approach to cope with this scenario is to make the SPT resistant against a given number of component failures, by adding to it a set of suitably selected edges from the underlying graph, so that the resulting structure will still contain an SPT of the surviving network. If we prepare ourselves to resist against a set of at most f failing edges in G , then the corresponding structure will be named an *f-edge-fault-tolerant* (f -EFT) SPT. Unfortunately, it can be seen that even if $f = 1$ and $m = \Theta(n^2)$, then $\Theta(m)$ additional edges may be needed, as will be shown in the full version of this paper. Thus, to sparsify such a structure, it makes sense to resort to *approximate* shortest paths from the source, that are stretched at most by a factor $\sigma > 1$, for any possible set of failures that has to be handled.

In this paper, we show how to build¹ an efficient structure of this sort. Moreover, we show that it is possible to transform such a structure into an efficient *oracle* that will allow to quickly switch to the alternative paths in case a set of failures will take place.

1.1 Related Work

In the recent past, several single and multiple edge/vertex-fault-tolerant *approximate* SPT (ASPT) structures have been devised. More formally, we say that a spanning subgraph H of G is an f -EFT σ -ASPT if it satisfies the following condition: For each set of edges $F \subseteq E(G)$ of size at most f , all the distances from the source s in the subgraph $H - F = (V(G), E(H) \setminus F, w)$ are at most σ times longer than the corresponding distances in $G - F$. Similar definitions can be given for the *vertex-fault-tolerant* (VFT) case.

A natural counterpart of fault-tolerant SPT structures are fault-tolerant σ -stretched *single-source distance oracles* (σ -SSDO in the following), i.e., *compact* data structures that can be built with a *low* preprocessing time, and that are able to *quickly* return σ -approximate distances/paths from the source following a set of failures. Converting a fault-tolerant SPT into a corresponding SSDO with the very same stretch, and additionally having a small size and a fast query time, is a quite natural process, because of its practical usage: computing the alternative post-failure distances/paths on the structure may indeed be very time consuming. However, such a conversion process is not straightforward, in general, since it requires to exploit distance-related information that are instead implicit in the underlying structure, and this has to be done by optimizing the trade-off between the size and the query time of the oracle.

Turning back our attention to fault-tolerant SPT structures, their study originated in [4], where the authors built in $O(m \log n + n \log^2 n)$ time a 1-VFT 3-SSDO of size $O(n \log n)$, and, for *unweighted* graphs, in $O(m\sqrt{n/\varepsilon})$ time a 1-VFT $(1 + \varepsilon)$ -SSDO of size $O(\frac{n}{\varepsilon^3} + n \log n)$, for any $\varepsilon > 0$, both having a distance (resp., path) query time of $O(1)$ (resp., proportional to the path's size). In such a paper, the authors observe explicitly that, as a result of

¹ Throughout this introduction, all the discussed structures are poly-time computable, even if we may omit to specify the actual running time.

independent interest, the latter oracle (but actually the former as well) can be converted into a corresponding structure (i.e., a spanning subgraph), having the same size and stretch. For the weighted case, the obtained 1-VFT 3-ASPT of size $O(n \log n)$ was then substantially improved in [8], where the authors showed the existence of a 1-E/VFT $(1 + \varepsilon)$ -ASPT of size $O(\frac{n \log n}{\varepsilon^2})$, for any $\varepsilon > 0$ (without providing a corresponding oracle).

Concerning unweighted graphs, Parter and Peleg in [30] presented a 1-E/VFT *Breadth-First Search tree* (BFS) of size $O(n \cdot \min\{ecc(s), \sqrt{n}\})$, where $ecc(s)$ denotes the eccentricity of the source vertex s in G , namely a structure containing *exact* shortest paths from the source after a single edge/vertex failure. In the same paper, the authors also exhibit a corresponding lower bound of $\Omega(n^{3/2})$ for the size of a 1-E/VFT BFS. Then, in [31], the same authors presented a set of lower and upper bounds to the size of fault-tolerant (σ, β) -ABFS, where a further additive distortion β is allowed to the distances. More precisely, they showed that for every $\beta \in [1, O(\log n)]$, there exists a graph G and a source vertex $s \in V(G)$ such that a corresponding 1-EFT $(1, \beta)$ -ABFS requires $\Omega(n^{1+\varepsilon(\beta)})$ edges, for some function $\varepsilon(\beta) \in (0, 1)$. Moreover, they also constructed a 1-EFT $(1, 4)$ -ABFS of size $O(n^{4/3})$. Finally, assuming at most $f = O(1)$ edge failures can take place, they showed the existence of (i) an f -EFT $(3(f + 1), (f + 1) \log n)$ -ABFS of size $O(fn)$, and (ii) an f -EFT $(3f + 4)$ -ABFS of size $O(fn \log^{f+1} n)$. These structures will be exactly our touchstone in this paper, since they are the only ones concerned with multiple-edge-failure single-source shortest paths.

1.2 Our Results

In this paper, we present the following main results:

- An f -EFT $(2|F| + 1)$ -ASPT of size $O(fn)$ that is able to handle the failure of any set $F \subseteq E(G)$ of at most f edges. This considerably improves w.r.t. to its direct competitors, namely the structures presented in [31]: our structure has a size that is never worse, a lower stretch, works on weighted graphs, and handles an arbitrary (i.e., even non-constant) number of failures. Moreover, our construction is simpler and can be computed quickly in $O(fm \alpha(m, n))$ time, where α is the inverse of the Ackermann's function.
- A corresponding f -EFT $(2|F| + 1)$ -SSDO of size $O(\min\{m, fn\} \log^2 n)$, that has a query time for a post-failure distance from the source of $O(|F|^2 \log^2 n)$, and is also able to report the corresponding path in the same time plus the path size. The preprocessing time is $O(fm \alpha(m, n) + fn \log^3 n)$. Notice that if one is willing to use $O(m \log^2 n)$ space, then our oracle will be prepared to handle any number of edge failures (i.e., up to m).

Interestingly enough, the former result is obtained by posing a simple yet surprising relationship between the structure of the replacement paths and the *minimum spanning forest* (MSF) of an ad-hoc auxiliary graph. This approach is also useful to develop the latter result, that is indeed obtained through an efficient updating of an MSF after that a *batch* of k edge modifications (i.e., edge insertions, deletions and weight changes) are simultaneously performed. For this problem indeed we provide the following result:

- a *sensitivity oracle*² of size $O(m \log^2 n)$, that can be built in $O(m \log^3 n)$ time, and is able to report in $O(k^2 \log^2 n)$ time the (at most $2k$) edges either exiting from or entering into the MSF. As a result of independent interest, it is worth noticing that our oracle can be used to efficiently maintain an MSF under relatively short sequences of *non-simultaneous* updates. Indeed, observe that a sequence $\lambda = \langle \lambda_1, \dots, \lambda_h \rangle$ of updates can

² We use this noun for the oracle in accordance with its functionality of only reporting the updates in the MSF.

be managed through h sequential queries to the oracle, where the i -th query will involve the modifications to the starting MSF induced by the batch of the first i updates. This way, we spend $O(h^2 \log^2 n)$ time to handle each single update. Hence, as the fastest long-standing algorithm for the classic (and clearly more general) *fully-dynamic* MSF problem has a worst-case cost of $O(\sqrt{n})$ per update [20], it follows that for $h = o(\sqrt[4]{n}/\log n)$, our oracle should be preferred, since it will manage each update in $o(\sqrt{n})$ time. Notice also that a comparison with other known online/offline algorithms for maintaining an MSF that are more efficient in an amortized sense, like for instance those given in [19, 25, 26], is unfeasible, since they need to start from an empty graph to guarantee their bounds (or, they need long sequences of updates to become efficient). Thus, when starting from an arbitrary graph, as it happens in our setting, a single update operation could even cost them $\Theta(n)$ time!

Finally, we point out that we are also able to prove a lower bound of $\Omega(n^{1+\frac{1}{k}})$ on the size of any f -EFT σ -ASPT with $f \geq \log n$ and $\sigma < \frac{3k+1}{k+1}$, that holds if the Erdős' girth conjecture is true. Our lower bound shows that, in contrast to the single-edge failure case, it is not possible to obtain a stretch arbitrary close to 1 with size $\tilde{O}(n)$ when the number of faults is more than $\log n$. We look at the problem of understanding if this can be done for constant $f > 1$ as an interesting open problem. Due to space limitations, this result will be given in the full version of the paper.

1.3 Other Related Work on Fault-Tolerant Single/Multiple-Source Structures/Oracles

Besides the papers mentioned before, several other research efforts have been devoted to structures and oracles for tolerating single/multiple failures in single-source shortest paths. An early work on the topic is [27], where the authors were concerned with the computation of *best swap edges* (w.r.t. several swap functions) for the failure of each and every edge in an SPT. As a by-product of their results, it can be easily seen that by adding to an SPT the (at most) $n - 1$ best swap edges w.r.t. to the new distance from s to the root of the subtree disconnected from s after an edge failure, then a 1-EFT 3-ASPT is obtained. Interestingly, such a structure can be easily converted into a 1-EFT 3-SSDO of size $O(n)$ and query time $O(1)$. Recently, in [15], the authors faced the special case of *shortest-path failures*, in which the failure of a set F of at most f adjacent edges along any source-leaf path has to be tolerated. They proposed an f -EFT $(2k - 1)(2|F| + 1)$ -ASPT of size $O(kn f^{1+1/k})$, where $|F|$ denotes the size of the actual failing path, and $k \geq 1$ is a parameter of choice. Notice that this result is subsumed by ours. Moreover, they also provided a conversion to a corresponding oracle, and for the special case of $f = 2$, they gave an ad-hoc solution of size $O(n \log n)$ and with stretch 3. For directed graphs with integer positive edge weights bounded by M , in [23] the authors showed how to build efficiently in $\tilde{O}(Mn^\omega)$ time a randomized 1-EFT 1-SSDO of size $\Theta(n^2)$ and with $O(1)$ query time, where returned distances are exact w.h.p., and $\omega < 2.373$ denotes the matrix multiplication exponent.

Concerning unweighted graphs, in [8] the authors showed that an *ordinary* (i.e., non fault-tolerant) (σ, β) -*spanner* (i.e., where distances/paths between arbitrary pairs of nodes are at most (σ, β) -stretched) of size $O(g(n))$ can be used to build a 1-EFT (resp., VFT) (σ, β) -ABFS of the same size (resp., of size $O(g(n) + n \log n)$). This result is useful for building sparse 1-VFT $(1, \beta)$ -ABFS structures by making use of the vast literature on *additive* $(1, \beta)$ -spanners (e.g., [5, 11]). Finally, Parter in [29] presented a 2-EFT BFS having $O(n^{5/3})$ edges, which is tight.

Another research stream related to our work is that on *multi-source* (MS) fault-tolerant structures, for which we look at distances/paths from a set $S \subseteq V(G)$ of sources. Here, results are known only for unweighted graphs. In [30] the authors gave an algorithm to compute a 1-EFT MSBFS of size $O(\sqrt{|S|} n^{3/2})$, which is tight. Then, in [8] it was shown that an ordinary (σ, β) -spanner of size $O(g(n))$ can be used to build a 1-EFT (σ, β) -AMSBFS of size $O(g(n) + n|S|)$, and similarly for the vertex case of size $O(g(n) + n \log n|S|)$.

1.4 More Related Work on (Fault-Tolerant) Spanners/Oracles

For the sake of completeness, we also give some hints on the large body of literature on the related topic of (fault-tolerant) spanners and distance oracles.

On weighted graphs, the currently best known construction is, for any $f \geq 1$ and any integer parameter $k \geq 1$, the f -EFT (resp., VFT) $(2k - 1)$ -spanner of size $O(f n^{1+1/k})$ (resp., $\tilde{O}(f^2 k^{f+1} n^{1+1/k})$) given in [13]. For the vertex-failure case, this has been then improved in a randomized sense in [16], where the expected size was reduced to $\tilde{O}(f^{2-1/k} n^{1+1/k})$. For a comparison, the sparsest known $(2k - 1)$ -multiplicative ordinary spanner has size $O(n^{1+1/k})$ [2], and this is believed to be asymptotically tight due to the girth conjecture of Erdős [21]. Then, in [3] it was introduced the resembling concept of 1-EFT *resilient* spanners, i.e., spanners such that whenever any edge in G fails, then the relative distance increases in the spanner are very close to those in G .

Ordinary (i.e., fault-free) *all-pairs distance oracles* (APDO) on weighted graphs were introduced in a seminal work by Thorup and Zwick [32] (who also coined the term *oracle*), followed by a sequel of papers (among the others, we mention [12, 18] for the currently best bounds). In a fault-tolerant setting, in [6] the authors built (on directed graphs) a 1-E/VFT 1-APDO of size $\tilde{O}(n^2)$ and with query time $O(1)$. For two failures, in [17] the authors built, still on directed graphs, a 2-E/VFT 1-APDO of size $\tilde{O}(n^2)$ and with query time $O(\log n)$. Concerning multiple-edge failures, in [14] the authors built, for any integer $k \geq 1$, an f -EFT $(8k - 2)(f + 1)$ -APDO of size $O(fk n^{1+1/k} \log(nW))$, where W is the ratio of the maximum to the minimum edge weight in G , and with a query time of $\tilde{O}(|F| \log \log d)$, where F is the actual set of failing edges, and d is the distance between the queried pair of nodes in $G - F$.

On unweighted graphs, it makes instead sense to study fault-tolerant additive spanners. In particular, Braunschvig et al. [9] proposed the following general approach to build an f -EFT additive spanner: Let A be an f -EFT σ -spanner, and let B be an ordinary $(1, \beta)$ -spanner. Then $H = A \cup B$ is an f -EFT $(1, 2f(2\beta + \sigma - 1) + \beta)$ -spanner. Recently, in [7] the corresponding analysis has been refined yielding a better additive bound of $2f(\beta + \sigma - 1) + \beta$. Finally, for other results on single edge/vertex failures spanners/oracles on unweighted graphs, we refer the reader to [4, 28, 7].

2 An f -EFT $(2|F| + 1)$ -ASPT and a Corresponding Oracle

In this section we show how to compute an f -EFT $(2|F| + 1)$ -ASPT H of G . When up to f edges can fail, it is easy to see that whenever G is $(f + 1)$ -edge-connected, H must contain $\Omega(fn)$ edges even if we are only interested in preserving the connectivity of G , since the degree of each vertex must be at least equal to $f + 1$. Here we show that $|E(H)| = O(fn)$ edges also suffice if we seek to preserve distances that are at most $(2f + 1)$ -stretched w.r.t. the surviving part of G .

Let $d_X(u, u')$ and $\pi_X(u, u')$ denote the distance and the shortest path between nodes u and u' in any subgraph X of G , respectively. When $u = s$, we will simply write $d_X(u')$ and $\pi_X(u')$. If π is a path, $\pi[u, u']$ will denote the subpath of π between $u, u' \in V(\pi)$.

Algorithm 1: Algorithm for computing an f -EFT $(2|F| + 1)$ -ASPT of G .

```

1  $T \leftarrow$  compute an SPT of  $G$ 
2 for  $(u, v) \in E(G)$  do
3   if  $(u, v) \in E(T)$  then  $w'(u, v) \leftarrow 0$ 
4   else  $w'(u, v) \leftarrow d_T(u) + w(u, v) + d_T(v)$ 
5  $G' \leftarrow (V(G), E(G), w')$ 
6  $G_0 \leftarrow G'$ 
7 for  $i = 0, \dots, f$  do
8    $M_i \leftarrow$  edges of an MSF of  $G_i$  (w.r.t.  $w'$ )
9    $G_{i+1} \leftarrow G_i \setminus M_i$ 
10  $H \leftarrow$  subgraph of  $G$  containing the edges in  $\bigcup_{i=0}^f M_i$ 
11 return  $H$ 

```

For any given integer f , Algorithm 1 returns an f -EFT $(2|F| + 1)$ -ASPT of G . First, it computes an SPT T of G that is used to assign a weight to the edges of an auxiliary graph $G' = (V(G), E(G), w')$. More precisely, the weight of an edge e of G' is 0 if e is also in T , otherwise it is equal to the sum of the corresponding edge weight in G and the distances in T between s and the endpoints of e . Then, $f + 1$ MSFs M_0, \dots, M_f of G' are iteratively computed: when we compute the i -th forest, we remove its edges M_i from G' before computing the $(i + 1)$ -th forest, so that the sets M_i are pairwise disjoint. The sought subgraph H contains all the edges of the sets M_i . Notice that M_0 coincides with $E(T)$.

We now argue that H is indeed an f -EFT $(2|F| + 1)$ -ASPT of G . Fix a vertex t and let $\pi = \pi_{G-F}(t)$ be the shortest path from s to t in the surviving graph $G - F$.³ The path π traverses the vertices of several trees in the forest $T - F$. We say that an edge is *new* if its endpoints belong to two different trees in $T - F$. Let N be the set of new edges in π .

Now consider an MSF M of the graph $H - F$ (w.r.t. w'). This is also an MSF of the graph $G' - F$ (w.r.t. w') as shown by the following lemma.

► **Lemma 1.** *For every $F \subseteq E(G)$ with $|F| \leq f$, any MSF M of $H - F$ (w.r.t. w') is also an MSF of $G' - F$ (w.r.t. w').*

Proof. In what follows, whenever ties arise we break them by prioritizing the edges in H . First we show that, given any cut-set⁴ C of G' , H contains the $\min\{|C|, f + 1\}$ lightest edges of C . Indeed, for any set M_i , consider the set $C_i = C \setminus \bigcup_{j=0}^{i-1} M_j$. Either C_i is non empty, and therefore M_i contains the lightest edge in C_i , or $C_i = \emptyset$ which means that each edge in C belongs to some set M_j and hence to H .

Let M' be an MSF of $G' - F$. We prove the claim by showing that each edge $e \in E(M')$ must also belong to M . Let C' be the cut-set of G' that contains e and every edge $e' \in E(G')$ that forms a cycle with e in $M' \cup \{e'\}$. Since e is the lightest edge of $C' \setminus F$, it is within the $f + 1$ lightest edges of C' . As a consequence $e \in E(H - F)$, and it also belongs to M as it is the lightest edge in $C' \cap E(H - F)$. ◀

Let $\pi' = \pi_M(s, t)$ and notice that π' traverses each tree of the forest $T - F$ at most once since edges in $E(T)$ have weight 0 in H . Once again, let N' be the set of new edges of π' . We now provide an upper bound to the distance $d_{H-F}(t)$ using the path π' :

³ We assume that such a path exists, as otherwise $d_{G-F}(t) = +\infty$ which implies $d_{H-F}(t) = +\infty$, and we are done.

⁴ A cut-set of a graph X is a subset of $E(X)$ whose removal increases the number of connected components of X .

► **Lemma 2.** $d_{H-F}(t) \leq w(\pi') \leq \sum_{e \in N'} w'(e) + d_G(t)$.

Proof. Let M be an MSF of the graph $H - F$ (w.r.t. w'). The first inequality is trivial as $\pi' = \pi_M(s, t)$ is a path (not necessarily shortest) between s and t in (a subgraph of) $H - F$, hence we focus on proving the second inequality.

Let T_0, \dots, T_h be the trees of $T - F$ traversed by π' , in order, and let $e'_i = (v_{i-1}, u_i)$ be the new edge in π' connecting a vertex v_{i-1} of T_{i-1} to a vertex u_i of T_i . In such a way we have $N' = \{e'_1, \dots, e'_h\}$. We call r_i the vertex in $V(T_i) \cap V(\pi')$ that has the lowest depth in T_i .⁵ According to this definition, r_0 coincides with s , r_h is the lowest common ancestor between u_h and t , and r_i is the lowest common ancestor between u_i and v_i , for every $0 < i < h$.

We prove by induction on i that $w(\pi'[s, r_i]) \leq \sum_{j=1}^i w'(e'_j)$. The base case $i = 0$ is trivially true. Now suppose that the inductive hypothesis holds for i , we prove it also for $i + 1$:

$$\begin{aligned} w(\pi'[s, r_{i+1}]) &= w(\pi'[s, r_i]) + d_{T_i}(r_i, v_i) + w(e'_{i+1}) + d_{T_{i+1}}(u_{i+1}, r_{i+1}) \\ &\leq \sum_{j=1}^i w'(e'_j) + d_T(v_i) + w(e'_{i+1}) + d_T(u_{i+1}) = \sum_{j=1}^i w'(e'_j) + w'(e'_{i+1}) = \sum_{j=1}^{i+1} w'(e'_j). \end{aligned}$$

We now use the fact that $d_{T_h}(r_h, t) = d_T(r_h, t) = d_G(r_h, t)$ to prove the claim:

$$w(\pi') = w(\pi'[s, r_h]) + w(\pi'[r_h, t]) \leq \sum_{j=1}^h w'(e'_j) + d_{T_h}(r_h, t) \leq \sum_{j=1}^h w'(e'_j) + d_G(t). \quad \blacktriangleleft$$

Next lemma shows that the weights of the new edges of π' are, in turn, upper bounded by the weight of some new edge of the path π .

► **Lemma 3.** For each $e' \in N'$, we have $w'(e') \leq \max_{e \in N} w'(e)$.

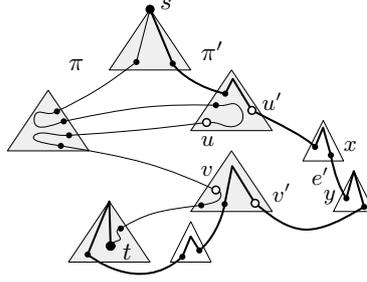
Proof. Let $e' = (x, y)$ be an arbitrary edge in N' . W.l.o.g., we assume that the path π' traverses the vertices s, x, y, t in this order. We recall that the path π' traverses each tree in $T - F$ at most once, i.e., all the vertices of π' that belong to the same tree in $T - F$ must be contiguous in π' . Moreover, as e' is new, x and y belong to two different trees in $T - F$.

Let Z be the set of trees of the forest $T - F$ that are traversed by the path π . Let u' be the last vertex of $\pi'[s, x]$ that belongs to a tree, say T_u , of Z (see Figure 1). Observe that u' is always defined since s belongs to some tree of Z . In a similar way, let v' be the first vertex of $\pi'[y, t]$ that belongs to a tree, say T_v , of Z . Again, observe that v' is always defined as t belongs to some tree of Z other than that containing s , and so $T_u \neq T_v$, and finally notice that $e' \in E(\pi'[u', v'])$. We know that π traverses both T_u and T_v (in some order), so we let π^* be the minimal (w.r.t. inclusion) subpath of π with one endpoint, say u , in $V(T_u)$, and the other endpoint, say v , in $V(T_v)$.

Let $N^* = E(\pi^*) \cap N$ be the set of new edges in π^* . Notice that $N^* \neq \emptyset$ as $T_u \neq T_v$, and that adding the edges in N^* (weighted according to w') to M forms (at least) a cycle C containing both e' and an edge in N^* , say e^* . Since M is an MSF of $G' - F$, as shown by Lemma 1, we have that $w'(e') \leq w'(e^*) \leq \max_{e \in N} w'(e)$. ◀

Finally, next lemma relates the weights w' of the new edges of π to distances in the surviving graph $G - F$.

⁵ We think of T_i as rooted in the vertex of $V(T_i)$ which is closest to s in T .



■ **Figure 1** The forest $T - F$ obtained by deleting the failed edges in F from T . The path π is the shortest path between s and t in $G - F$, while π' (in bold) is the unique path in M between the same vertices. Gray trees contain a vertex of π and are therefore in Z . Edges having endpoints in different trees are *new*.

► **Lemma 4.** For $e \in N$, $w'(e) \leq 2d_{G-F}(t)$.

Proof. Let $e = (u, v)$ with $d_{G-F}(u) \leq d_{G-F}(v)$. Since e lies on the shortest path $\pi = \pi_{G-F}(s, t)$, we can write:

$$w'(e) = d_T(u) + w(e) + d_T(v) \leq d_{G-F}(v) + d_T(v) \leq 2d_{G-F}(v) \leq 2d_{G-F}(t). \quad \blacktriangleleft$$

We are now ready to prove the main result of this section:

► **Theorem 5.** The graph H returned by Algorithm 1 is an f -EFT $(2|F| + 1)$ -ASPT of G . Moreover, Algorithm 1 requires $O(fm\alpha(m, n))$ time and $O(m)$ space.

Proof. First, observe that $\pi' = \pi_M(s, t)$ contains at most $|F|$ new edges. Indeed all the edges in $T - F$ have weight 0, while the remaining edges have a positive weight. This means that $E(T - F) \subseteq E(M)$. As $T - F$ has no more than $|F| + 1$ connected components, we have that at most $|F|$ other edges – which are not in $E(T - F)$ – can belong to M .

By using the above fact in conjunction with Lemmas 2–4, we can write:

$$\begin{aligned} d_{H-F}(t) &\leq w(\pi') \leq \sum_{e \in N'} w'(e) + d_G(t) \leq |F| \max_{e \in N'} w'(e) + d_G(t) \\ &\leq |F| \max_{e \in N} w'(e) + d_G(t) \leq 2|F|d_{G-F}(t) + d_{G-F}(t) = (2|F| + 1)d_{G-F}(t). \end{aligned} \quad (1)$$

We recall that this holds for every vertex $t \in V(G)$. Concerning the computational complexity of Algorithm 1, we make use of Chazelle’s algorithm [10] – that computes an MSF in $O(m\alpha(m, n))$ time and linear space – to compute the $f + 1$ MSFs M_0, \dots, M_f . ◀

2.1 A Corresponding Oracle

In this section we show how to build an oracle that, given a positively real-weighted graph G and a distinguished source vertex s , is able to answer queries of the form: *Given a set F of at most f edge failures, and a destination node t in G , report a $(2|F| + 1)$ -approximate path/distance from s to t in $G - F$.*

We first compute an SPT T of G and a f -EFT $(2|F| + 1)$ -ASPT H of G , as shown in the previous section. Then, the oracle is composed of three ingredients:

- the tree T and all the distances $d_T(v) = d_G(v)$ from s to any vertex $v \in V(G)$;
- an MSF sensitivity oracle Q of H w.r.t. the weights w' , built as shown in Section 3;
- an oracle to answer *lowest common ancestor* (LCA) queries between two vertices in T .

Such an oracle can be built in linear time and has a constant query time [24].

The resulting size is $O(fn \log^2 n)$ and the time required to build our oracle is $O(fm\alpha(m, n) + fn \log^3 n)$. Interestingly, if we do not know the value of f in advance, we can build, in $O(m \log^3 n)$ time, an oracle of size $O(m \log^2 n)$ that is able to report $(2|F| + 1)$ approximate paths/distances, for *any* number $|F|$ of faults.

We will make use of the following additional property of our MSF oracle Q , that will be shown in Section 3: *Q can report, in $O(|F|^2 \log^2 n)$ time, all the new edges (and their weights), on the unique path from s to t in the updated MSF, in order.*

Answering a Path Query

To return a $(2|F| + 1)$ -approximate path between s and t , it suffices to report the path $\pi' = \pi_M(s, t)$, as shown by Equation (1).

We query the MSF oracle Q for the new edges on the unique path from s to t in the updated MSF. Let $\langle e'_1, \dots, e'_h \rangle$ be these new edges, in order, with $e'_i = (v_{i-1}, u_i)$. For $0 < i < h$, let r_i be the LCA between u_i and v_i , and let r_h be the LCA between u_h and t . We now have all the pieces to reconstruct and return the path π' . Indeed, if we let $\pi'_i = \pi_T(u_i, r_i) \circ \pi_T(r_i, v_i)$, the following holds:

$$\pi' = \pi_T(s, v_0) \circ e'_1 \circ \pi'_1 \circ e'_2 \circ \pi'_2 \circ \dots \circ \pi'_{h-1} \circ e'_h \circ \pi_T(u_h, r_h) \circ \pi_T(r_h, t) \quad (2)$$

where each subpath is entirely in T and all the endpoints are known. The whole procedure requires $O(|F|^2 \log^2 n)$ time to perform the query on Q , $O(|F|)$ time for the LCA queries, and $O(|\pi'|)$ time to reconstruct the path. The overall query time is therefore $O(|F|^2 \log^2 n + |\pi'|)$.

Answering a Distance Query

To report the length of a $(2|F| + 1)$ -approximate path from s to t , we can replace each subpath in Equation (2) with the corresponding distance, in order to obtain:

$$w(\pi') = d_T(v_0) + \sum_{i=1}^{h-1} \left(w(e'_i) + d_T(u_i, r_i) + d_T(r_i, v_i) + w(e'_{i+1}) \right) + d_T(u_h, r_h) + d_T(r_h, t).$$

The above quantity can be computed in $O(h) = O(|F|)$ time, once we know the edges e_1, \dots, e_h and we notice that $w(e'_i) = w'(e'_i) - d_T(v_{i-1}) - d_T(u_i)$, and that if x is a descendant of r_i in T , then $d_T(r_i, x) = d_T(x) - d_T(r_i)$. The overall query time is thus $O(|F|^2 \log^2 n)$.

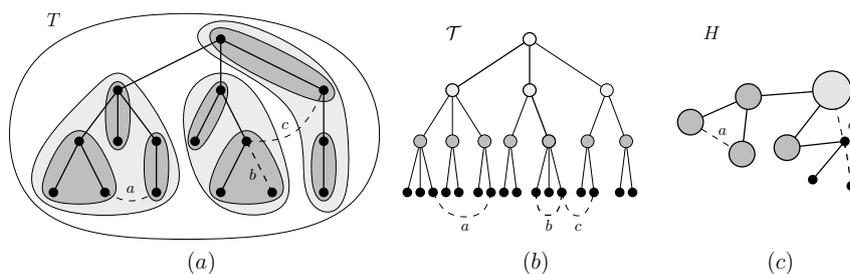
3 A Minimum Spanning Tree Sensitivity Oracle

In this section we present an oracle that, given a real-weighted graph G with n vertices and m edges, along with any *minimum spanning tree* (MST) T of G , is able to answer queries of the form:

“Given a set of k edge updates on G (i.e., edge insertions, deletions and weight modifications), let T' be the new MST of G . What are the edges in the symmetric difference of $E(T)$ and $E(T')$?”⁶

In other words, the oracle can report all the edges of T that leave the MST as a consequence of the updates, along with all the new edges in T' that enter the MST in their place. The

⁶ For the sake of avoiding technicalities, we assume that each edge is subject to at most a single update and we also assume that the graph G always remains connected, so that we simply talk about an MST instead of an MSF of G . For instance, this can be easily guaranteed by adding a dummy vertex $x \notin V(G)$ that is connected to all the vertices of $V(G)$ with edges of large weights.



■ **Figure 2** Hierarchical clustering of the vertices of T . (a) shows all the sets of \mathcal{C} , where black vertices are singletons. (b) shows the tree \mathcal{T} associated with the clustering \mathcal{C} . (c) shows the graph H whose vertices are computed by Algorithm 2. The edges in F are depicted using dashed lines.

oracle requires $O(m \log^2 n)$ space and can be built $O(m \log^2 n)$ space and $O(m \log^3 n)$ time, while a query involving k updates can be answered in $O(k^2 \log^2 n)$ time and space.

Our oracle exploits the fact that, when few updates are to be handled, the changes in the resulting MST will be small. This implies that large portions of T and T' will coincide, and knowing these portions would allow us to save a considerable amount of work compared to the time needed to recompute T' from scratch. To this aim, we build a structure that maintains a set of connected subtrees of T at different levels of granularity.

In details, we will use a hierarchical clustering of the vertices of T . Our clustering is inspired by the construction of *topology trees* given in [22]. In [22], the author solves the *dynamic* MST problem by using a collection of topology trees that are built on top of an auxiliary graph representing shrunk components of G . We use our clustering in a different way and, as we do not need to support permanent updates of G , we are also able to simplify the construction. Due to space limitations, the full description of our construction will be given in the extended version of the paper, while here we provide a sketch of it.

We start by describing the properties of our hierarchical clustering. Let Δ be the maximum degree⁷ of a vertex in T . Each cluster C will have a level $\ell(C) \in \{0, \dots, L\}$, and we will call \mathcal{C}_i the set of clusters of level i . Our clustering will guarantee that:

- P₁**. Clusters of each level i are a partition of the vertices of T , i.e., they are pairwise disjoint and $\bigcup_{C \in \mathcal{C}_i} C = V(G)$;
- P₂**. The vertices in each cluster induce a connected component of T ;
- P₃**. Clusters of level 0 are singletons, i.e., they contain a single vertex of T ;
- P₄**. There is only one cluster of level L (and it coincides with $V(T)$);
- P₅**. Each cluster of level $i \geq 1$ is the union of at least 2 and at most Δ clusters of level $i - 1$.

It follows from the above properties that a cluster of level i contains at least 2^i vertices, and hence $L \leq \log n$. Figure 2 (a) shows an example of such a clustering. This hierarchy can be represented by a tree \mathcal{T} of height L rooted in the unique cluster in \mathcal{C}_L . The children of a cluster of level $i \geq 1$ in \mathcal{T} are the clusters of level $i - 1$ it contains (see Figure 2 (b)).

For each pair of clusters C, C' with $C \neq C'$ we maintain an ordered set $E(C, C')$ containing all the edges of $E(G)$ with one endpoint in C and the other in C' . This set is ordered according to edge weights in a non-decreasing fashion. Let $C(u)$ be the set of the $L + 1$ clusters of the hierarchy (one for each level) that contain vertex u . It is easy to see that an edge

⁷ In order to compute the clustering, the tree T will be rooted. We still define the degree of a vertex v in T to be the number of edges that are incident v , including the edge from v to its parent in T , if any.

Algorithm 2: Algorithm for computing the set of vertices (i.e., clusters) of H .

```

1  $R \leftarrow \mathcal{C}_L$ 
2 for  $(u, v) \in F$  do
3   while  $\text{root}_{\mathcal{T}}(u) = \text{root}_{\mathcal{T}}(v)$  do
4      $C \leftarrow \text{root}_{\mathcal{T}}(u)$ 
5      $\mathcal{T} \leftarrow \mathcal{T} \setminus \{C\}$ 
6      $R \leftarrow (R \setminus \{C\}) \cup \text{children}_{\mathcal{T}}(C)$  // Split  $C$ 
7 return  $R$ 

```

$(u, v) \in E(G)$ appears in at most $|C(u)| \cdot |C(v)| = O(\log^2 n)$ sets, and hence the overall number of elements in the sets is at most $O(m \log^2 n)$.

We now describe how a query can be answered. In order to do so, it is useful to split each weight update operation involving an edge e into two separate operations, namely the deletion of e followed by its reinsertion with the new (updated) weight. By doing so, all the operations in F are now either insertions or deletions. For the sake of clarity, we first consider the case in which all the updates F are edge deletions, and we will show later how this can be extended to deal also with edge insertions.

Handling Edge Deletions

In order to handle deletions, we use Algorithm 2 to construct an auxiliary graph H whose vertices are clusters. The algorithm will compute a set R of clusters of \mathcal{T} that will coincide with $V(H)$. Initially R contains the unique cluster in \mathcal{C}_L that is the root of \mathcal{T} and represents the whole tree T . At each time, the set of clusters in R , although of different levels, will always form a partition of the vertices of T . The algorithm proceeds iteratively, by considering one after the other the edges of F . When an edge (u, v) is considered, if u and v belong to the same cluster C of R , we *split* C , i.e., we remove C from \mathcal{T} and R , and we add to R all the clusters of level $\ell(C) - 1$ contained in C . In this way \mathcal{T} is always a forest and R contains the roots of its trees.

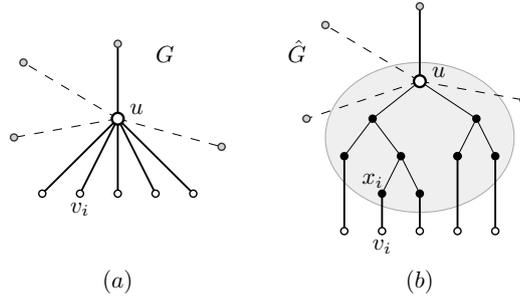
In the end, H is such that all the edges in F have their endpoints into different clusters of $V(H)$. Moreover, as each edge in F can produce at most L splits, and each split operation can increase the number of vertices by at most $\Delta - 1$, we have that H contains at most $fL(\Delta - 1) = O(\Delta f \log n)$ vertices (see Figure 2 (c)).

To construct the set $E(H)$ we consider all the pairs C, C' of vertices in $V(H)$. For each of these pairs we examine the edges in $E(C, C')$, in order, and we select the first edge e so that $e \notin F$, if any. Then, if e exists, we add the edge (C, C') to H with weight $w(e)$.

We can now compute an MST \tilde{T} of H in time $O(\Delta^2 f^2 \log^2 n)$ by using any standard MST algorithm. Finally, we look at the edges of \tilde{T} and we answer the query by returning the edges in $E(\tilde{T})$ that are not in $E(T)$. Notice also that, once \tilde{T} has been computed, it is easy to report all the edges in $E(\tilde{T}) \setminus E(T)$ that belong to the unique path between any two vertices in the updated MST. This kind of query can still be answered in $O(\Delta^2 f^2 \log^2 n)$ time, and it is needed by our fault-tolerant ASPT oracle of Section 2.1.

Handling General Edge Updates

It turns out that the complexity of the problem lies in handling the edge-deletion operations. Indeed, once this has been done, the remaining edge-insertion operations can be easily performed. To this aim, we reorganize the batch by first performing all the delete operations, and we make use of a *top-tree* [1], i.e., a data structure that dynamically maintains a



■ **Figure 3** Reducing the degree of vertices in T : on the left side, the tree T (solid edges) embedded in G , on the right side the superimposition of the binary tree to T in order to get a maximum degree of 3. Thin solid edges have weight 0, while the weight of (x_i, v_i) is $w(u, v_i)$.

(weighted) forest under edge-insertion (*link*) and edge-deletion (*cut*) operations. Moreover, given two vertices u and v , top-trees are able to report the heaviest edge that lies on the path between u and v in the current forest. Each of these operations can be performed in $O(\log \eta)$ time where η is the number of vertices of the forest.

The idea is to maintain the current MST T' by using a top-tree that is initialized when the oracle is built to represent the tree T . This takes $O(n \log n)$ time. Then, we perform all the edge-deletion operations (as already described), while updating the top-tree accordingly (this requires $O(|F| \log n)$ time since the number of needed link and cut operations is $O(|F|)$).

Now we handle the insertions one by one. In order to insert a new edge $e = (u, v)$, we search for the heaviest edge e' of the path connecting u and v in T' . If e' is heavier than e , we cut e' from T' and we link the two resulting components by adding the edge e . It is easy to see that this procedure requires an overall time of $O(|F| \log n)$.

By keeping track of all the $O(|F|)$ updates in the MST T' , we can easily answer a query consisting of both edge-insertion and edge-deletion operations in $O(\Delta^2 f^2 \log^2 n)$ time.

Reducing the Degree of T

So far, the complexity of our oracle depends on the maximum degree Δ of a vertex in T . However, using standard techniques (see, e.g., [22]), we now show that the updates on the original graph G and its MST T can be mapped onto an auxiliary graph \hat{G} with weight function \hat{w} and a corresponding MST \hat{T} , such that \hat{G} has asymptotically the same size of G , and each vertex of \hat{G} has a degree at most 3 in \hat{T} .

Initially \hat{G} , \hat{w} , and \hat{T} coincide with G , w , and T , respectively. We iteratively search for a vertex u in \hat{T} that has more than 2 children, and we lower its degree. Let $\text{children}_{\hat{T}}(u) = \{v_1, \dots, v_h\}$, we proceed as follows: we remove all the edges in $\{(u, v_i) : 1 \leq i \leq h\}$ from both \hat{G} and \hat{T} , then we add to both \hat{G} and \hat{T} a binary tree whose root coincides with u , and that has exactly h leaves x_1, \dots, x_h . We assign weight $\hat{w}(e) = 0$ to all the edges e of this tree. Finally, we add to \hat{G} and \hat{T} an edge (x_i, v_i) for each $1 \leq i \leq h$, and we set $\hat{w}(x_i, v_i) = w(u, v_i)$. An example of such a transformation is shown in Figure 3.

Each time we have to perform a weight update or delete operation on an edge (u, v_i) of G , we instead perform it on the corresponding edge (x_i, v_i) . Insertions and operations involving edges in $E(G) \setminus E(T)$ do not require any special care. In a similar way, whenever the answer of a query contains an edge (x_i, v_i) , we replace it with the corresponding edge (u, v_i) . Clearly, $O(n)$ vertices and edges are added by this process, and hence $|V(G)| = \Theta(|V(\hat{G})|)$ and $|E(G)| = \Theta(|E(\hat{G})|)$.

Once the maximum degree of the tree has been reduced to a constant, the query time of our oracle becomes $O(f^2 \log^2 n)$. To achieve such a query time, however, we must be careful in our implementation as it will be discussed in the full version of the paper.

References

- 1 Stephen Alstrup, Jacob Holm, Kristian de Lichtenberg, and Mikkel Thorup. Maintaining information in fully dynamic trees with top trees. *ACM Transactions on Algorithms*, 1(2):243–264, 2005. URL: <http://doi.acm.org/10.1145/1103963.1103966>, doi:10.1145/1103963.1103966.
- 2 Ingo Althöfer, Gautam Das, David P. Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9:81–100, 1993. URL: <http://dx.doi.org/10.1007/BF02189308>, doi:10.1007/BF02189308.
- 3 Giorgio Ausiello, Paolo Giulio Franciosa, Giuseppe Francesco Italiano, and Andrea Ribichini. On resilient graph spanners. In *ESA*, pages 85–96, 2013.
- 4 Surender Baswana and Neelesh Khanna. Approximate shortest paths avoiding a failed vertex: Near optimal data structures for undirected unweighted graphs. *Algorithmica*, 66(1):18–50, 2013. URL: <http://dx.doi.org/10.1007/s00453-012-9621-y>, doi:10.1007/s00453-012-9621-y.
- 5 Surender Baswana, Kavitha Telikepalli, Kurt Mehlhorn, and Seth Pettie. Additive spanners and (α, β) -spanners. *ACM Transactions on Algorithms*, 7(1):5, 2010. URL: <http://dx.doi.org/10.1145/1868237.1868242>, doi:10.1145/1868237.1868242.
- 6 Aaron Bernstein and David R. Karger. A nearly optimal oracle for avoiding failed vertices and edges. In *STOC*, pages 101–110, 2009.
- 7 Davide Bilò, Fabrizio Grandoni, Luciano Gualà, Stefano Leucci, and Guido Proietti. Improved purely additive fault-tolerant spanners. In *ESA*, pages 167–178, 2015.
- 8 Davide Bilò, Luciano Gualà, Stefano Leucci, and Guido Proietti. Fault-tolerant approximate shortest-path trees. In *ESA*, pages 137–148, 2014. URL: http://dx.doi.org/10.1007/978-3-662-44777-2_12, doi:10.1007/978-3-662-44777-2_12.
- 9 Gilad Braunschvig, Shiri Chechik, and David Peleg. Fault tolerant additive spanners. In *WG*, pages 206–214, 2012. URL: http://dx.doi.org/10.1007/978-3-642-34611-8_22, doi:10.1007/978-3-642-34611-8_22.
- 10 Bernard Chazelle. A minimum spanning tree algorithm with inverse-ackermann type complexity. *J. ACM*, 47(6):1028–1047, 2000. URL: <http://doi.acm.org/10.1145/355541.355562>, doi:10.1145/355541.355562.
- 11 Shiri Chechik. New additive spanners. In *SODA*, pages 498–512, 2013. URL: <http://dx.doi.org/10.1137/1.9781611973105.36>, doi:10.1137/1.9781611973105.36.
- 12 Shiri Chechik. Approximate distance oracles with constant query time. In *STOC*, pages 654–663, 2014. URL: <http://doi.acm.org/10.1145/2591796.2591801>, doi:10.1145/2591796.2591801.
- 13 Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. Fault-tolerant spanners for general graphs. In *STOC*, pages 435–444, 2009.
- 14 Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. f -sensitivity distance oracles and routing schemes. In *ESA*, pages 84–96, 2010.
- 15 Annalisa D’Andrea, Mattia D’Emidio, Daniele Frigioni, Stefano Leucci, and Guido Proietti. Path-fault-tolerant approximate shortest-path trees. In *SIROCCO*, pages 224–238, 2015. URL: http://dx.doi.org/10.1007/978-3-319-25258-2_16, doi:10.1007/978-3-319-25258-2_16.
- 16 Michael Dinitz and Robert Krauthgamer. Fault-tolerant spanners: better and simpler. In *PODC*, pages 169–178, 2011.

- 17 Ran Duan and Seth Pettie. Dual-failure distance and connectivity oracles. In *SODA*, pages 506–515, 2009. URL: <http://dl.acm.org/citation.cfm?id=1496770.1496826>.
- 18 Michael Elkin and Seth Pettie. A linear-size logarithmic stretch path-reporting distance oracle for general graphs. In *SODA*, pages 805–821, 2015. URL: <http://dx.doi.org/10.1137/1.9781611973730.55>, doi:10.1137/1.9781611973730.55.
- 19 David Eppstein. Offline algorithms for dynamic minimum spanning tree problems. *J. Algorithms*, 17(2):237–250, 1994. URL: <http://dx.doi.org/10.1006/jagm.1994.1033>, doi:10.1006/jagm.1994.1033.
- 20 David Eppstein, Zvi Galil, Giuseppe F. Italiano, and Amnon Nissenzweig. Sparsification-a technique for speeding up dynamic graph algorithms (extended abstract). In *FOCS*, pages 60–69, 1992. URL: <http://dx.doi.org/10.1109/SFCS.1992.267818>, doi:10.1109/SFCS.1992.267818.
- 21 Paul Erdős. Extremal problems in graph theory. In *Theory of Graphs and its Applications*, pages 29–36, 1964.
- 22 Greg N. Frederickson. Data structures for on-line updating of minimum spanning trees, with applications. *SIAM J. Comput.*, 14(4):781–798, 1985. URL: <http://dx.doi.org/10.1137/0214055>, doi:10.1137/0214055.
- 23 Fabrizio Grandoni and Virginia Vassilevska Williams. Improved distance sensitivity oracles via fast single-source replacement paths. In *FOCS*, pages 748–757, 2012.
- 24 Dov Harel and Robert Endre Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM J. Comput.*, 13(2):338–355, 1984. URL: <http://dx.doi.org/10.1137/0213024>, doi:10.1137/0213024.
- 25 Monika Rauch Henzinger and Valerie King. Maintaining minimum spanning forests in dynamic graphs. *SIAM J. Comput.*, 31(2):364–374, 2001. URL: <http://dx.doi.org/10.1137/S0097539797327209>, doi:10.1137/S0097539797327209.
- 26 Jacob Holm, Kristian de Lichtenberg, and Mikkel Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *J. ACM*, 48(4):723–760, 2001. URL: <http://doi.acm.org/10.1145/502090.502095>, doi:10.1145/502090.502095.
- 27 Enrico Nardelli, Guido Proietti, and Peter Widmayer. Swapping a failing edge of a single source shortest paths tree is good and fast. *Algorithmica*, 35(1):56–74, 2003.
- 28 Merav Parter. Vertex fault tolerant additive spanners. In *DISC*, pages 167–181, 2014. URL: http://dx.doi.org/10.1007/978-3-662-45174-8_12, doi:10.1007/978-3-662-45174-8_12.
- 29 Merav Parter. Dual failure resilient BFS structure. In *PODC*, pages 481–490, 2015.
- 30 Merav Parter and David Peleg. Sparse fault-tolerant BFS trees. In *ESA*, pages 779–790, 2013. URL: http://dx.doi.org/10.1007/978-3-642-40450-4_66, doi:10.1007/978-3-642-40450-4_66.
- 31 Merav Parter and David Peleg. Fault tolerant approximate BFS structures. In *SODA*, pages 1073–1092, 2014.
- 32 Mikkel Thorup and Uri Zwick. Approximate distance oracles. *J. ACM*, 52(1):1–24, 2005. URL: <http://doi.acm.org/10.1145/1044731.1044732>, doi:10.1145/1044731.1044732.