

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

## Information and Computation

[www.elsevier.com/locate/yinco](https://www.elsevier.com/locate/yinco)

# Constant-time dynamic weight approximation for minimum spanning forest

Monika Henzinger<sup>a,1</sup>, Pan Peng<sup>b,\*</sup>

<sup>a</sup> Faculty of Computer Science, University of Vienna, Austria

<sup>b</sup> School of Computer Science and Technology, University of Science and Technology of China, China

## ARTICLE INFO

### Article history:

Received 22 January 2021

Received in revised form 27 August 2021

Accepted 22 September 2021

Available online xxxx

### Keywords:

Dynamic graph algorithms

Minimum spanning forest

Sublinear-time algorithms

Cell-probe lower bounds

## ABSTRACT

We give two fully dynamic algorithms that maintain a  $(1 + \varepsilon)$ -approximation of the weight  $M$  of a minimum spanning forest (MSF) of an  $n$ -node graph  $G$  with edges weights in  $[1, W]$ , for any  $\varepsilon > 0$ . (1) Our *deterministic* algorithm takes  $O(W^2 \log W / \varepsilon^3)$  *worst-case* update time, which is  $O(1)$  if both  $W$  and  $\varepsilon$  are constants. (2) Our randomized (Monte-Carlo style) algorithm works with high probability and runs in *worst-case*  $O(\log W / \varepsilon^4)$  update time if  $W = O((m^*)^{1/6} / \log^{2/3} n)$ , where  $m^*$  is the minimum number of edges in the graph throughout all the updates. It works even against an adaptive adversary.

We complement our algorithmic results with two cell-probe lower bounds for dynamically maintaining an approximation of the weight of an MSF of a graph.

© 2021 Elsevier Inc. All rights reserved.

## 1. Introduction

Minimum spanning forest (MSF) is a fundamental and well studied graph problem in computer science. Given an edge-weighted graph  $G$ , an MSF is a subgraph of  $G$  that forms a spanning forest and has minimum weight among all spanning forests of  $G$ , where the weight of a spanning forest is the sum of the edge weights of the forest. In this paper, we study dynamic algorithms for maintaining and extracting information regarding the MSF of a dynamically changing graph.

A (fully) dynamic graph algorithm is a data structure that provides information about a graph property while the graph is being modified by *edge updates* such as edge insertions or deletions. When designing a dynamic graph algorithm the goal is to minimize the time per update or query operation. Formally, a *fully dynamic* graph algorithm is an algorithm that maintains some information of a graph  $G$  which is undergoing an arbitrary sequence of the following operations: 1) **Insert**( $u, v, w$ ): insert the edge  $(u, v)$  with weight  $w$  in  $G$ ; 2) **Delete**( $u, v$ ): delete the edge  $(u, v)$  from  $G$ . (If the considered graph is unweighted, then the weight  $w$  of an insertion is always set to be 1.) The algorithm can further support different **Query** operations depending on the specific graph property. Dynamically maintaining the (minimum) spanning forest (or tree) (e.g., [1–9]) and dynamically testing connectivity between any pair of vertices (e.g., [2,5,3,10–13]) have played a funda-

\* Corresponding author.

E-mail addresses: [monika.henzinger@univie.ac.at](mailto:monika.henzinger@univie.ac.at) (M. Henzinger), [ppeng@ustc.edu.cn](mailto:ppeng@ustc.edu.cn) (P. Peng).

<sup>1</sup> The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013)/ERC Grant Agreement no. 340506.

<https://doi.org/10.1016/j.ic.2021.104805>

0890-5401/© 2021 Elsevier Inc. All rights reserved.

**Table 1**

The best known dynamic algorithms for the  $(1 + \varepsilon)$ -approximate MSF problem on a graph with edge weights from  $[1, W]$ . Those algorithms maintain a spanning tree whose weight is a  $(1 + \varepsilon)$ -approximation of the weight of an MSF. For the Monte Carlo algorithm, there is some  $1/\text{poly}(n)$  probability of answering a query incorrectly. For the Las Vegas algorithm, the answers to the queries are always correct, while the update time is a random variable.

	worst-case	amortized
deterministic	$n^{o(1)}$ [14]	$\min\{O(\frac{\log W \log^2 n}{\varepsilon \log \log n}), O(\frac{\log^4 n}{\log \log n})\}$ [15,3] [6]
Monte Carlo	$O(\log W \log^4 n/\varepsilon)$ [12,3]	
Las Vegas		expected $O(\log W (\log n)(\log \log n)^2/\varepsilon)$ [13,3]

mental role in the area of dynamic graph algorithms. The currently best dynamic algorithms for MSF take  $O(\log^4 n / \log \log n)$  amortized<sup>2</sup> time per update [6] (that improves upon [5]) and  $O(n^{o(1)})$  worst-case time [9,14].

An approximate version of this problem has also been studied: A  $(1 + \varepsilon)$ -approximation<sup>3</sup> of the weight  $M$  of an MSF is a value  $M'$  such that  $(1 - \varepsilon) \cdot M \leq M' \leq (1 + \varepsilon) \cdot M$ . In the  $(1 + \varepsilon)$ -approximate MSF problem, the fully dynamic algorithm maintains an MSF of weight  $M'$  that is a  $(1 + \varepsilon)$ -approximation of the weight of an MSF in the current graph. This problem was first studied by Henzinger and King [3], who also provided a reduction that uses a dynamic algorithm  $\mathcal{B}$  for a spanning tree to obtain a dynamic algorithm  $\mathcal{A}$  for the  $(1 + \varepsilon)$ -approximate MSF problem, by invoking the algorithm  $\mathcal{B}$  on  $\log W / \varepsilon$  subgraphs of a graph  $G$  with edge weights in  $[1, W]$ , for some  $W \geq 1$ . By combining this reduction and best known algorithms for dynamic connectivity (and spanning forest), we have the following table Table 1 summarizing the state-of-the-art. We remark that all these algorithms require at least logarithmic update time per operation.

It is natural to consider a relaxed version of the above problems, by asking how to dynamically maintain an exact or approximate value, i.e., not the spanning tree, of the weight of an MSF with much faster update time. We first note that maintaining the exact value of the weight of an MSF dynamically cannot be done in constant time per operation: There is a lower bound of Patrascu and Demaine [16], who showed that in the cell-probe model many fundamental graph properties, such as asking whether the graph is connected or maintaining the exact value of the weight of an MSF, require  $\Omega(\log n)$  time per operation, where  $n$  is the number of nodes in the graph. Their lower bound even holds for the unweighted case, i.e. for  $W = 1$ . However, their lower bound does not apply to maintaining an approximate value of the weight of an MSF, leading to the following open question: *Can a  $(1 + \varepsilon)$ -approximation of the weight of an MSF be maintained in time  $o(\log n)$ ?*

We answer this question positively in this paper. Our main contributions are two dynamic algorithms for maintaining a  $(1 + \varepsilon)$ -approximation of the weight of an MSF, that bypass the  $\Omega(\log n)$  time lower bound barrier. Even stronger, they have constant update times for a wide range of parameters. More specifically, our first algorithmic result is as follows.

**Theorem 1.1.** *There exists a fully dynamic, deterministic algorithm that maintains an estimator  $\bar{M}$  that  $(1 + \varepsilon)$ -approximates the weight  $M$  of a MSF of a graph with edge weights from  $[1, W]$ . The worst-case time per update operation is  $O(W^2 \cdot \log W \cdot \varepsilon^{-3})$ .*

For constant  $W$  and  $\varepsilon$  this is a deterministic worst-case  $O(1)$  time bound. Our algorithm does not require the initial graph to be empty. In contrast, for maintaining the exact value of the weight of an MSF, the best known deterministic algorithm has  $O(n^{o(1)})$  worst-case time [14] (which actually maintains the MSF, not just the value). In comparison, the already mentioned lower bound of  $\Omega(\log n)$  by [16] also applies for maintaining the exact weight of an MSF even for  $W = 1$ .

We then show that any deterministic data structure that dynamically maintains the  $(1 + \varepsilon)$ -approximate weight of an MSF requires super constant time per operation, if  $W \geq (\log n)^{\omega_n(1)}$ , where  $\omega_n(1)$  is any function that goes to infinity if  $n$  goes to infinity. The lower bound on the time per operation for this problem is established in the cell-probe model.

**Theorem 1.2.** *Let  $W \geq (\log n)^{\omega_n(1)}$ . Let  $G$  be a dynamic graph with edge weights in  $[1, W]$ . Any deterministic data structure that dynamically maintains the weight  $M$  of an MSF of a graph  $G$  within a multiplicative factor of  $(1 + \varepsilon)$  for any constant  $\varepsilon > 0$ , or an additive error less than  $W/2$  must perform  $\omega_n(1)$  cell probes, where each cell has size  $O(\log n)$ .*

Our second algorithmic result is a randomized dynamic algorithm that works even against an adaptive adversary.<sup>4</sup> We let  $m^*$  be the minimum number of edges of the graph throughout all the updates. In this paper, “With high probability” means “with probability at least  $1 - \frac{1}{n^c}$  for some constant  $c \geq 1$ ”.

<sup>2</sup> An algorithm is said to have amortized update time of  $\alpha$  if for any  $t$ , after  $t$  updates the total update time is at most  $\alpha t$ . It is said to have worst-case update time of  $\alpha$  if every update time is at most  $\alpha$ .

<sup>3</sup> Note that we can easily translate this definition of  $(1 + \varepsilon)$ -approximation of  $M$  to the standard one, i.e., a value  $M''$  such that  $M \leq M'' \leq (1 + \varepsilon)M$ . This can be achieved by multiplying the estimate  $M'$  by  $1/(1 - \varepsilon)$  and replacing  $\varepsilon$  by  $\varepsilon' = \min\{\varepsilon/4, \frac{1}{2}\}$ .

<sup>4</sup> That is, the adversary sees the answers to all query operations before deciding which edge to update next. We say an adversary is *oblivious*, if it fixes the sequence of edge insertions and deletions and is oblivious to the randomness of algorithm.

**Table 2**

Our algorithms for the problem of maintaining a  $(1 + \varepsilon)$ -approximation of the weight of an MSF of a graph with edge weights from  $[1, W]$ . Our algorithms do not maintain a spanning tree. Note that our algorithms are faster than the corresponding algorithms in Table 1 for a wide range of parameters.

	worst-case
deterministic	$O(W^2 \cdot (\log W) \cdot \varepsilon^{-3})$ Theorem 1.1
Monte Carlo	$O\left(\frac{\log W}{\varepsilon} + \frac{W^3(\log W) \log(W/\varepsilon)(\log(\log(W/\varepsilon)) + \log n)}{\varepsilon^4 \sqrt{m^*}}\right)$ Theorem 1.3

**Theorem 1.3.** *Let  $\varepsilon \in (0, 1)$  and  $W \geq 1$ . There exists a fully dynamic algorithm that with high probability, maintains an estimator  $\overline{M}$  that is a  $(1 + \varepsilon)$ -approximation of the weight  $M$  of MSF of a graph  $G$  with edge weights from  $[1, W]$ . The worst-case time per update operation is*

$$O\left((\log W)/\varepsilon + W^3(\log W) \log(W/\varepsilon) (\log(\log(W/\varepsilon)) + \log n)/(\varepsilon^4 \sqrt{m^*})\right),$$

where  $m^* \geq 1$  is the minimum number of edges in  $G$  throughout all the updates. Our algorithm works against an adaptive adversary.

From the above, we can easily obtain a randomized algorithm that runs in sub-logarithmic time (for constant  $\varepsilon > 0$ ) for the setting that  $W$  might not be constant.

**Corollary 1.4.** *Let  $\varepsilon \in (\frac{1}{n^4}, 1)$ . There exists a fully dynamic algorithm that, with high probability, maintains an estimator  $\overline{M}$  that is a  $(1 + \varepsilon)$ -approximation of the weight  $M$  of an MSF of a graph  $G$  with edge weights from  $[1, W]$  such that  $W = O((m^*)^{1/6}/\log^{2/3} n)$ , and has*

- $O((\log W)/\varepsilon^4)$  worst-case time per update operation;
- $O(1)$  worst-case time per update operation, if  $W$  and  $\varepsilon$  are constant;
- $o(\log n)$  worst-case time per update operation, if  $W = 2^{o(\log n)}$  and  $\varepsilon$  is constant.

The algorithm works against an adaptive adversary.

**Proof.** By applying Theorem 1.3 with the assumptions that  $\varepsilon \in (\frac{1}{n^4}, 1)$  and that  $W = O((m^*)^{1/6}/\log^{2/3} n)$ , we obtain a corresponding algorithm with worst-case update time

$$\begin{aligned} & O\left(\frac{\log W}{\varepsilon} + \frac{W^3(\log W) \log(W/\varepsilon)(\log(\log(W/\varepsilon)) + \log n)}{\varepsilon^4 \cdot \sqrt{m^*}}\right) \\ &= O\left(\frac{\log W}{\varepsilon} + \frac{\sqrt{m^*}(\log W) \log(W/\varepsilon)(\log n)}{\varepsilon^4 \cdot \sqrt{m^*} \cdot \log^2 n}\right) = O\left(\frac{\log W}{\varepsilon^4}\right), \end{aligned}$$

where we used the fact that  $m^* \leq O(n^2)$ . This finishes the proof of the corollary.  $\square$

The requirement that  $W$  needs to be small compared to the minimum number of edges in the graph might look artificial. However, we present the following lower bound for any randomized data structure that dynamically maintains the approximate weight of an MSF, which shows that some requirement of this type is necessary.

**Theorem 1.5.** *For any constant  $\varepsilon, \alpha > 0$  and  $W = n^\alpha$ , any data structure that dynamically maintains with high probability the weight of an MSF of a graph  $G$  with edge weights in  $[1, W]$  and  $W = \Omega(\varepsilon m^*)$  within a multiplicative factor of  $(1 + \varepsilon)$ , or an additive error less than  $W/2$ , must perform  $\Omega(\log n)$  cell probes, where each cell has size  $O(\log n)$ .*

This stands in interesting contrast to the above Corollary 1.4, which has update time  $o(\log n)$ , but requires  $W$  to be  $\min\{O((m^*)^{1/6}/\log^{2/3} n), 2^{o(\log n)}\}$ .

We remark that main technical contributions are our algorithmic results given by Theorem 1.1 and 1.3, which are summarized in Table 2, for ease of comparison.

### 1.1. Our techniques

Both our deterministic and randomized dynamic graph algorithms for this problem use an approach developed in the area of property testing: Build an efficient algorithm for estimating the number of connected components (CCs) in a graph

and apply it to suitable subgraphs of  $G$  [17–19]. More specifically, we build constant-time dynamic algorithms that estimate the number of CCs with appropriate additive error, apply them to  $O(\log W/\varepsilon)$  many subgraphs, and then use an extension of the formula in [17] to disconnected graphs to estimate the weight of an MSF. We stress that most of previous work uses the formula for *connected* graphs, while in our setting, it is more natural to remove the requirement that the dynamic graph is always connected. Thus, we extend the formula to work for disconnected graphs (see Lemma 2.1). To use this formula our algorithms need to estimate the number of CCs with an additive error  $\varepsilon' \cdot \text{nis}(G)$ , where  $\text{nis}(G)$  is the number of *non-isolated* vertices in  $G$  (see below why this is crucial). Our randomized dynamic algorithm for this problem achieves such an error in time  $O(\max\{1, \log(1/\varepsilon') \log n / (\varepsilon'^3 \cdot \sqrt{m^*})\})$  with high probability, and our deterministic algorithm achieves the same error  $\varepsilon' \cdot \text{nis}(G)$  in time  $O((1/\varepsilon')^2)$ .

**The randomized algorithm.** The *randomized* algorithm for estimating the number of connected components is based on the following simple approach (used previously e.g. [20]): *Whenever (1) there exists a static algorithm that in time  $T$  estimates a desired parameter (here the number of CCs) with an additive error of  $\text{ERR}$  and (2) each update operation changes the value of a desired parameter only by an additive value up to  $+/-\delta$  (here 1), then running the static algorithm every  $\frac{\text{ERR}}{\delta}$  update operations leads to a dynamic algorithm with additive error of at most  $2\text{ERR}$  and amortized time  $O(\frac{T\delta}{\text{ERR}})$  per update.*

To turn this into a *worst-case* time bound observe that we can run the static algorithm “in the background” executing  $O(\frac{T\delta}{\text{ERR}})$  steps of the static algorithm at every update operation. This increases the additive error only by a constant factor. We make use of the static (constant-time) algorithm of [21] (that improves upon [17]) for estimating the number of CCs with additive error  $\varepsilon'n$  as a subroutine. However, directly invoking this static algorithm we can obtain only a dynamic estimator for the number of CCs with an additive error  $\varepsilon'n^{2/3} \log^{2/3} n$  with  $O(1/\varepsilon'^3)$  update time (see Appendix A). This is not sufficient for the MSF approximation algorithm: Specifically, we need to be able to dynamically approximate the number of CCs with an additive error  $\varepsilon' \cdot \text{nis}(G)$  with fast update time and  $\text{nis}(G)$  can be arbitrarily small. To achieve this smaller additive error, we carefully choose different values of  $\text{ERR}$  throughout all the updates and sample the non-isolated vertices uniformly at random in the dynamic graph. We give more detailed discussions on some other technical difficulties (e.g., we need to “synchronize” the updates in  $G$  and its subgraphs to achieve the desired update time) and how we handle them in Section 4.

Sampling non-isolated vertices uniformly at random in a dynamic graph is exactly the problem solved by  $\ell_0$ -sampling in streaming algorithms (see e.g. [22]). However, all such algorithms, while only using  $O(\text{poly} \log n)$  space, require time  $\Omega(\log n)$  per operation. We give a relatively simple data structure (see Lemma 4.2) that allows to sample all non-zero entries in a dynamically changing vector of size  $n$  in constant time (no matter how small their number might be), albeit with space  $O(n)$ . We believe that our data structure might be of independent interest.

Our randomized algorithm uses “fresh” random bits whenever it invokes the static (constant-time) algorithm on the current graph, and does not reuse any information computed before. Furthermore, our analysis assumes that the adversary changes the graph in the worst possible way, i.e., changes the number of CCs by at most 1 in each update. Thus, our algorithm works against an adaptive adversary, i.e. an adversary that sees the answers to all queries *before* deciding on the next update operation.

**The deterministic algorithm.** To design a *deterministic worst-case* dynamic algorithm we cannot simply invoke the static constant-time algorithm: that algorithm is inherently randomized as it is designed with the goal of reading the smallest possible portion of the graph. Instead we carefully implement the random local exploration that underlies the static randomized algorithm in a deterministic way. Our key observations are that (1) *we only need to count the number of CCs which are small in size, i.e., which consist of up to  $1/\varepsilon'$  vertices, as the number of larger CCs is at most  $\varepsilon' \cdot \text{nis}(G)$*  and (2) *these counts can be maintained in worst-case time  $O(1/\varepsilon'^2)$  after each update by exploring a neighborhood of  $O(1/\varepsilon')$  vertices “near” the endpoints of the updated edge.*

Both the randomized and the deterministic MSF algorithm run their respective CCs estimation algorithms on each of the  $O(\log W/\varepsilon)$  relevant subgraphs with  $\varepsilon' = \Theta(\varepsilon/W)$ . Using the above-mentioned formula we obtain an estimate  $\bar{M}$  that approximates the weight  $M$  of MSF with an additive error of  $\varepsilon \text{nis}(G)/4$ . As the weight of any MSF is at least  $\text{nis}(G)/2$ , this additive error is at most  $\varepsilon M/2$ , i.e., a  $(1 + \varepsilon)$ -approximation of  $M$ . Whenever a query on  $M$  is asked, we return  $\bar{M}$  in constant time. For our deterministic algorithm for MSF, the time per edge update is  $O(1/\varepsilon'^2) = O(W^2/\varepsilon^2)$  for each of the  $O(\log W/\varepsilon)$  subgraphs, resulting in a worst-case  $O(W^2 \log W/\varepsilon^3)$  update time. The running time of our randomized algorithm for MSF can be analyzed analogously.

We remark that our algorithms are much simpler than the fastest dynamic (exact or approximate) MSF algorithms: the algorithms of [6,5,12,13,23] maintain a hierarchical decomposition with  $O(\log n)$  levels (in the approximate setting this is even  $O(\log n \log W)$  levels), the algorithm of [9] maintains a decomposition of the graph into expanders and a “remaining” part.

**The lower bounds.** Both the lower bounds given in Theorem 1.2 and Theorem 1.5 are proven via simple reductions from the previous cell probe lower bounds for dynamic connectivity. Let us take the proof of Theorem 1.2 for example. We give a reduction, similar to the one in [24], from a problem called the *parity prefix sum problem* to the problem of dynamically maintaining the weight of an MSF of a graph with large maximum edge weights within a multiplicative factor of  $1 + \varepsilon$ . That is, using our reduction, one can use an algorithm for the latter problem to solve the former problem, for which a cell-probe lower bound is long known [25]. This gives a corresponding lower bound for the problem of dynamically maintaining a

$(1 + \varepsilon)$ -approximation of the weight of an MSF of a graph. The proof of Theorem 1.5 is similar, except that we use a variant of reduction from the problem of dynamic connectivity, whose cell-probe lower bound is given in [16].

## 1.2. Other related work

There are few other graph problems for which constant-time dynamic algorithms are known: (a) maximal matching (randomized) [26], (b)  $(2 + \varepsilon)$ -approximate vertex cover (deterministic) [27], (c)  $(2k - 1)$ -stretch spanner of size  $O(n^{1+\frac{1}{k}} \log^2 n)$  for constant  $k$  (randomized) [28], and (d)  $(\Delta + 1)$ -vertex coloring [29,30]. All these are *amortized* time bounds and almost all these algorithms except the coloring algorithm in [29] maintain a sophisticated hierarchical graph decomposition, which makes them rather impractical.

Fully dynamic algorithms for minimum spanning forest and connectivity testing have also been investigated for some special classes of graphs (e.g., [31,32]). There are also works on approximating the weight of minimum spanning tree (or forest) in the query model [17,18] in sublinear time and data streaming model [19,33,34] in sublinear space.

## 2. Preliminaries

Let  $G$  be a (static) undirected graph without parallel edges and with edge weights in  $[1, W]$ . Let  $M$  be the weight of an MSF of  $G$ . We use  $\text{ncc}(G)$  to denote the number of CCs of  $G$ ,  $\text{nis}(G)$  to denote the number of *non-isolated* vertices of  $G$ , and  $\text{size}$  of a CC to denote the number of vertices in the CC.

Our algorithms exploit a relation between the weight of an MSF of a graph  $G$  and the number of CCs of some subgraphs of  $G$ .

For any  $\varepsilon > 0$ , we let  $r = \lceil \log_{1+\varepsilon} W \rceil$ ,  $\ell_i = (1 + \varepsilon)^i$ , and  $\lambda_i = (1 + \varepsilon)^{i+1} - (1 + \varepsilon)^i$ . For any  $i \geq 0$ , we let  $G^{(i)}$  denote the subgraph of  $G$  spanned by all edges with weights at most  $\ell_i$ , and let  $c^{(i)}$  denote the number of CCs in  $G^{(i)}$ . We will make use of the following lemma, whose proof is almost the same as in [19] (see also [17,18]), except that we are considering a graph that is not necessarily connected. We present the proof here for the sake of completeness.

**Lemma 2.1** ([17–19]). *Let  $\varepsilon \in (0, 1)$ . Let  $G$  be a weighted graph and let  $M$ ,  $W$ , and  $c^{(i)}$  be defined as above. Then*

$$M \leq n - c^{(r)} \cdot (1 + \varepsilon)^r + \sum_{i=0}^{r-1} \lambda_i \cdot c^{(i)} \leq (1 + \varepsilon)M. \quad (1)$$

**Proof.** We let  $G'$  be the graph that is obtained by rounding each edge weight up to the nearest power of  $(1 + \varepsilon)$ . Note that the weight  $M'$  of an MSF of  $G'$  satisfies that  $M \leq M' \leq (1 + \varepsilon)M$ .

Now we invoke the Kruskal's algorithm on  $G'$  to compute  $M'$ . The algorithm will add  $n - c^{(0)}$  edges of weight  $\ell_0 = 1$ ,  $c^{(0)} - c^{(1)}$  edges of weight  $\ell_1$ , and so on. Thus

$$M' = n - c^{(0)} + \sum_{i=0}^{r-1} \ell_{i+1} (c^{(i)} - c^{(i+1)}) = n - c^{(r)} \ell_r + \sum_{i=0}^{r-1} (\ell_{i+1} - \ell_i) \cdot c^{(i)}.$$

This proves the statement of the lemma.  $\square$

## 3. A deterministic dynamic algorithm

In this section, we present our deterministic dynamic algorithm for maintaining the weight  $M$  of a minimum spanning forest of a graph  $G$  without parallel edges and with edge weights in  $[1, W]$ , and give the proof of Theorem 1.1.

We first give a deterministic dynamic algorithm for estimating the number of connected components (CCs) with appropriate additive error.

**Theorem 3.1.** *Let  $\varepsilon > 0$ . There exists a fully dynamic and deterministic algorithm that preprocesses a potentially non-empty graph in  $O(\frac{n}{\varepsilon})$  time, and maintains an estimator  $\bar{c}$  s.t.,  $|\bar{c} - \text{ncc}(G)| \leq \varepsilon \cdot \text{nis}(G)$  with worst-case  $O(1/\varepsilon^2)$  update time per operation.*

We remark that in the above theorem, the initial graph can be an arbitrary graph, and the performance guarantee holds even if the algorithm is not aware of the value  $\text{nis}(G)$ . The proof of the above theorem is deferred to the end of this section. By combining Theorem 3.1 and the relation in Lemma 2.1, we are ready to give the algorithm DETERESTWMSF for maintaining an estimator of the weight of an MST.

DETERESTWMSF( $G, \varepsilon, W$ )  $\triangleright$  **Dynamically Maintaining an estimator  $\overline{M}$  of the weight of an MSF**

1. Let  $r = \lceil \log_{1+\varepsilon/2} W \rceil$ ,  $\ell_i = (1 + \varepsilon/2)^i$ , and  $\lambda_i = (1 + \varepsilon/2)^{i+1} - (1 + \varepsilon/2)^i$ .
2. Let  $G^{(i)}$  denote the subgraph of the current graph  $G$  spanned by all edges with weights at most  $\ell_i$ . Let  $c^{(i)}$  be the number of CCs of  $G^{(i)}$ .
3. For each  $1 \leq i \leq r$ , invoke the dynamic algorithm from Theorem 3.1 to maintain an estimator  $\overline{c}_i$  for  $c^{(i)}$  with  $\varepsilon' = \varepsilon/(12W)$ .
4. Define  $\overline{M} := n - \overline{c}_r \cdot (1 + \varepsilon/2)^r + \sum_{i=0}^{r-1} \lambda_i \cdot \overline{c}_i$ .

Now we are ready to prove Theorem 1.1, which is restated in the following for the sake of readability.

**Theorem 1.1.** *There exists a fully dynamic, deterministic algorithm that maintains an estimator  $\overline{M}$  that  $(1 + \varepsilon)$ -approximates the weight  $M$  of a MSF of a graph with edge weights from  $[1, W]$ . The worst-case time per update operation is  $O(W^2 \cdot \log W \cdot \varepsilon^{-3})$ .*

**Proof.** Let  $r, \ell_i, \lambda_i, G^{(i)}, c^{(i)}, \overline{M}$  be defined as in the above algorithm DETERESTWMSF( $G, \varepsilon, W$ ). Note that  $\text{nis}(G) \geq \text{nis}(G^{(i)})$ , since  $G^{(i)}$  is a subgraph of  $G$  for any  $1 \leq i \leq r$ .

Since  $\overline{c}_i$  is the estimator of  $c^{(i)}$  obtained by invoking the dynamic algorithm in Theorem 3.1 with approximation parameter  $\varepsilon' = \varepsilon/(12W)$ , we have that for each  $1 \leq i \leq r$ ,

$$|\overline{c}_i - c^{(i)}| \leq \varepsilon \cdot \text{nis}(G^{(i)})/(12W) \leq \varepsilon \cdot \text{nis}(G)/(12W)$$

Let  $X = n - c^{(r)} \cdot (1 + \varepsilon/2)^r + \sum_{i=0}^{r-1} \lambda_i \cdot c^{(i)}$ . Recall that  $M$  is the weight of an MSF of the current graph  $G$ . We have

$$\begin{aligned} |\overline{M} - X| &\leq (1 + \varepsilon/2)^r |\overline{c}_r - c^{(r)}| + \sum_{i=0}^{r-1} \lambda_i |\overline{c}_i - c^{(i)}| \\ &\leq (1 + \varepsilon/2)^r \cdot (\varepsilon \cdot \text{nis}(G)/(12W)) + (\varepsilon/2) \sum_{i=0}^{r-1} (1 + \varepsilon/2)^i \cdot (\varepsilon \cdot \text{nis}(G)/(12W)) \\ &< (1 + \varepsilon/2)^r (\varepsilon \cdot \text{nis}(G)/(12W)) + (1 + \varepsilon/2)^r (\varepsilon \cdot \text{nis}(G)/(12W)) \\ &\leq 2(1 + \varepsilon/2)^{1+\log_{1+\varepsilon/2} W} \cdot (\varepsilon \cdot \text{nis}(G)/(12W)) \\ &\leq 2 \cdot (3/2) \cdot (\varepsilon \cdot \text{nis}(G)/12) \leq \varepsilon \cdot \text{nis}(G)/4 \leq \varepsilon \cdot M/2, \end{aligned}$$

where the last inequality follows from the fact that  $M \geq \text{nis}(G)/2$ , as each non-isolated vertex is incident to at least one edge (of weight at least 1) of any MSF in the simple graph  $G$ .

By Lemma 2.1 with approximation parameter  $\varepsilon/2$ , we have that  $M \leq X \leq (1 + \varepsilon/2)M$ . Thus,  $\overline{M}$  is a  $(1 + \varepsilon)$ -approximation of  $M$ .

Now note that for each  $1 \leq i \leq r$ , for estimating  $c^{(i)}$  we invoke the dynamic algorithm in Theorem 3.1 with  $\varepsilon' = \varepsilon/(12W)$ , which has worst-case time  $O(1/\varepsilon'^2) = O(W^2/\varepsilon^2)$  per update operation. Therefore, the worst-case time per update operation of DETERESTWMSF( $G, \varepsilon, W$ ) is  $\sum_{i=1}^r O(\frac{W^2}{\varepsilon^2}) = O(\frac{r}{\varepsilon^2} W^2) = O(\frac{W^2 \cdot \log W}{\varepsilon^3})$ .  $\square$

### 3.1. Proof of Theorem 3.1

In the following, we present the proof of Theorem 3.1, which we restate below.

**Theorem 3.1.** *Let  $\varepsilon > 0$ . There exists a fully dynamic and deterministic algorithm that preprocesses a potentially non-empty graph in  $O(\frac{n}{\varepsilon})$  time, and maintains an estimator  $\overline{c}$  s.t.,  $|\overline{c} - \text{ncc}(G)| \leq \varepsilon \cdot \text{nis}(G)$  with worst-case  $O(1/\varepsilon^2)$  update time per operation.*

**Proof.** For a graph  $G$ , we let  $\text{nsc}(G)$  denote the number of CCs of size at most  $1/\varepsilon$  in  $G$ . We first observe that to approximate  $\text{ncc}(G)$  of a dynamic graph  $G$  with an additive error  $\varepsilon \cdot \text{nis}(G)$ , it suffice to compute and maintain  $\text{nsc}(G)$ . This is true as the total number of CCs of size larger than  $\frac{1}{\varepsilon}$  is at most  $\varepsilon \cdot \text{nis}(G)$ , where  $\text{nis}(G)$  is the number of non-isolated vertices of  $G$ . In the following, we first give a static algorithm for computing  $\text{nsc}(G_0)$ , for any initial graph  $G_0$ , and then we show how to maintain  $\text{nsc}(G)$  dynamically.

*Preprocessing.* We first give a static algorithm STATICNCC for computing the number of small CCs of any initial graph  $G_0$ . We maintain a set of  $1/\varepsilon$  counters  $\text{cnt}_\ell$ , where  $\text{cnt}_\ell$  denotes the number of CCs of size  $\ell$ . Initially, all the counters are set to 0 and all vertices are marked *unvisited*. We recursively choose an arbitrary *unvisited* vertex  $v$ , mark it as *visited* and start a



BFS at  $v$  which runs until (1) it has reached (e.g. discovered an edge to)  $1/\varepsilon + 1$  unvisited vertices, (2) it reaches a visited vertex, or (3) the BFS terminates because whole connected component (of size at most  $1/\varepsilon$ ) containing  $v$  has been explored. Then we mark all the newly discovered vertices as *visited* and update the counters accordingly. More precisely, the static and dynamic algorithms are given in the following two tables.

STATICNCC( $G_0, \varepsilon$ ) ▷ **A static algorithm for the number of CCs of size at most  $1/\varepsilon$**

1. Initialize  $\text{cnt}_\ell = 0$ , for each  $1 \leq \ell \leq 1/\varepsilon$ . Mark all vertices as unvisited.
2. While there exists some unvisited vertex  $v$ :
  - (a) Do BFS from  $v$  until (i)  $1/\varepsilon + 1$  unvisited vertices have been reached, or (ii) any visited vertex has been reached, or (iii) no more new vertices can be reached. Mark all the newly discovered vertices in the search as visited.
  - (b) If (iii) occurs, and  $\ell$  vertices have been reached for some  $\ell \leq 1/\varepsilon$ , then increment  $\text{cnt}_\ell$  by 1.
3. Define the estimator  $\bar{c} := \sum_{\ell=1}^{1/\varepsilon} \text{cnt}_\ell$ .

Note that by definition of the algorithm and the estimator  $\bar{c}$ , it holds that  $\bar{c} = \text{nsc}(G_0)$ .

**Lemma 3.2.** *The algorithm STATICNCC( $G, \varepsilon$ ) can be implemented in  $O(n \cdot \frac{1}{\varepsilon})$  time for any  $n$ -vertex graph  $G$ .*

**Proof.** Note that it suffices to bound the time of *exploring* each CC  $C$ , i.e., until all the vertices inside  $C$  have been marked as visited. Note that  $\text{cnt}_\ell$  is exactly the number of CCs of size  $\ell$ , for  $\ell \leq 1/\varepsilon$  and consider two cases, which together show the  $O(n/\varepsilon)$  bound.

1. If  $|C| = \ell \leq \frac{1}{\varepsilon}$ , then the total time for exploring  $C$  is  $O(\ell^2)$ . In this case, we note that the total time for exploring CCs of size at most  $1/\varepsilon$  is  $\sum_{\ell=1}^{1/\varepsilon} \text{cnt}_\ell \cdot O(\ell^2) \leq \sum_{\ell=1}^{1/\varepsilon} \text{cnt}_\ell \cdot \ell \cdot O(1/\varepsilon) = O(n/\varepsilon)$ , where the last equation follows from the fact that  $\sum_{\ell=1}^{1/\varepsilon} \text{cnt}_\ell \cdot \ell \leq n$ .
2. If  $|C| > 1/\varepsilon$ , let  $S = \{v_1, v_2, \dots, v_b\}$  denote the set of vertices from which we start a BFS in  $C$ . Note that for each  $i \leq b$ , the number of newly discovered vertices from vertex  $v_i$  is at most  $1/\varepsilon + 1$  by the description of our algorithm. Let  $t_j$  denote the number of vertices in  $S$  from which the BFS discovers exactly  $j$  new vertices, for each  $j \leq 1/\varepsilon + 1$ . Then  $|C| = \sum_{j=1}^{1/\varepsilon+1} t_j \cdot j$ . Furthermore, we note that for each  $j \geq 1$ , it takes time  $O(j \cdot \frac{1}{\varepsilon})$  for the BFS to discover exactly  $j$  new vertices, as we will only scan at most  $\frac{1}{\varepsilon} + 1$  neighbors for each of these new vertices. Thus, the total time of exploring  $C$  is  $\sum_{j=1}^{1/\varepsilon+1} t_j \cdot O(j \cdot \frac{1}{\varepsilon}) \leq O(1/\varepsilon) \cdot \sum_{j=1}^{1/\varepsilon+1} t_j \cdot j = O(|C|/\varepsilon)$ . Thus, the total time of exploring CCs of size at least  $1/\varepsilon + 1$  is  $\sum_{C:|C| \geq 1/\varepsilon+1} O(|C|/\varepsilon) = O(n/\varepsilon)$ , where the last equation follows from the fact that  $\sum_{C:|C| \geq 1/\varepsilon+1} |C| \leq n$ . □

*Handling edge updates.* Now we give the details of our dynamic algorithm DETERDYNAMICNCC( $G, \varepsilon$ ) for updating the counter  $\bar{c}$  by running a limited BFS from the two endpoints of the updated edge in the graph before and after the update.

**Correctness.** As aforementioned, it suffices to show that throughout all the updates,  $\bar{c} = \text{nsc}(G)$ , where  $G$  is the current graph. We give details as follows.

For any edge insertion  $(u, v)$ , we know that the number  $\text{nsc}$  (of CCs of size at most  $1/\varepsilon$ ) can change by at most 2. More precisely, it changes if and only if at least one of  $s_u^{(0)}, s_v^{(0)}$  is at most  $1/\varepsilon$  and  $u, v$  do not belong to the same CC before the edge insertion, where  $s_u^{(0)}$  and  $s_v^{(0)}$  are the sizes of the explored subgraphs (before the edge insertion) starting from  $u$  and  $v$ , respectively, that we compute in the algorithm. Furthermore, if Step 2a happens, i.e., exactly one of  $s_u^{(0)}$  and  $s_v^{(0)}$ , say  $s_u^{(0)}$ , is no larger than  $1/\varepsilon$ , then a small CC merges into a large CC, and thus  $\text{nsc}$  decreases by 1. If Step 2b happens (i.e.,  $s_u^{(0)}, s_v^{(0)}$  are smaller than  $1/\varepsilon$ ): if Step 2(b)i happens, i.e.,  $s_u^{(1)}$  is larger than  $1/\varepsilon$ , then two small CCs merge into a CC of size larger than  $1/\varepsilon$  and thus  $\text{nsc}$  decreases by 2; if Step 2(b)ii happens, i.e.,  $s_u^{(1)}$  is no larger than  $1/\varepsilon$  and  $s_u^{(1)} = s_u^{(0)}$ , then  $u, v$  belong to the same CC before  $(u, v)$  was inserted and thus  $\text{nsc}$  remains unchanged; if Step 2(b)iii happens, i.e.,  $s_u^{(1)}$  is no larger than  $1/\varepsilon$  and  $s_u^{(1)} \neq s_u^{(0)}$ , then two small CCs merge into a CC of size no larger than  $1/\varepsilon$  and thus  $\text{nsc}$  decreases by 1. By the description of our algorithm, after the insertion  $(u, v)$ , the maintained  $\bar{c}$  still satisfies that  $\bar{c} = \text{nsc}(G')$ , where  $G'$  is the updated graph.

The case for edge deletions can be analyzed similarly.

**Running time.** Now we analyze the running time of our dynamic algorithm. We note that for each update (either insertion or deletion), we only need to execute  $O(1)$  BFS calls, each of which will explore at most  $O(1/\varepsilon)$  vertices (and thus  $O(1/\varepsilon^2)$  edges). Thus, the worst-case update time is  $O(1/\varepsilon^2)$ .

Theorem 3.1 now follows from the above analysis of the static algorithm STATICNCC and the dynamic algorithm DETERDYNAMICNCC. □

DETERDYNAMICNCC( $G, \varepsilon$ )  $\triangleright$  **Maintaining an estimator for  $\text{ncc}(G)$  of a dynamic graph  $G$**

1. Preprocessing: run the algorithm STATICNCC( $G_0, \varepsilon$ ) to find the  $\bar{c}$ , the number of CCs of  $G_0$  of size at most  $1/\varepsilon$ .
2. Handling an edge insertion  $(u, v)$ : perform three BFS calls: two from  $u$  and  $v$ , respectively, in the graph before the insertion of  $(u, v)$ , and one from  $u$  in the graph after the insertion. Stop the BFS once  $1/\varepsilon + 1$  vertices have been reached or no more new vertices can be reached. Let  $s_u^{(0)}, s_v^{(0)}, s_u^{(1)}$  denote the sizes of the corresponding explored subgraphs.
  - (a) If exactly one of  $s_u^{(0)}$  and  $s_v^{(0)}$ , say  $s_u^{(0)}$ , is no larger than  $1/\varepsilon$ , then decrement  $\bar{c}$  by 1.
  - (b) If both of  $s_u^{(0)}, s_v^{(0)}$  are smaller than  $1/\varepsilon$ :
    - i. if  $s_u^{(1)}$  is larger than  $1/\varepsilon$ , then decrement  $\bar{c}$  by 2;
    - ii. if  $s_u^{(1)}$  is no larger than  $1/\varepsilon$  and  $s_u^{(1)} = s_u^{(0)}$ , then keep  $\bar{c}$  unchanged;
    - iii. if  $s_u^{(1)}$  is no larger than  $1/\varepsilon$  and  $s_u^{(1)} \neq s_u^{(0)}$ , then decrement  $\bar{c}$  by 1.
3. Handling an edge deletion  $(u, v)$ : perform three BFS calls: one from  $u$  in the graph before the deletion of  $(u, v)$ , and two from  $u$  and  $v$ , respectively, in the graph after the deletion. Stop the BFS once  $1/\varepsilon + 1$  vertices have been reached or no more new vertices can be reached. Let  $s_u^{(0)}, s_u^{(1)}, s_v^{(1)}$  denote the sizes of the corresponding explored subgraphs.
  - (a) If exactly one of  $s_u^{(1)}$  and  $s_v^{(1)}$ , say  $s_u^{(1)}$ , is no larger than  $1/\varepsilon$ , then increment  $\bar{c}$  by 1.
  - (b) If both of  $s_u^{(1)}, s_v^{(1)}$  are smaller than  $1/\varepsilon$ :
    - i. if  $s_u^{(0)}$  is larger than  $1/\varepsilon$ , then increment  $\bar{c}$  by 2;
    - ii. if  $s_u^{(0)}$  is no larger than  $1/\varepsilon$  and  $s_u^{(0)} = s_u^{(1)}$ , then keep  $\bar{c}$  unchanged;
    - iii. if  $s_u^{(0)}$  is no larger than  $1/\varepsilon$  and  $s_u^{(0)} \neq s_u^{(1)}$ , then increment  $\bar{c}$  by 1.

#### 4. A randomized dynamic algorithm

Now we present our randomized dynamic algorithm  $\mathcal{A}$  for estimating the weight of an MSF of  $G$  and prove Theorem 1.3. To do so, as for the previous deterministic algorithm, we give dynamic algorithms  $\mathcal{B}_j$  that approximate the number of CCs of subgraphs  $G^{(j)}$  (as defined before) of  $G$  within an additive error  $O(\varepsilon \text{nis}(G)/W)$  with small worst-case update time.

We first consider a dynamic algorithm whose input is a dynamic graph  $H$  and an additional (integer) parameter  $T$  that is also changing along with each edge update. That is, starting with an initial graph and parameter, at each timestamp  $i \geq 1$ , the graph  $H$  is updated by one edge insertion or deletion, and the parameter  $T$  is updated to a new value. We give such a dynamic algorithm for estimating the number of CCs of  $H$  with an additive error parametrized by  $T$ . We will show the following theorem.

**Theorem 4.1.** *Let  $\varepsilon', p \in (0, 1)$ . Consider a dynamic algorithm that takes as input a dynamic graph  $H$  and an additional (integer) parameter  $T$ , called  $T$ -parameter, that is also changing along with each edge update. Suppose that (a)  $T \geq \text{nis}(H)$  throughout all the updates and (b) the  $T$ -parameters at timestamps  $i$  and  $i - 1$  differ by at most 2, for  $i \geq 1$ . Then there exists a fully dynamic algorithm RANDDYNAMICNCC( $H, T, \varepsilon', p$ ) that with probability at least  $1 - p$ , maintains an estimator  $\bar{c}$  for the number  $\text{ncc}$  of CCs of a dynamic graph  $H$  such that  $|\bar{c} - \text{ncc}(H)| \leq \varepsilon' \cdot T$ . The worst-case time per update operation is  $O(\max\{1, \log(1/\varepsilon') \log(1/p)/(e'^3 T^*)\})$ , where  $T^*$  is the minimum of the  $T$ -parameters over all updates. Our algorithm works against an adaptive adversary.*

We defer the proof of Theorem 4.1 to the end of this section. To approximate the weight of an MSF of a dynamic graph  $G$ , our algorithm  $\mathcal{B}_j$  for approximating the number of CCs of  $G^{(j)}$  is obtained by invoking the algorithm RANDDYNAMICNCC from Theorem 4.1 on  $G^{(j)}$  with the  $T$ -parameter being specified to be the number  $\text{nis}(G)$  of non-isolated vertices in  $G$ . Note that by our choice of  $T$ -parameter and choosing  $p = 1/\text{poly}(n)$ , the algorithm  $\mathcal{B}_j$  approximates the number of CCs of each subgraph  $G^{(j)}$  within an additive  $\varepsilon' \text{nis}(G)$ , with worst-case  $O(\max\{1, \log(1/\varepsilon') \log(n)/(e'^3 \text{nis}^*)\})$  time per update operation, with high probability, where  $\text{nis}^*$  is the minimum number of non-isolated vertices in  $G$  throughout all the updates. The running time can be further guaranteed to be  $O(\max\{1, \log(1/\varepsilon') \log(n)/(e'^3 \sqrt{m^*})\})$  as it holds that  $\text{nis}^* > \sqrt{2m^*}$  (a fact that is proven in the proof of Theorem 1.3).

By combining the dynamic algorithm for the number of CCs given in Theorem 4.1 and the relation of the weight of MSF of a graph  $G$  and the number of CCs in its subgraphs as given in Lemma 2.1, we are ready to present our randomized algorithm RANDESTWMSF for the weight of an MST of  $G$ . For each  $1 \leq j \leq r = \lceil \log_{1+\varepsilon/2} W \rceil$ , it first initializes a data structure  $\mathcal{B}_j$  for maintaining an estimator  $\bar{c}_j$  of the number of CCs of each  $G^{(j)}$  using the initial graph  $G_0$ ,  $\text{ncc}(G_0^{(j)})$ ,  $T = \text{nis}(G_0)$ , and corresponding parameters  $\varepsilon' = \varepsilon/(24W)$ ,  $p = p'/r$ . Then for each update in  $G$ , the algorithm first updates each  $G^{(j)}$



accordingly and then invoke  $\mathcal{B}_j$  with  $T$ -parameter  $\text{nis}(G)$  to update  $\bar{c}_j$ . To guarantee that the  $T$ -parameter for  $G^{(j)}$  fulfills the precondition of Theorem 4.1, we introduce the use of self-loops to “synchronize” the updates in  $G$  and  $G^{(j)}$ , for each  $j \leq r$ . That is, for each update on the graph  $G$ , it is guaranteed that there will be one or two corresponding updates on each  $G^{(j)}$ , for each  $j \leq r$  (see Footnote 5 for more explanations). The algorithm  $\text{RANDESTWMSF}$  is formally described as follows.

$\text{RANDESTWMSF}(G, \varepsilon, p', W) \triangleright$  **Dynamically Maintaining an estimator  $\bar{M}$  of the weight of an MSF**

1. Let  $r = \lceil \log_{1+\varepsilon/2} W \rceil$ ,  $\ell_j = (1 + \varepsilon/2)^j$ , and  $\lambda_j = (1 + \varepsilon/2)^{j+1} - (1 + \varepsilon/2)^j$ .
2. Let  $G^{(j)}$  denote the subgraph of the current graph  $G$  spanned by all edges with weights at most  $\ell_j$ . Let  $c^{(j)}$  be the number of CCs of  $G^{(j)}$ .
3. For each  $1 \leq j \leq r$ , initialize an estimator  $\bar{c}_j$  for  $c^{(j)}$  to be  $\bar{c}_j = \text{ncc}(G_0^{(j)})$ , where  $G_0$  is the initial graph of  $G$ ; initialize a data structure  $\mathcal{B}_j$  for  $\text{RANDDYNAMICNCC}(H, T, \varepsilon', p)$  from Theorem 4.1 with  $p = p'/r$ ,  $\varepsilon' = \varepsilon/(24W)$ ,  $H = G_0^{(j)}$  and  $T = \text{nis}(G_0)$ .
4. For each  $i \geq 1$  and the  $i$ -th update in  $G$ :
  - for each  $j$  with  $1 \leq j \leq r$ :
    - (a) if the  $i$ -th update in  $G$  leads to an update in  $G^{(j)}$ , perform the same update to  $H$  in  $\mathcal{B}_j$ , and set the corresponding  $T$ -parameter to be  $\text{nis}(G)$ , where  $G$  is the graph after the  $i$ -th update;
    - (b) if the  $i$ -th update with edge endpoints  $u, v$  in  $G$  does not lead to an update in  $G^{(j)}$ , perform two consecutive updates to  $H$  in  $\mathcal{B}_j$ :
      - i. insert self-loop<sup>5</sup>  $(u, u)$  to  $H$ , and set the corresponding  $T$ -parameter to be  $\text{nis}(G')$ , where  $G'$  is the union of the graph  $G$  after the  $i$ -th update together with the self-loop  $(u, u)$ ;
      - ii. delete self-loop  $(u, u)$  from  $H$ , and set the corresponding  $T$ -parameter to be  $\text{nis}(G)$ , where  $G$  is the graph after the  $i$ -th update.
    - (c) update  $\mathcal{B}_j$  and  $\bar{c}_j$  according to the above edge update(s) in  $H$ .
5. Define  $\bar{M} := n - \bar{c}_r \cdot (1 + \varepsilon/2)^r + \sum_{j=0}^{r-1} \lambda_j \cdot \bar{c}_j$ .

Now we restate Theorem 1.3 and give its proof.

**Theorem 1.3.** *Let  $\varepsilon \in (0, 1)$  and  $W \geq 1$ . There exists a fully dynamic algorithm that with high probability, maintains an estimator  $\bar{M}$  that is a  $(1 + \varepsilon)$ -approximation of the weight  $M$  of MSF of a graph  $G$  with edge weights from  $[1, W]$ . The worst-case time per update operation is*

$$O\left((\log W)/\varepsilon + W^3(\log W) \log(W/\varepsilon) (\log(\log W/\varepsilon) + \log n)/(\varepsilon^4 \sqrt{m^*})\right),$$

where  $m^* \geq 1$  is the minimum number of edges in  $G$  throughout all the updates. Our algorithm works against an adaptive adversary.

**Proof.** Let  $p'$  be a parameter that will be specified later. Let  $r, \ell_j, \lambda_j, G^{(j)}, c^{(j)}, \bar{M}$  be defined as in the algorithm  $\text{RANDESTWMSF}(G, \varepsilon, p', W)$ . Note that  $\text{nis}(G) \geq \text{nis}(G^{(j)})$ , since  $G^{(j)}$  is a subgraph of  $G$  for any  $1 \leq j \leq r$ . Since  $G$  is simple, we know that  $M \geq \text{nis}(G)/2$ , as each non-isolated vertex is incident to at least one edge (of weight at least 1) of any MSF.

Note that  $\bar{c}_j$  is the estimator of  $c^{(j)}$  obtained by invoking the dynamic algorithm in Theorem 4.1 with  $p = p'/r$ ,  $\varepsilon' = \varepsilon/(24W)$  and the above specified dynamic graph  $H$  and  $T$ -parameter. As each update in  $G$  changes  $\text{nis}(G)$  by at most 2,  $T$ -parameter changes by at most 2 for each update in  $G^{(j)}$  (and thus fulfills requirement (b) of Theorem 4.1) if we guarantee that each update corresponding to edge  $(u, v)$  in  $G$  leads to at least one update in  $G^{(j)}$ . In the case that the weight of  $(u, v)$  is at most  $\ell_j$ , then  $(u, v)$  belongs to  $G^{(j)}$ , and, thus, the requirement is naturally fulfilled.

If the weight of  $(u, v)$  is larger than  $\ell_j$ , we execute two operations that cancel each other. Specifically we insert a self-loop  $(u, u)$  at Step 4(b)i and then immediately delete it at Step 4(b)ii, so that the graph  $G^{(j)}$  remains correct, but the  $T$ -parameters in these two steps also fulfill requirements (a) and (b). More precisely, after adding the self-loop  $(u, u)$ , it holds that the  $T$ -parameter is at least  $\text{nis}(H)$  as  $H$  is the number of non-isolated vertices of a subgraph of  $G$  with the self-loop  $(u, u)$ ; furthermore, the change of the  $T$ -parameter is at most 1 as the self-loop can change the number of non-isolated

<sup>5</sup> Note that if we do not introduce the self-loops, then it is not the case that each edge update in  $G$  leads to an update in every graph  $G^{(j)}$ . In particular, if the updated edge of  $G$  has weight larger than  $\ell_j$ , then the graphs  $G^{(i)}$  with  $\ell_i < \ell_j$  stay the same, i.e., they are *not* updated. This further leads to the issue that when one edge update occurs on  $G^{(i)}$ , the  $T$ -parameter for  $G^{(i)}$  may have changed by more than 2, i.e., the precondition of Theorem 4.1 is no longer fulfilled. This could happen if before this update in  $G^{(i)}$ ,  $G$  was updated many times (where each updated edge had weight more than  $\ell_i$ ) and thus the  $T$  parameter (i.e.,  $\text{nis}(G)$ ) of  $G^{(i)}$  has changed by more than 2.

vertices by at most 1, so the requirement (b) holds for this operation. After deleting the self-loop  $(u, u)$ , the requirement (a) is still guaranteed, as  $H$  will be exactly  $G^{(j)}$ , which implies that the corresponding  $T$ -parameter is at least  $\text{nis}(G) \geq \text{nis}(H)$ ; the requirement (b) still holds as deleting one self-loop can change the  $T$ -parameter by at most 1.

For simplicity, we slightly abuse the notation by letting  $T_i$  denote the specified  $T$ -parameter in the algorithm corresponding to the  $i$ -th update of  $G$  (i.e., it corresponds to the  $T$ -parameter in either Step 4a, 4(b)i or 4(b)ii). By Theorem 4.1 and the fact that  $\text{ncc}(H) = \text{ncc}(G^{(j)})$  (as  $H$  is either  $G^{(j)}$  itself or  $G^{(j)}$  with an additional self-loop  $(u, u)$ , for some vertex  $u$ ), the algorithm computes an estimator  $\bar{c}_j$  for  $c^{(j)}$  such that with probability  $1 - p'/r$ , it holds that

$$|\bar{c}_j - c^{(j)}| \leq \varepsilon' \cdot T_i \leq \varepsilon/(24W) \cdot (\text{nis}(G) + 1) \leq \varepsilon/(12W) \cdot \text{nis}(G),$$

where the second inequality holds as  $T_i \leq \text{nis}(G) + 1$  throughout the update, and the last inequality follows as  $\text{nis}(G) \geq \text{nis}^* \geq 1$ .

Thus with probability at least  $1 - r \cdot \frac{p'}{r} = 1 - p'$ , the above inequality holds for all  $j$  such that  $1 \leq j \leq r$ . Assume in the following that this event holds. Let  $X = n - c^{(r)} \cdot (1 + \varepsilon/2)^r + \sum_{j=0}^{r-1} \lambda_j \cdot c^{(j)}$ . By the same calculation as in the proof of Theorem 1.1, we have that

$$|\bar{M} - X| \leq \varepsilon \text{nis}(G)/4 \leq \varepsilon M/2$$

Now by Lemma 2.1 with approximation parameter  $\varepsilon/2$ , it holds that  $M \leq X \leq (1 + \varepsilon/2)M$ . Thus  $\bar{M}$  is a  $(1 + \varepsilon)$ -approximation of  $M$ .

Now note that for each  $j$ , it always holds that for the  $i$ -th update in  $G$  satisfies that  $T_i \geq \text{nis}(G) \geq \text{nis}^*$ , and thus the worst-case time spent per update for computing  $\bar{c}_j$  is

$$O(\max\{1, \log(1/\varepsilon') \log(1/p)/(\varepsilon'^3 \text{nis}^*)\}) = O(1 + W^3 \log(W/\varepsilon) \log(r/p')/(\varepsilon^3 \text{nis}^*)).$$

In total, the worst-case time per update operation of  $\text{RANDESTWMSF}(G, \varepsilon, p', W)$  is

$$\begin{aligned} & O\left(\sum_{j=1}^r (1 + W^3 \log(W/\varepsilon) \log(r/p')/(\varepsilon^3 \text{nis}^*))\right) \\ &= O((\log W)/\varepsilon + W^3 \log W \log(W/\varepsilon) \log(\log W/(\varepsilon p'))/(\varepsilon^4 \text{nis}^*)). \end{aligned}$$

Now we note that  $\text{nis}^* > \sqrt{2m^*}$ , where  $m^*$  is the minimum number of edges of the graph throughout all the updates. This is true as for the graph  $G^*$  with minimum non-isolated vertices, i.e.,  $\text{nis}(G^*) = \text{nis}^*$ , the total number of edges in  $G^*$  is both at most  $(\text{nis}^*)(\text{nis}^* - 1)/2 < (\text{nis}^*)^2/2$  and also at least  $m^*$ .

Therefore, by setting  $p' = 1/n^c$  for any constant  $c > 1$ , we can guarantee that the worst-case update time of the algorithm  $\text{RANDESTWMSF}(G, \varepsilon, p', W)$  is

$$O((\log W)/\varepsilon + W^3 (\log W) \log(W/\varepsilon) (\log(\log W/\varepsilon) + \log n)/(\varepsilon^4 \sqrt{m^*})),$$

and the algorithm succeeds with probability at least  $1 - p' = 1 - 1/n^c$ .

The algorithm works against an adaptive adversary as each of the algorithms from Theorem 4.1 works against an adaptive adversary, and the MSF algorithm returns only a weighted sum of the values returned by each of these algorithms.  $\square$

#### 4.1. Proof of Theorem 4.1

In the following, we present the proof of Theorem 4.1. We use  $H$  to denote the considered graph.

**Theorem 4.1.** *Let  $\varepsilon', p \in (0, 1)$ . Consider a dynamic algorithm that takes as input a dynamic graph  $H$  and an additional (integer) parameter  $T$ , called  $T$ -parameter, that is also changing along with each edge update. Suppose that (a)  $T \geq \text{nis}(H)$  throughout all the updates and (b) the  $T$ -parameters at timestamps  $i$  and  $i - 1$  differ by at most 2, for  $i \geq 1$ . Then there exists a fully dynamic algorithm  $\text{RANDDYNAMICNCC}(H, T, \varepsilon', p)$  that with probability at least  $1 - p$ , maintains an estimator  $\bar{c}\bar{c}$  for the number  $\text{ncc}$  of CCs of a dynamic graph  $H$  such that  $|\bar{c}\bar{c} - \text{ncc}(H)| \leq \varepsilon' \cdot T$ . The worst-case time per update operation is  $O(\max\{1, \log(1/\varepsilon') \log(1/p)/(\varepsilon'^3 T^*)\})$ , where  $T^*$  is the minimum of the  $T$ -parameters over all updates. Our algorithm works against an adaptive adversary.*

##### 4.1.1. A data structure for dynamically sampling non-zero entries

We first give a data structure to maintain some data structures so that the algorithm can perform *non-isolated vertex-sample queries*, i.e., to sample a vertex from the set of all non-isolated vertices uniformly at random *at any time*. Our dynamic algorithm for estimating the weight of an MSF will make use of this data structure.

We first present a more general data structure to sample non-zero entries from a dynamic array and then show how to use it to support non-isolated vertex-sample queries. Given a set  $V$  of  $n$  elements (here vertices), numbered from 0 to  $n - 1$ , each element  $u$  with an associated number  $d_u$  (here degree), we show how to support the following operations in constant time with preprocessing time  $O(n)$ :

- **Update**( $u, \delta$ ): add  $\delta$  to  $d_u$ , where  $\delta$  can be positive or negative.
- **Non-zero sample**( $\cdot$ ): return an element that is chosen uniformly at random from all elements  $u$  with  $d_u \neq 0$ .

**Lemma 4.2.** *There exists a data structure that supports **Non-zero sample**( $\cdot$ ) queries to a fully dynamic array with constant update time.*

**Proof.** Let us call an element  $u$  of  $V$  with  $d_u \neq 0$  a *non-zero* element. We implement the data structure by using two arrays and a counter:

1. Keep the number  $\text{nis}$  of non-zero elements of  $V$ .
2. Keep an array  $\mathcal{A}$  of size  $n$ , where only the first  $\text{nis}$  entries are used, such that (i) each entry in  $\mathcal{A}$  stores a non-zero element  $u$  together with  $d_u$  and (ii) each non-zero element of  $V$  is stored in  $\mathcal{A}$  within the first  $\text{nis}$  entries.
3. Keep an array  $\mathcal{P}$  of size  $n$ , which has an entry for each element of  $V$ , such that if an element  $u$  is stored in  $\mathcal{A}[i]$  (i.e.  $u$  is non-zero), then  $\mathcal{P}[u] = i$ ; and if an element  $u$  is not stored in  $\mathcal{A}$  (i.e.  $d_u = 0$ ), then  $\mathcal{P}[u] = -1$ . Thus  $\mathcal{P}$  contains indices corresponding to the positions of elements in  $\mathcal{A}$  or  $-1$ .

During preprocessing we initialize both arrays, set all entries of  $\mathcal{P}$  to  $-1$ , and set  $\text{nis} = 0$ . Then insert each  $u$  whose initial value  $d_u \neq 0$  by calling **Update**( $u, d_u$ ).

**Handling an Update**( $u, \delta$ ) **operation.** Whenever an **Update**( $u, \delta$ ) operation is executed, we check if  $\mathcal{P}[u] > -1$ .

1. If  $\mathcal{P}[u] > -1$ , then  $u$  is stored in  $\mathcal{A}$  and  $\mathcal{P}[u]$  contains the index of  $u$  in  $\mathcal{A}$ . Thus we add  $\delta$  to  $d_u$ , which is retrieved and then stored in the entry  $\mathcal{A}[\mathcal{P}[u]]$ . If the resulting value  $d_u \neq 0$ , this completes the update operation. If, however, the resulting value  $d_u = 0$ , let  $v$  be the element stored in  $\mathcal{A}[\text{nis}]$ . We copy into  $\mathcal{A}[\mathcal{P}[u]]$  all information of element  $v$ , which we retrieve from  $\mathcal{A}[\text{nis}]$ . Then we set  $\mathcal{P}[v] = \mathcal{P}[u]$ , set  $\mathcal{P}[u] = -1$ , and decrement  $\text{nis}$ .
2. If  $\mathcal{P}[u] = -1$ , then we increment  $\text{nis}$  by 1, set  $d_u = \delta$ , store  $u$  and  $d_u$  in  $\mathcal{A}[\text{nis}]$ , and set  $\mathcal{P}[u] = \text{nis}$ .

**Handling Non-zero sample operation.** To implement this operation, we pick a random integer  $j$  between 0 and  $\text{nis} - 1$  and return the element from  $\mathcal{A}[j]$ .  $\square$

**Supporting non-isolated vertex-sample queries in dynamic graphs.** Next we show how to use the above data structure to support non-isolated vertex-sample query throughout all the updates. Whenever an edge  $(u, v)$  is inserted, for each  $x \in \{u, v\}$ , we call **Update**( $x, 1$ ). Whenever an edge  $(u, v)$  is deleted, for each  $x \in \{u, v\}$ , we call **Update**( $x, -1$ ). To sample a non-isolated vertex, we call **Non-zero sample**( $\cdot$ ).

#### 4.1.2. The dynamic algorithm

*A static algorithm for estimating the number of CCs.* We will make use of a static algorithm, which we briefly describe here. Recall that  $\text{ncc}(H)$  denotes the number of CCs of a graph  $H$ . The following lemma was proven by Berenbrink et al. [21] (which improves upon the result in [17]) and it gives a constant-time algorithm for estimating  $\text{ncc}(H)$ . It is assumed that the algorithm can perform a *vertex-sample query*, which allows it to sample (in constant time) a vertex uniformly at random from  $V$ , and can make queries to the adjacency list of the graph. Note that these two queries for accessing a static graph can be supported by maintaining an array of vertices and the adjacency list of the graph, respectively.

**Lemma 4.3** ([21]). *Let  $\varepsilon > 0$  and  $0 < p < 1$ . Suppose the algorithm has access to the adjacency list of a graph  $H$  and can perform vertex-sample queries. Then there exists an algorithm that with probability at least  $1 - p$ , returns an estimate that approximates  $\text{ncc}(H)$  with an additive error  $\varepsilon n$ . The running time of the algorithm is  $O(1/\varepsilon^2 \log(1/\varepsilon) \log(1/p))$ .*

We remark that the algorithm in [21] simply samples (uniformly at random)  $O(1/\varepsilon^2)$  vertices, performs a BFS starting from each sampled vertex (for a number of steps) and then makes decisions based on the explored subgraphs. Let  $H_{\text{nis}}$  denote the subgraph induced by all non-isolated vertices in a graph  $H$ . Note that if the algorithm is able to perform a non-isolated vertex-sample query, i.e., the algorithm can sample a vertex uniformly at random from the vertex set of  $H_{\text{nis}}$ , then one can approximate the number  $\text{ncc}(H_{\text{nis}})$  of CCs in the subgraph  $H_{\text{nis}}$  with an additive error  $\varepsilon |V(H_{\text{nis}})| = \varepsilon \text{nis}(H)$ , where  $\text{nis}(H)$  is the number of non-isolated vertices in  $H$ . This is true as we can simply treat  $H[N]$  as the input graph in the algorithm from Lemma 4.3. This also proves the following lemma, Lemma 4.4, which will be invoked by our dynamic algorithm.

**Lemma 4.4.** *Let  $\varepsilon > 0$  and  $0 < p < 1$ . Suppose the algorithm has access to the adjacency list of a static graph  $H$  and can sample a vertex uniformly at random from the set of all non-isolated vertices in  $H$ . Then there exists an algorithm that with probability at least  $1 - p$ , returns an estimate  $\bar{b}$  that approximates  $\text{ncc}(H_{\text{nis}})$  with an additive error  $\varepsilon \cdot \text{nis}(H)$ . The running time of the algorithm is  $O(1/\varepsilon^2 \log(1/\varepsilon) \log(1/p))$ .*

*From static to dynamic.* The idea of our dynamic algorithm is as follows. By slightly abusing notation, let  $H_i$  denote the graph  $H$  after the  $i$ -th update (and  $H_0$  denotes the initial graph). Our algorithm first maintains the number  $\Gamma$  of non-isolated

vertices in the current graph in the straightforward way: we can maintain the degree of each vertex, and count the number of vertices with non-zero degrees, which can be done in constant update time.

During initialization we set  $\Psi = T_0$  and  $\bar{c} = \text{ncc}(H_0)$ , which is calculated by running a static BFS traversal on  $H_0$ . Then it partitions the updates (and the corresponding graphs) into *phases*. For notational convenience, we let phase  $\mathcal{P}_0$  consist only of the graph  $H_0$ . We let  $H'_j$  be the graph corresponding to the last update in  $\mathcal{P}_j$ , for any  $j \geq 0$ . At the beginning of each phase  $\mathcal{P}_j$  (for  $j \geq 1$ ), we are given two parameters, an estimate  $\bar{c}$  and  $\Psi$ , which correspond to estimator of the number of CCs and the number of non-isolated vertices in the graph  $H'_{j-1}$  corresponding to the last update in  $\mathcal{P}_{j-1}$ . The phase  $\mathcal{P}_j$  consists of all graphs corresponding to the next  $\frac{\varepsilon' \Psi}{4}$  updates. The parameters  $\bar{c}$  and  $\Psi$  remain unchanged (i.e., they are as fixed at the beginning of the phase) until the last update in  $\mathcal{P}_j$ . More precisely, let  $T'_j$  be the  $T$ -parameter corresponding to the last update in  $\mathcal{P}_j$ . At the end of  $\mathcal{P}_j$ , the algorithm sets  $\Psi$  to  $T'_j$ , runs the static algorithm from Lemma 4.4 to obtain an estimate  $\bar{b}$  for  $\text{ncc}([H'_j]_{\text{nis}})$  and re-set the estimate  $\bar{c}$  to  $\bar{b} + n - \text{nis}(H'_j)$ . Then we start with a new phase  $\mathcal{P}_{j+1}$  consisting of the graphs corresponding to the next  $\varepsilon' \Psi / 4$  updates, where  $\Psi$  is as fixed at the beginning of this phase and then we repeat the above.

Throughout, we maintain the adjacency list of the dynamic graph in a trivial way and maintain the data structure from Section 4.1.1 for sampling a vertex uniformly at random from the set of all non-isolated vertices from the current graph. When asked a query on the number of CCs of the current graph, the algorithm returns  $\bar{c}$ . The description of our algorithm RANDDYNAMICNCC is given as follows, omitting the details for maintaining adjacency list, array of degrees, and data structures for sampling non-isolated vertices.

**RANDDYNAMICNCC( $H, T, \varepsilon', p$ )**  $\triangleright$  **Maintaining an estimator  $\bar{c}$  for  $\text{ncc}(H)$  of a dynamic graph  $H$  (with parameter  $T \geq \text{nis}(H)$ ) with additive error  $\leq \varepsilon' \cdot T$**

1. Preprocessing: Use BFS in the initial graph  $H_0$  to obtain  $\text{nis}(H_0)$  and  $\text{ncc}(H_0)$ . Start of the first phase. Initialize  $\Gamma = \text{nis}(H_0)$  and  $\bar{c} = \text{ncc}(H_0)$ ,  $\Psi = T_0$ . Let  $i = 1$ .
2. For the  $i$ -th update:
  - (a) update  $\Gamma$  to be  $\text{nis}(H_i)$
  - (b) if  $i \bmod (\frac{\varepsilon' \Psi}{4}) = 0$ , then  $\triangleright$  The end of a phase
    - i. compute an estimator  $\bar{b}$  for  $\text{ncc}([H_i]_{\text{nis}})$  by running the static algorithm in Lemma 4.4 on  $H_i$  with parameter  $\varepsilon = \frac{\varepsilon'}{4}$  and  $p$
    - ii. set  $\bar{c} = \bar{b} + n - \Gamma \triangleright n - \Gamma$  is the number of isolated nodes in  $H_i$
    - iii. set  $\Psi = T_i$
  - (c) set  $i = i + 1$

#### 4.1.3. Putting it together

Now we prove the correctness of the above dynamic algorithm, analyze its running time and then finish the proof of Theorem 4.1.

**Correctness.** Let  $j \geq 0$  and let  $\mathcal{P}_j$  denote the set of all the graphs in the  $j$ -th phase in the algorithm. In particular,  $\mathcal{P}_0 = \{H_0\}$ . Let  $H'_j$  and  $T'_j$  be the graph and the  $T$ -parameter, respectively, corresponding to the last update in  $\mathcal{P}_j$ . We first prove the following claims.

**Claim 4.5.** For any  $j \geq 0$  and the graph  $H = H'_j$ , our estimate satisfies that  $|\bar{c} - \text{ncc}(H'_j)| \leq \varepsilon' \text{nis}(H'_j) / 4 \leq \varepsilon' \cdot T'_j / 4$  with probability  $1 - p$ .

**Proof.** For  $j = 0$ , it holds that  $H'_0 = H_0$  and our estimate  $\bar{c} = \text{ncc}(H_0)$  by definition. Let  $j \geq 1$ . Our algorithm calls the static algorithm from Lemma 4.4, which returns with probability  $1 - p$  an estimator  $\bar{b}$  for  $\text{ncc}([H'_j]_{\text{nis}})$  such that  $|\bar{b} - \text{ncc}([H'_j]_{\text{nis}})| \leq \frac{\varepsilon' \text{nis}(H'_j)}{4}$ . This gives  $|\bar{c} - \text{ncc}(H'_j)| = |\bar{b} + n - \text{nis}(H'_j) - \text{ncc}(H'_j)| = |\bar{b} - \text{ncc}([H'_j]_{\text{nis}})| \leq \frac{\varepsilon' \text{nis}(H'_j)}{4} \leq \frac{\varepsilon' T'_j}{4}$ , where the second equation follows from the fact that  $\text{ncc}(H'_j)$  is the sum of  $\text{ncc}([H'_j]_{\text{nis}})$  and the number of isolated vertices,  $n - \text{nis}(H'_j)$ .  $\square$

**Claim 4.6.** For any  $j \geq 0$ , any  $H \in \mathcal{P}_{j+1} \setminus \{H'_{j+1}\}$  that is the graph right after the  $i$ -th update, it holds that  $|\bar{c} - \text{ncc}(H)| \leq \varepsilon' T_i$  with probability  $1 - p$ .

**Proof.** Note that after update  $f_j$ ,  $\Psi = T'_j \geq \text{nis}(H'_j)$ ; both  $\Psi$  and the estimate  $\bar{c}$  remain unchanged before the last update in  $\mathcal{P}_{j+1}$ . As each update changes  $\text{ncc}(H)$  by at most 1, with at least probability  $1 - p$ , it holds that  $|\bar{c} - \text{ncc}(H)| \leq \varepsilon' \text{nis}(H'_j)/4 + \varepsilon' \Psi/4 \leq \varepsilon' \Psi/2$  for all graph  $H \in \mathcal{P}_{j+1} \setminus \{H'_{j+1}\}$ . Note that  $|\Psi - T_i| \leq \frac{\varepsilon' \Psi}{2}$ , as each update (during the next  $\frac{\varepsilon' \Psi}{4} - 1$  updates) can change  $T_i$  by at most 2. Thus  $\Psi \leq \frac{1}{1 - \varepsilon'/2} T_i \leq (1 + \varepsilon') T_i \leq 2T_i$ . This implies that  $\bar{c}$  approximates  $\text{ncc}(H)$  with an additive error  $\varepsilon' T_i$ .  $\square$

Thus with probability  $1 - p$ , for any  $H$  (right after the  $i$ -th update), we have  $|\bar{c} - \text{ncc}(H)| \leq \varepsilon' T_i$ . Finally, note that the algorithm uses “fresh” random bits at the beginning of each phase, only needs to access to the current graph, and does not reuse any information computed in prior phases. Within each phase we performed a worst-case analysis, i.e., we assumed that the adversary changes the graph in the worst possible way, i.e., changing  $\text{ncc}(H)$  by 1 in each update. Thus, our algorithm works against an adaptive adversary.

**Running time.** For each phase that contains all graphs in  $\mathcal{P}_{j+1} \setminus \{H'_{j+1}\}$  with parameter  $\Psi = T'_j$ , the amortized time (over the phase) per update operation is

$$O(\max\{1, (1/\varepsilon'^2) \log(1/\varepsilon') \log(1/p)/(\varepsilon' \Psi)\}) = O(\max\{1, \log(1/\varepsilon') \log(1/p)/(\varepsilon'^3 \Psi)\}).$$

(Note that we always need to use  $O(1)$  time to update the adjacency list and other data structures so as to provide query access to the graph). Since  $T^*$  is the minimum value  $T_i$  over graphs throughout all the updates, then in any phase,  $\Psi \geq T^*$  and the amortized update time is  $O(\max\{1, \frac{\log(1/\varepsilon') \log(1/p)}{\varepsilon'^3 T^*}\})$ . The worst-case update time guarantee follows from the standard global rebuilding technique.

## 5. Lower bounds: Proofs of Theorem 1.2 and 1.5

### 5.1. A lower bound for deterministic data structure

Our first lower bound is built upon a construction by Henzinger and Fredman [24]. The *parity prefix sum* problem is defined as follows: Given an array  $A[1], \dots, A[n]$  with entries from  $\{0, 1\}$ , initialized with 0, build a data structure that executes an arbitrary sequence of the following operations: (1)  $\text{Add}(i)$ : Set  $A[i] = (A[i] + 1) \bmod 2$ ; (2)  $\text{Sum}(i)$ : return  $\sum_{j=1}^i A[j] \bmod 2$ . Fredman and Saks [25] showed the following theorem.

**Theorem 5.1** ([25]). *There is an amortized lower bound of  $\Omega(\log n / (\log \log n + \log b))$  update time per operation for the parity prefix sum problem in the cell probe model with word size  $b$ .*

Now we show the following.

**Theorem 5.2.** *Let  $G$  be a dynamic  $n$ -vertex graph with edge weights in  $[1, W]$ . Any data structure that dynamically maintains the weight  $M$  of an MSF of a graph  $G$  within an additive error less than  $W/2$  for any  $W \geq 1$ , or a multiplicative factor of  $(1 + \varepsilon)$ , for any  $W > 2\varepsilon n$  must perform  $\Omega(\log n / (\log \log n + \log b))$  cell probes, where each cell has size  $b$ .*

**Proof.** The proof is analogous to the lower bound for dynamic connectivity in [24]. Give a parity prefix sum problem build a graph with  $n$  vertices, labeled from  $1, \dots, n$ , one vertex called *even*, and one vertex called *odd*. At any point in time the graph maintains the following invariant: It consists of two paths, namely an *even* path containing all vertices  $i$  (called “even” vertices) such that  $\text{Sum}(i)$  returns 0, in order of their indices, and such that the vertex in the chain is connected to vertex *even*; and an analogous *odd* path containing all “odd” vertices, connected to vertex *odd*. It is shown that each  $\text{Add}(i)$  only leads to a constant number of edge updates in the graph. To determine which edges need to be deleted and inserted, a Van-Emde-Boas priority queue is maintained, which adds a cost of  $O(\log \log n)$  to each  $\text{Add}(i)$  operation. To answer a parity query  $\text{Sum}(i)$  return 1 iff vertex *odd* and vertex  $i$  are connected.

Now we show how to use this construction to give a lower bound for the dynamic approximate MSF weight problem. We maintain the same graph as above, and give each edge weight 1. Thus the weight of an MSF before and after every operation is  $n - 2$ . To answer a  $\text{Sum}(i)$  query we insert an edge of weight  $W$  from vertex *odd* to vertex  $i$ . If the answer is 1, then  $i$  and *odd* were already connected before the edge insertion and the weight of an MSF is unchanged. If the answer is 0, then  $i$  and *odd* were not connected before the edge insertion and, thus, the weight of an MSF is now  $n - 2 + W$ . Thus any dynamic MSF algorithm with multiplicative actor of  $(1 + \varepsilon)$ , for any  $W > 2\varepsilon n$  and any dynamic MSF algorithm with additive error less than  $W/2$  for any  $W \geq 1$ , can distinguish the two cases. Thus, they must perform  $\Omega(\log n / (\log \log n + \log b))$  cell probes, where each cell has size  $b$ .  $\square$

Now we are ready to prove Theorem 1.2.

**Theorem 1.2.** *Let  $W \geq (\log n)^{\omega_n(1)}$ . Let  $G$  be a dynamic graph with edge weights in  $[1, W]$ . Any deterministic data structure that dynamically maintains the weight  $M$  of an MSF of a graph  $G$  within a multiplicative factor of  $(1 + \varepsilon)$  for any constant  $\varepsilon > 0$ , or an additive error less than  $W/2$  must perform  $\omega_n(1)$  cell probes, where each cell has size  $O(\log n)$ .*

**Proof.** We let  $G'$  be a graph with  $n'$  vertices and edge weights in  $[1, W]$ , for  $n' = \lfloor \frac{W}{3\varepsilon} \rfloor$ . By Theorem 5.2, any data structure that dynamically maintains the weight  $M$  of an MSF of a graph  $G'$  with  $W > 2\varepsilon n'$  within an additive error less than  $W/2$  for any  $W \geq 1$ , or a multiplicative factor of  $(1 + \varepsilon)$  must perform  $\Omega(\log n' / (\log \log n' + \log b))$  cell probes, where each cell has size  $b$ .

Note that  $n' = (\log n)^{\omega_n(1)}/\varepsilon$ . Now assume that we have a data structure  $\mathcal{A}$  that dynamically maintains the  $(1 + \varepsilon)$ -approximation of the weight of an MSF of an  $n$ -vertex graph. (The case that  $\mathcal{A}$  has an additive error less than  $W/2$  can be analyzed similarly.) We use  $\mathcal{A}$  to the graph  $\tilde{G}$  which consists of  $G'$  and  $n - n'$  isolated vertices. Note that the weight of an MSF of  $\tilde{G}$  is the same as the weight of an MSF of  $G'$ . Thus  $\mathcal{A}$  maintains a  $(1 + \varepsilon)$ -approximation of the weight of an MSF of  $G'$ .

By setting  $b = \Theta(\log n)$ , we know that  $\mathcal{A}$  must perform  $\Omega(\log n' / (\log \log n' + \log b)) = \omega_n(1)$  cell probes, where each cell has size  $O(\log n)$ .  $\square$

## 5.2. A lower bound for randomized data structure

We let  $m^*(G)$  be the minimum number of edges of a dynamic graph  $G$  throughout all the updates. We first prove the following.

**Theorem 5.3.** *Let  $G$  be a dynamic graph with edge weights in  $[1, W]$ . Any data structure that dynamically maintains the weight  $M$  of an MSF of a graph  $G$  with  $W = \Omega(\varepsilon \cdot m^*(G))$  and  $m^*(G) = \Theta(n)$  within an additive error less than  $W/2$  for any  $W \geq 1$ , or a multiplicative factor of  $(1 + \varepsilon)$ , with probability at least  $1 - \frac{1}{n^c}$  for some constant  $c > 0$ , must perform  $\Omega((\log n)^2/b)$  cell probes, where each cell has size  $b$ .*

**Proof.** In [16] Patrascu and Demaine construct an  $n$ -node graph and show that there exists a sequence  $\mathcal{S}$  of  $T$  edge insertion, edge deletion, and query operations such that any data structure for dynamic connectivity must perform  $\Omega((\log n)^2/b)$  cell probes to process the sequence, where each cell has size  $b$ . This shows that the amortized number of cell probes per operation is  $\Omega((\log n)^2/b)$ .

The graph  $G$  in the proof of [16] consists of a  $\sqrt{n} \times \sqrt{n}$  grid, where each node in column 1 has exactly 1 edge to a node of column 2 and no other edges, each node in column  $i$ , with  $1 < i < \sqrt{n}$  has exactly 1 edge to a node of column  $i - 1$  and 1 edge to a node of column  $i + 1$  and no other edges, and each node in column  $\sqrt{n}$  has exactly 1 edge to a node of column  $\sqrt{n} - 1$  and no other edges. Thus, the graph consists of  $\sqrt{n}$  paths of length  $\sqrt{n} - 1$  and the edges between column  $i$  and  $i + 1$  for any  $1 \leq i < \sqrt{n}$  represent a permutation of the  $\sqrt{n}$  rows. The sequence  $\mathcal{S}$  consists of “batches” of  $O(\sqrt{n})$  edge updates, replacing the permutation of some column  $i$  by a new permutation for column  $i$ . Between the batches of updates are “batches” of connectivity queries, each consisting of  $\sqrt{n}$  connectivity queries and a parameter  $1 \leq k \leq \sqrt{n}$ , where the  $j$ -th query for  $1 \leq j \leq \sqrt{n}$  of each batch tests whether the  $j$ -th vertex of column 1 is connected with a specific vertex of column  $k$ .

We now show how to use this construction to give a lower bound for the dynamic approximate MSF problem. We add a new vertex  $s$  and an edge between  $s$  and every vertex in column 1 and give weight 1 to every edge. We now show how to modify each connectivity query  $(u, v)$  such that it consists of a constant number of edge updates and one query for the value of an MSF. Thus, in the resulting sequence  $\mathcal{S}'$  the number of query operations equals the number of query operations in  $\mathcal{S}$  and the number of update operations is linear in the number of update and query operations in  $\mathcal{S}$ . Thus the total number of operations in  $\mathcal{S}'$  is only a constant factor larger than the number of operations in  $\mathcal{S}$ , which, together with the result of [16], implies that the amortized number of cell probes per operation is  $\Omega((\log n)^2/b)$ .

We now show how to simulate a connectivity query  $(u, v)$ , where  $u$  is in column 1 and  $v$  is in column  $k$  for some  $1 \leq k \leq \sqrt{n}$ . To simulate a connectivity query  $(u, v)$  we (1) remove the edge from  $u$  to  $s$ , (2) add the edge  $(u, v)$  with weight  $W$  and then (3) ask a query for the weight  $M$  of an MSF. Afterwards we undo the changes to  $G$ . Note that if  $u$  and  $v$  are connected in  $G$  then the edge  $(u, v)$  does not belong to the MSF (and the graph is disconnected), otherwise it does. Furthermore, the value  $M$  will be  $n - 2$  if  $u$  and  $v$  are connected, and  $n - 2 + W$  if  $u$  and  $v$  are not connected. This implies that if our data structure maintains  $M$  within an additive error less than  $W/2$ , for any  $W \geq 1$ , or within a multiplicative factor of  $1 + \varepsilon$  for  $W > 2\varepsilon n$ , then we can also test connectivity. Then the statement of the theorem follows from the lower bound for testing connectivity. Finally, we note that  $m^*(G) = n - O(\sqrt{n}) = \Theta(n)$  by construction of the graph and thus  $W = \Omega(\varepsilon m^*(G)) > 2\varepsilon n$ . This finishes the proof.  $\square$

Now we are ready to give the proof of Theorem 1.5.

**Theorem 1.5.** *For any constant  $\varepsilon, \alpha > 0$  and  $W = n^\alpha$ , any data structure that dynamically maintains with high probability the weight of an MSF of a graph  $G$  with edge weights in  $[1, W]$  and  $W = \Omega(\varepsilon m^*)$  within a multiplicative factor of  $(1 + \varepsilon)$ , or an additive error less than  $W/2$ , must perform  $\Omega(\log n)$  cell probes, where each cell has size  $O(\log n)$ .*



**Proof.** The proof is similar to the proof of Theorem 1.2.

We let  $G'$  be a graph with  $n'$  vertices and edge weights in  $[1, W]$ , for  $n' = \lfloor \frac{W}{3\varepsilon} \rfloor$ . By Theorem 5.3, any data structure that dynamically maintains the weight  $M$  of an MSF of a graph  $G'$  with  $W = \Omega(\varepsilon \cdot m^*(G'))$  and  $m^*(G') = \Theta(n')$  within an additive error less than  $W/2$  for any  $W \geq 1$ , or a multiplicative factor of  $(1 + \varepsilon)$ , with probability at least  $1 - \frac{1}{(n')^c}$ , must perform  $\Omega((\log n')^2/b)$  cell probes, where each cell has size  $b$ .

Note that  $n' = \Theta(n^\alpha/\varepsilon)$  where  $\alpha > 0$  is a constant. Now assume that we have a data structure  $\mathcal{A}$  that dynamically maintains the  $(1 + \varepsilon)$ -approximation of the weight of an MSF of an  $n$ -vertex graph with probability at least  $1 - \frac{1}{n^{c'}} = 1 - \frac{1}{(n')^c}$  for some constant  $c' > 0$ . (The case that  $\mathcal{A}$  has an additive error less than  $W/2$  can be analyzed similarly.) We use  $\mathcal{A}$  to the graph  $\tilde{G}$  which is consisted of  $G'$  and  $n - n'$  isolated vertices. Note that the weight of an MSF of  $\tilde{G}$  is the same as the weight of an MSF of  $G'$ . Thus  $\mathcal{A}$  maintains a  $(1 + \varepsilon)$ -approximation of the weight of an MSF of  $G'$ . Note that  $m^*(G) = m^*(G')$ , and that  $W = \Omega(\varepsilon \cdot m^*(G')) = \Omega(\varepsilon \cdot m^*(G))$  and  $m^*(G) = m^*(G') = \Theta(n')$ .

By setting  $b = \Theta(\log n)$ , we know that  $\mathcal{A}$  must perform  $\Omega((\log n')^2/b) = \Omega(\log^2(W/\varepsilon)/b) = \Omega(\log n)$  cell probes, where each cell has size  $O(\log n)$ .  $\square$

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Appendix A. A note on dynamically estimating the number of CCs

*Estimating  $\text{ncc}(G)$  with an additive error  $\varepsilon n^{O(1)}$ .* We note that similar to our previous algorithm (in Section 4) for estimating  $\text{ncc}(G)$  with an additive error  $\varepsilon T(G)$ , if we simply invoke the static algorithm from Lemma 4.4 on the current graph  $H = G$  with parameters  $\text{nis}(H) = n$  and  $\text{ncc}(H_{\text{nis}}) = \text{ncc}(G)$  to obtain an estimator  $\overline{cc}$  and re-compute the estimator every  $\Theta(\varepsilon n)$  updates, then the corresponding algorithm always maintain an estimator for  $\text{ncc}(G)$  with an additive error  $\varepsilon n$ . That is, we have the following theorem.

**Theorem A.1.** *Let  $1 > \varepsilon > 0$  and  $0 < p < 1$ . There exists a fully dynamic algorithm that with probability at least  $1 - p$ , maintains an estimator  $\overline{cc}$  for the number  $\text{ncc}$  of CCs of a graph  $G$  s.t.,  $|\overline{cc} - \text{ncc}(G)| \leq \varepsilon \cdot n$ . The worst-case time per update operation is  $O(\max\{1, \frac{\log(1/\varepsilon)\log(1/p)}{\varepsilon^3 n}\})$ .*

The following is a direct corollary of the above theorem.

**Corollary A.2.** *Let  $\varepsilon > 0$  and let  $c$  be any constant such that  $c \geq 1$ . There exists a fully dynamic algorithm that with probability at least  $1 - \frac{1}{n^c}$ , maintains an estimator  $\overline{cc}$  for the number  $\text{ncc}$  of CCs of a graph  $G$  s.t.,  $|\overline{cc} - \text{ncc}(G)| \leq \varepsilon n^{2/3} \log^{2/3} n$ . The worst-case time per update operation is  $O(\varepsilon^{-3})$ .*

**Proof.** If  $\varepsilon < n^{-\frac{2}{3}}$ , then for each update, one can use the naive BFS algorithm to exactly compute  $\text{ncc}(G)$ , which runs in time  $O(n^2) = O(\varepsilon^{-3})$ . If  $\varepsilon \geq n^{-\frac{2}{3}}$ , we can apply Theorem A.1 with parameters  $p = \frac{1}{n^c}$ ,  $\varepsilon' = \varepsilon n^{-1/3} \cdot \log^{2/3} n$ , to obtain an  $\overline{cc}$  for  $\text{ncc}(G)$  with an additive error  $\varepsilon n^{2/3} \log^{2/3} n$ . The corresponding dynamic algorithm has update time  $O(\varepsilon^{-3} \cdot (1 + \frac{\log(1/\varepsilon)}{\log n})) = O(\varepsilon^{-3})$ .  $\square$

**Remark.** We cannot expect to be able to get a constant-time algorithm for maintaining the number of connected components with an additive error of 1 or a multiplicate error of 2: Any such algorithm would be able to decide whether the graph is connected or not, contradicting the  $\Omega(\log n)$  lower bound for dynamically maintaining whether a graph is connected [16].

## References

- [1] G.N. Frederickson, Data structures for on-line updating of minimum spanning trees (preliminary version), in: STOC, 1983, pp. 252–257.
- [2] D. Eppstein, Z. Galil, G.F. Italiano, A. Nissenzweig, Sparsification - a technique for speeding up dynamic graph algorithms, J. ACM 44 (5) (1997) 669–696.
- [3] M.R. Henzinger, V. King, Randomized fully dynamic graph algorithms with polylogarithmic time per operation, J. ACM 46 (4) (1999) 502–516.
- [4] M.R. Henzinger, V. King, Maintaining minimum spanning trees in dynamic graphs, in: International Colloquium on Automata, Languages, and Programming, Springer, 1997, pp. 594–604.
- [5] J. Holm, K. De Lichtenberg, M. Thorup, Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity, J. ACM 48 (4) (2001) 723–760.
- [6] J. Holm, E. Rotenberg, C. Wulff-Nilsen, Faster fully-dynamic minimum spanning forest, in: Algorithms-ESA 2015, Springer, 2015, pp. 742–753.
- [7] C. Wulff-Nilsen, Fully-dynamic minimum spanning forest with improved worst-case update time, in: Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, ACM, 2017, pp. 1130–1143.
- [8] D. Nanongkai, T. Saranurak, Dynamic spanning forest with worst-case update time: adaptive, Las Vegas, and  $O(n^{1/2-\varepsilon})$ -time, in: Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, ACM, 2017, pp. 1122–1129.

- [9] D. Nanongkai, T. Saranurak, C. Wulff-Nilsen, Dynamic minimum spanning forest with subpolynomial worst-case update time, in: Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on, IEEE, 2017, pp. 950–961.
- [10] M. Thorup, Near-optimal fully-dynamic graph connectivity, in: Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21–23, 2000, Portland, OR, USA, 2000, pp. 343–350.
- [11] B.M. Kapron, V. King, B. Mountjoy, Dynamic graph connectivity in polylogarithmic worst case time, in: SODA, 2013, pp. 1131–1141.
- [12] D. Gibb, B. Kapron, V. King, N. Thorn, Dynamic graph connectivity with improved worst case update time and sublinear space, preprint, arXiv:1509.06464, 2015.
- [13] S. Huang, D. Huang, T. Kopelowitz, S. Pettie, Fully dynamic connectivity in  $O(\log n(\log \log n)^2)$  amortized expected time, in: Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16–19, 2017, 2017, pp. 510–520.
- [14] J. Chuzhoy, Y. Gao, J. Li, D. Nanongkai, R. Peng, T. Saranurak, A deterministic algorithm for balanced cut with applications to dynamic connectivity, flows, and beyond, preprint, arXiv:1910.08025, 2019.
- [15] C. Wulff-Nilsen, Faster deterministic fully-dynamic graph connectivity, in: Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2013, pp. 1757–1769.
- [16] M. Patrascu, E.D. Demaine, Logarithmic lower bounds in the cell-probe model, SIAM J. Comput. 35 (4) (2006) 932–963.
- [17] B. Chazelle, R. Rubinfeld, L. Trevisan, Approximating the minimum spanning tree weight in sublinear time, SIAM J. Comput. 34 (6) (2005) 1370–1379.
- [18] A. Czumaj, C. Sohler, Estimating the weight of metric minimum spanning trees in sublinear time, SIAM J. Comput. 39 (3) (2009) 904–922.
- [19] K.J. Ahn, S. Guha, A. McGregor, Analyzing graph structure via linear measurements, in: Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, Society for Industrial and Applied Mathematics, 2012, pp. 459–467.
- [20] M. Gupta, R. Peng, Fully dynamic  $(1 + \epsilon)$ -approximate matchings, in: Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on, IEEE, 2013, pp. 548–557.
- [21] P. Berenbrink, B. Krayenhoff, F. Mallmann-Trenn, Estimating the number of connected components in sublinear time, Inf. Process. Lett. 114 (11) (2014) 639–642.
- [22] H. Jowhari, M. Sağlam, G. Tardos, Tight bounds for lp samplers, finding duplicates in streams, and related problems, in: Proceedings of the Thirtieth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, ACM, 2011, pp. 49–58.
- [23] M.R. Henzinger, V. King, Maintaining minimum spanning forests in dynamic graphs, SIAM J. Comput. 31 (2) (2001) 364–374.
- [24] M.R. Henzinger, M.L. Fredman, Lower bounds for fully dynamic connectivity problems in graphs, Algorithmica 22 (3) (1998) 351–362.
- [25] M. Fredman, M. Saks, The cell probe complexity of dynamic data structures, in: Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing (STOC), ACM, 1989, pp. 345–354.
- [26] S. Solomon, Fully dynamic maximal matching in constant update time, in: Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on, IEEE, 2016, pp. 325–334.
- [27] S. Bhattacharya, J. Kulkarni, Deterministically maintaining a  $(2 + \epsilon)$ -approximate minimum vertex cover in  $O(1/\epsilon^2)$  amortized update time, in: Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2019, pp. 1872–1885.
- [28] S. Baswana, S. Khurana, S. Sarkar, Fully dynamic randomized algorithms for graph spanners, ACM Trans. Algorithms 8 (4) (2012) 35.
- [29] M. Henzinger, P. Peng, Constant-time dynamic  $(\delta + 1)$ -coloring, in: 37th International Symposium on Theoretical Aspects of Computer Science, STACS 2020, 2020, pp. 53:1–53:18.
- [30] S. Bhattacharya, F. Grandoni, J. Kulkarni, Q.C. Liu, S. Solomon, Fully dynamic  $(\delta + 1)$ -coloring in constant update time, preprint, arXiv:1910.02063, 2019.
- [31] D. Eppstein, G.F. Italiano, R. Tamassia, R.E. Tarjan, J. Westbrook, M. Yung, Maintenance of a minimum spanning forest in a dynamic plane graph, J. Algorithms 13 (1) (1992) 33–54.
- [32] D. Eppstein, Z. Galil, G.F. Italiano, T.H. Spencer, Separator based sparsification: I. planarity testing and minimum spanning trees, J. Comput. Syst. Sci. 52 (1) (1996) 3–27.
- [33] Z. Huang, P. Peng, Dynamic graph stream algorithms in  $o(n)$  space, Algorithmica 81 (5) (2019) 1965–1987.
- [34] P. Peng, C. Sohler, Estimating graph parameters from random order streams, in: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2018, pp. 2449–2466.