

# Average Sensitivity of Graph Algorithms\*

Nithin Varma  
University of Haifa  
nvarma@bu.edu

Yuichi Yoshida  
National Institute of Informatics  
yyoshida@nii.ac.jp

## Abstract

In modern applications of graph algorithms, where the graphs of interest are large and dynamic, it is unrealistic to assume that an input representation contains the full information of a graph being studied. Hence, it is desirable to use algorithms that, even when provided with only a (large) subgraph, output solutions that are close to the solutions output when the whole graph is available. We formalize this feature by introducing the notion of average sensitivity of graph algorithms, which is the average earth mover’s distance between the output distributions of an algorithm on a graph and its subgraph obtained by removing an edge, where the average is over the edges removed and the distance between two outputs is the Hamming distance.

In this work, we initiate a systematic study of average sensitivity. After deriving basic properties of average sensitivity such as composition, we provide efficient approximation algorithms with low average sensitivities for concrete graph problems, including the minimum spanning forest problem, the global minimum cut problem, the minimum  $s$ - $t$  cut problem, and the maximum matching problem. In addition, we prove that the average sensitivity of our global minimum cut algorithm is almost optimal, by showing a nearly matching lower bound. We also show that every algorithm for the 2-coloring problem has average sensitivity linear in the number of vertices. One of the main ideas involved in designing our algorithms with low average sensitivity is the following fact; if the presence of a vertex or an edge in the solution output by an algorithm can be decided locally, then the algorithm has a low average sensitivity, allowing us to reuse the analyses of known sublinear-time algorithms and local computation algorithms. Using this fact in conjunction with our average sensitivity lower bound for 2-coloring, we show that every local computation algorithm for 2-coloring has query complexity linear in the number of vertices, thereby answering an open question.

---

\*A preliminary version of this manuscript appeared in the proceedings of SODA 2021 [51].

# 1 Introduction

In modern applications of graph algorithms, where the graphs of interest are large and dynamic, it is unrealistic to assume that an input representation contains the full information of a graph being studied. For example, consider a social network, where a vertex corresponds to a user of the social network service and an edge corresponds to a friendship relation. It is reasonable to assume that users do not always update new friendship relations on the social network service, and that sometimes they do not fully disclose their friendship relations because of security or privacy reasons. Hence, we can only obtain an approximation  $G'$  to the true social network  $G$ . This brings out the need for algorithms that can extract information on  $G$  by solving a problem on  $G'$ . Moreover, as the solutions output by a graph algorithm are often used in applications such as detecting communities [38, 39], ranking nodes [43], and spreading influence [23], the solutions output by an algorithm on  $G'$  should be close to those output on  $G$ .

We adopt a popular model that the  $n$ -vertex input graph  $G'$  at hand is a randomly chosen (large) subgraph of an unknown true graph  $G$  [1, 26, 40]. Intuitively, a deterministic algorithm  $\mathcal{A}$  is said to be *stable-on-average* if the Hamming distance  $d_{\text{Ham}}(\mathcal{A}(G), \mathcal{A}(G'))$  is small, where  $\mathcal{A}(G)$  and  $\mathcal{A}(G')$  are outputs of  $\mathcal{A}$  on  $G$  and  $G'$ , respectively. Here, outputs are typically vertex sets or edges sets and we assume that they are represented appropriately using binary strings. More specifically, for an integer  $k \geq 1$ , we say that the  *$k$ -average sensitivity* of a deterministic algorithm  $\mathcal{A}$  is

$$\mathbb{E}_{\{e_1, \dots, e_k\} \sim \binom{E}{k}} [d_{\text{Ham}}(\mathcal{A}(G), \mathcal{A}(G - \{e_1, \dots, e_k\}))] \quad (1)$$

for every graph  $G = (V, E)$ , where  $\{e_1, \dots, e_k\}$  is sampled uniformly at random from  $\binom{E}{k}$ , the set of all subsets of  $E$  of cardinality  $k$ , and where  $G - F$  for a set of edges  $F \subseteq E$  denotes the subgraph obtained from  $G$  by removing  $F$ . When  $k = 1$ , we call the  $k$ -average sensitivity simply *average sensitivity*. We say that algorithms with low average sensitivity are *stable-on-average*. Although we focus on graphs here, we note that our definition can also be extended to the study of combinatorial objects other than graphs such as strings and constraint satisfaction problems.

An algorithm that outputs the same solution regardless of the input has the least possible average sensitivity, even though it is definitely useless. Hence, the key question in a study of average sensitivity is to reveal trade-offs between solution quality and average sensitivity for various problems.

**Example 1.1.** Consider the algorithm that, given a graph  $G = (V, E)$  on  $n$  vertices, outputs the set of vertices of degree at least  $n/2$ . As removing an edge changes the degree of exactly two vertices, the sensitivity of this algorithm is at most 2.

**Example 1.2.** Consider the  $s$ - $t$  shortest path problem, where given a graph  $G = (V, E)$  and two vertices  $s, t \in V$ , we are to output the set of edges in a shortest path from  $s$  to  $t$ . Since the length of a shortest path is always bounded by  $n$ , where  $n$  is the number of vertices, every deterministic algorithm has average sensitivity  $O(n)$ . Indeed, there exists a graph for which this trivial upper bound is tight. Think of a cycle of even length  $n$  and two vertices  $s, t$  in diametrically opposite positions. Consider an arbitrary deterministic algorithm  $\mathcal{A}$ , and assume that it outputs a path  $P$  (of length  $n/2$ ) among the two shortest paths from  $s$  to  $t$ . With probability half, an edge in  $P$  is removed, and  $\mathcal{A}$  must output the other path  $Q$  (of length  $n/2$ ) from  $s$  to  $t$ . Hence, the average sensitivity must be  $1/2 \cdot (n/2) = \Omega(n)$ . In this sense, there is no deterministic algorithm with nontrivial average sensitivity for the  $s$ - $t$  shortest path problem.

We also generalize our definition of average sensitivity to apply to randomized algorithms. Let  $\mathcal{A}(G)$  denote the output distribution of  $\mathcal{A}$  on  $G$ . Let  $d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}(G'))$  denote the earth mover's distance between  $\mathcal{A}(G)$  and  $\mathcal{A}(G')$ , where the distance between two outputs is measured by the Hamming distance. Specifically,  $d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}(G'))$  is equal to  $\min_{\mathcal{D}} [\mathbb{E}_{(x,y) \sim \mathcal{D}} [d_{\text{Ham}}(x, y)]]$ , where  $\mathcal{D}$  denotes a distribution over pairs  $(x, y)$  of outputs of  $\mathcal{A}$  such that the left and right marginals of  $\mathcal{D}$  are equal to  $\mathcal{A}(G)$  and  $\mathcal{A}(G')$ , respectively. Then, for an integer  $k \geq 1$ , the *k-average sensitivity* of a randomized algorithm  $\mathcal{A}$  is

$$\mathbb{E}_{\{e_1, \dots, e_k\} \sim \binom{E}{k}} [d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}(G - \{e_1, \dots, e_k\}))] \quad (2)$$

where  $\{e_1, \dots, e_k\}$  is sampled uniformly at random from  $\binom{E}{k}$ . Note that when the algorithm  $\mathcal{A}$  is deterministic, (2) matches the definition of the average sensitivity for deterministic algorithms.

**Remark 1.3.** The *k-average sensitivity* of an algorithm  $\mathcal{A}$  with respect to the total variation distance can be defined as  $\mathbb{E}_{\{e_1, \dots, e_k\} \sim \binom{E}{k}} [d_{\text{TV}}(\mathcal{A}(G), \mathcal{A}(G - \{e_1, \dots, e_k\}))]$ , where  $d_{\text{TV}}(\cdot, \cdot)$  denotes the total variation distance between two distributions. It is easy to observe that, if the *k-average sensitivity* of an algorithm with respect to the total variation distance is at most  $\gamma(G)$ , then its *k-average sensitivity* is bounded by  $H \cdot \gamma(G)$ , where the  $H$  is the maximum over Hamming weights of all solutions output (with nonzero probability) by running  $\mathcal{A}$  on  $G = (V, E)$  and on all the graphs in  $\{G - \{e_1, \dots, e_k\} : \{e_1, \dots, e_k\} \in \binom{E}{k}\}$ .

**Example 1.4.** Randomness does not help improve the average sensitivity of algorithms for the *s-t* shortest path problem. Think of the cycle graph given in Example 1.2, and suppose that a randomized algorithm  $\mathcal{A}$  outputs the *s-t* paths  $P$  and  $Q$  with probability  $p$  and  $q = 1 - p$ , respectively. Then, the average sensitivity is  $p \cdot 1/2 \cdot (n/2) + q \cdot 1/2 \cdot (n/2) = \Omega(n)$ .

## 1.1 Basic properties of average sensitivity

Our definition of average sensitivity has many nice properties. In this section, we discuss some such properties of average sensitivity that are useful in the design of our stable-on-average algorithms. We denote by  $\mathcal{G}$  the (infinite) set consisting of all graphs. Given a graph  $G = (V, E)$  and  $e \in E$ , we use  $G - e$  as a shorthand for  $G - \{e\}$ . We use  $n$  and  $m$  to denote the number of vertices and edges in the input graph, respectively.

**Bounds on *k-average sensitivity* from bounds on average sensitivity.** This is one of the most important properties of our definition of average sensitivity. It essentially says that a bound on the (1-)average sensitivity of an algorithm can be used to obtain a bound on the *k-average sensitivity* of that algorithm for  $k \geq 1$ . In other words, it is enough to analyze the average sensitivity of an algorithm with respect to the removal of a single edge.

**Theorem 1.5.** *Let  $\mathcal{A}$  be an algorithm for a graph problem with average sensitivity at most  $f(n, m)$ . Then, for any integer  $k \geq 1$ , the algorithm  $\mathcal{A}$  has *k-average sensitivity* at most  $\sum_{i=1}^k f(n, m - i + 1)$ .*

In particular, if the average sensitivity of an algorithm is bounded from above by a nondecreasing function of the number of edges, then its *k-average sensitivity* is at most  $k$  times the upper bound on its average sensitivity.

**Sequential composition.** Another useful feature of our definition of average sensitivity is that one can obtain a stable-on-average algorithm by sequentially applying several stable-on-average subroutines. The following two *sequential composition theorems* formalize this feature.

**Theorem 1.6** (Sequential composition). *Consider two randomized algorithms  $\mathcal{A}_1 : \mathcal{G} \rightarrow \mathcal{S}_1, \mathcal{A}_2 : \mathcal{G} \times \mathcal{S}_1 \rightarrow \mathcal{S}_2$ . Suppose that the average sensitivity of  $\mathcal{A}_1$  with respect to the total variation distance is  $\gamma_1(G)$  and the average sensitivity of  $\mathcal{A}_2(\cdot, S_1)$  is  $\beta_2^{(S_1)}(G)$  for any  $G \in \mathcal{G}, S_1 \in \mathcal{S}_1$ . Let  $\mathcal{A} : \mathcal{G} \rightarrow \mathcal{S}_2$  be a randomized algorithm obtained by composing  $\mathcal{A}_1$  and  $\mathcal{A}_2$ , that is,  $\mathcal{A}(G) = \mathcal{A}_2(G, \mathcal{A}_1(G))$  for  $G \in \mathcal{G}$ . Then, the average sensitivity of  $\mathcal{A}$  on a graph  $G \in \mathcal{G}$  is  $\mathsf{H} \cdot \gamma_1(G) + \mathbb{E}_{S_1 \sim \mathcal{A}_1(G)} \left[ \beta_2^{(S_1)}(G) \right]$ , where  $\mathsf{H}$  denotes the maximum over Hamming weights of all the solutions output (with nonzero probability) by running  $\mathcal{A}$  on  $G$  and all of the graphs in  $\{G - e : e \in E\}$ .*

Our second composition theorem is for the average sensitivity with respect to the total variation distance. This is also useful for analyzing the average sensitivity with respect to the earth mover's distance, as it can be bounded by the average sensitivity with respect to the total variation distance times the maximum over Hamming weights of solutions output, as in Remark 1.3.

**Theorem 1.7** (Sequential composition w.r.t. the TV distance). *Consider  $\ell$  randomized algorithms  $\mathcal{A}_i : \mathcal{G} \times \prod_{j=1}^{i-1} \mathcal{S}_j \rightarrow \mathcal{S}_i$  for  $i \in \{1, \dots, \ell\}$ . Suppose that, for each  $i \in \{1, \dots, \ell\}$ , the average sensitivity of  $\mathcal{A}_i(\cdot, S_1, \dots, S_{i-1})$  is  $\gamma_i(G)$  with respect to the total variation distance for every  $G \in \mathcal{G}$  and  $S_1 \in \mathcal{S}_1, \dots, S_{i-1} \in \mathcal{S}_{i-1}$ . Consider a sequence of computations  $S_1 = \mathcal{A}_1(G), S_2 = \mathcal{A}_2(G, S_1), \dots, S_\ell = \mathcal{A}_\ell(G, S_1, \dots, S_{\ell-1})$  for  $G \in \mathcal{G}$ . Let  $\mathcal{A} : \mathcal{G} \rightarrow \mathcal{S}_\ell$  be a randomized algorithm that performs this sequence of computations on input  $G$  and outputs  $S_\ell$ . Then, the average sensitivity of  $\mathcal{A}$  on a graph  $G \in \mathcal{G}$  with respect to the total variation distance is at most  $\sum_{i=1}^{\ell} \gamma_i(G)$ .*

**Parallel composition.** It is often the case that there are multiple algorithms that solve the same problem albeit with different average sensitivity guarantees. Such stable-on-average algorithms can be combined via *parallel composition*, where we run these algorithms according to a distribution determined by the input graph. The advantage of parallel composition is that the average sensitivity of the resulting algorithm might be better than that of the component algorithms.

**Theorem 1.8** (Parallel composition). *Let  $\mathcal{A}_1, \dots, \mathcal{A}_\ell$  be algorithms for a graph problem with average sensitivities  $\beta_1(G), \dots, \beta_\ell(G)$ , respectively. Let  $\mathcal{A}$  be an algorithm that, given a graph  $G$ , runs  $\mathcal{A}_i$  with probability  $\rho_i(G)$  for  $i \in \{1, \dots, \ell\}$ , where  $\sum_{i \in \{1, \dots, \ell\}} \rho_i(G) = 1$ . Let  $\mathsf{H}$  denote the maximum over Hamming weights of all solutions output (with nonzero probability) by running  $\mathcal{A}$  on  $G$  and on all the graphs in  $\{G - e : e \in E\}$ . Then the average sensitivity of  $\mathcal{A}$  is at most  $\sum_{i \in \{1, \dots, \ell\}} \rho_i(G) \cdot \beta_i(G) + \mathsf{H} \cdot \mathbb{E}_{e \sim E} \left[ \sum_{i \in \{1, \dots, \ell\}} |\rho_i(G) - \rho_i(G - e)| \right]$ .*

In this paper, we use the above theorem extensively to combine algorithms with different average sensitivities.

## 1.2 Connection to sublinear-time algorithms

We show a relationship between the average sensitivity of an algorithm and the query complexity of a sublinear-time algorithm [41, 16, 52] that simulates oracle access to the solution output by the former algorithm. Roughly speaking, we show, in Theorem 1.10, that the average sensitivity of an algorithm  $\mathcal{A}$  is bounded by the query complexity of another algorithm  $\mathcal{O}$ , which we call a

*solution oracle*, where  $\mathcal{O}$  queries the edges of the graph  $G$  and simulates oracle access to the solution produced by  $\mathcal{A}$  on input  $G$ . We first formalize the notion of a solution oracle.

**Definition 1.9** (Solution Oracle). Consider a deterministic algorithm  $\mathcal{A} : \mathcal{G} \rightarrow \mathcal{S}$  for a graph problem, where each solution output by  $\mathcal{A}$  is a subset of the set of edges of the input graph. An algorithm  $\mathcal{O}$  is a *solution oracle* for  $\mathcal{A}$  if  $\mathcal{O}$  satisfies:

- $\mathcal{O}$  has access to a graph  $G = (V, E)$ , which is represented as adjacency lists, via *neighbor queries*, where each query is of the form  $(v, i)$  for  $v \in V$  and  $i \in [|V|]$  and the answer is the  $i$ -th neighbor of vertex  $v$  in its adjacency list (and a special symbol if  $i$  is larger than the degree of  $v$ ),
- given an edge  $e \in E$  as input,  $\mathcal{O}$  queries  $G$  and outputs whether  $e$  is contained in the solution obtained by running  $\mathcal{A}$  on  $G$ .

Given a randomized algorithm  $\mathcal{A}$  and an associated function  $r : \mathbb{N} \rightarrow \mathbb{N}$ , the solution oracle  $\mathcal{O}$  of  $\mathcal{A}$  first generates a random string  $\pi \in \{0, 1\}^{r(|V|)}$  and then runs the solution oracle  $\mathcal{O}_\pi$  of the deterministic algorithm  $\mathcal{A}_\pi$  obtained by fixing the randomness of  $\mathcal{A}$  to  $\pi$ .

Note that an analogous definition can be made for algorithms that output a subset of vertices.

**Theorem 1.10** (Sublinearity implies low average sensitivity). *Consider a randomized algorithm  $\mathcal{A} : \mathcal{G} \rightarrow \mathcal{S}$  for a graph problem, where each solution output by  $\mathcal{A}$  is a subset of the set of edges in the input graph. Assume that there exists a solution oracle  $\mathcal{O}$  for  $\mathcal{A}$  such that  $\mathcal{O}$  makes at most  $q(G)$  queries to  $G$  in expectation, where this expectation is taken over the random coins of  $\mathcal{O}$  and over input edges  $e \in E$ . Then,  $\mathcal{A}$  has average sensitivity at most  $q(G)$ . Moreover, given the promise that the input graphs satisfy  $|E| \geq |V|$ , the statement applies also to algorithms for which each solution is a subset of the set of vertices in the input graph.*

We use Theorem 1.10 to design a stable-on-average matching algorithm (Theorem 6.5) based on a sublinear-time matching algorithm due to Yoshida et al. [52].

Closely related to Theorem 1.10 is Corollary 1.12, which says that if a graph problem has a *local computation algorithm* (LCA), then one can design a stable-on-average algorithm for that problem. LCAs, whose definition we give below, were introduced by Rubinfeld et al. [48] and has been widely studied ever since [2, 11, 17, 27, 29, 30, 31, 32, 33, 44, 47]. For more information on LCAs, we refer the interested reader to an excellent survey on the topic by Levi and Medina [28].

**Definition 1.11** (Local Computation Algorithm (LCA)). Consider a graph problem  $\mathcal{P} : \mathcal{G} \rightarrow \mathcal{S}$ , where the output to the problem is a subset of edges of the input graph. Let  $\delta : \mathbb{N} \rightarrow [0, 1]$  and  $q, r : \mathbb{N} \rightarrow \mathbb{N}$ . A  $(q, r, \delta)$ -LCA for  $\mathcal{P}$  is an algorithm  $\mathcal{L}$  that, given query access to a graph  $G = (V, E)$  (as in Definition 1.9), first generates a random string  $\pi \in \{0, 1\}^{r(|V|)}$ , and satisfies:

- given an input  $e \in E$ , the algorithm  $\mathcal{L}$  makes at most  $q(|V|)$  queries to  $G$  and answers whether  $e$  is part of a solution to the problem  $\mathcal{P}$  on graph  $G$ , and
- the answers of  $\mathcal{L}$  to all possible input edges are consistent with a single feasible solution to  $\mathcal{P}$  on  $G$ .

For every graph  $G$ , the probability (over the choice of random string) that there exists an input edge for which  $\mathcal{L}$  makes more than  $q(|V|)$  queries is at most  $\delta(|V|)$ .

We mention that  $\mathcal{L}$  is not allowed to perform any preprocessing on the graph. Additionally, the same set of edges is queried by  $\mathcal{L}$  when the same edge is given as input multiple times.

Note that one can have an analogous definition of LCAs for graph problems where each solution is a subset of vertices. We have the following result which is a direct corollary of Theorem 1.10.

**Corollary 1.12** (LCAs imply stable-on-average algorithms). *Consider a graph problem  $\mathcal{P} : \mathcal{G} \rightarrow \mathcal{S}$ . Let  $\delta : \mathbb{N} \rightarrow [0, 1]$  and  $q, r : \mathbb{N} \rightarrow \mathbb{N}$ . If  $\mathcal{P}$  has a  $(q, r, \delta)$ -LCA  $\mathcal{L}$ , then, there exists an algorithm  $\mathcal{A}$  for  $\mathcal{P}$ , that on input  $G = (V, E)$ , has an average sensitivity of at most  $q(|V|) + |E| \cdot \delta(|V|)$ .*

Theorem 1.10 and Corollary 1.12 cement the intuition that strong locality guarantees for solutions output by an algorithm imply that the removal of edges from a graph affects only the presence of a few elements (edges or vertices) in the solution, which in turn implies low average sensitivity. On the contrapositive side, Corollary 1.12 implies that a lower bound on the average sensitivity of algorithms for a problem implies a lower bound on the query complexity of an LCA (with failure probability  $o(1/n^2)$ ) for the same problem, where  $n$  denotes the number of vertices. Exploiting this result, we show that every LCA for 2-coloring has query complexity  $\Omega(n)$ , thereby answering an open question raised by Czumaj et al. [7]. We believe that this connection has the potential to shed more light on fundamental limits of LCAs and is of independent interest.

### 1.3 Stable-on-average algorithms for concrete problems

We summarize, in Table 1, the average sensitivity bounds that we obtain for various concrete problems. We use  $n$ ,  $m$ ,  $\text{OPT}$  to denote the number of vertices, the number of edges, and the optimal value.

All of our algorithms run in polynomial time, and for  $k \geq 1$ , upper bounds on  $k$ -average sensitivity of these algorithms can be easily obtained using Theorem 1.5. Except in the case of our algorithm for the minimum spanning forest problem, our stable-on-average algorithms are all randomized. Our lower bounds hold for any (randomized) algorithm to solve the respective problems with the specified approximation guarantee.

For the minimum spanning forest problem, we show that the classical Kruskal’s algorithm [24] has average sensitivity  $O(n/m)$ , which is at most 1, and is quite small considering that Kruskal’s algorithm is deterministic and that the spanning forest can have  $\Omega(m)$  edges. We also show a matching lower bound of  $\Omega(n/m)$  for the minimum spanning forest problem, implying that the average sensitivity of Kruskal’s algorithm is optimal. In contrast, we show that Prim’s algorithm can have average sensitivity  $\Omega(m)$  for a natural (and deterministic) rule of breaking ties among edges.

For the global minimum cut problem, we show that every algorithm that outputs the exact mincut has to have average sensitivity  $\Omega(n)$ . However, by allowing for a multiplicative approximation guarantee of  $2 + \epsilon$  for  $\epsilon > 0$ , we design a global minimum cut algorithm with average sensitivity  $n^{O(\frac{1}{\epsilon \text{OPT}})}$ . If  $\text{OPT} = \Omega(\log n)$ , the average sensitivity of our algorithm is  $O(1)$ , which is quite small. We also prove a nearly tight lower bound on the average sensitivity of any algorithm that guarantees a purely multiplicative approximation to the minimum cut size. In particular, when  $\text{OPT}$  is  $o(\log n)$ , our lower bound matches, up to a polylogarithmic factor, our upper bound for 3-approximating the minimum cut size.

Our lower bound of  $\Omega(n)$  on the average sensitivity of algorithms that output the exact global minimum cut also applies to the minimum  $s$ - $t$  cut problem. In contrast, we show that it is possible

Table 1: Our results. Here  $n$ ,  $m$ ,  $\text{OPT}$  denote the number of vertices, the number of edges, and the optimal value, respectively, and  $\epsilon \in (0, 1)$  is an arbitrary constant. The notation  $\tilde{O}(\cdot)$  hides polylogarithmic factors in  $n$ . Approximation guarantees are multiplicative unless specified otherwise.

Problem	Output	Approximation Guarantee	Average Sensitivity	Reference
Minimum Spanning Forest	Edge set	1	$O(\frac{n}{m})$	Sec. 3
		$< \infty$	$\Omega(\frac{n}{m})$	Sec. 3
Global Minimum Cut	Vertex set	$2 + \epsilon$	$n^{O(\frac{1}{\epsilon \text{OPT}})}$	Sec. 4.1
		1	$\Omega(n)$	Sec. 4.2
		$< \infty$	$\Omega\left(\frac{n \text{OPT}}{\text{OPT}^2}\right)$	Sec. 4.2
Minimum $s$ - $t$ Cut	Vertex set	additive $O(n^{2/3})$	$O(n^{2/3})$	Sec. 5
Maximum Matching	Edge set	1/2	1	Sec. 6.2
		$1 - \epsilon$	$\tilde{O}\left(\left(\frac{\text{OPT}}{\epsilon^3}\right)^{\frac{1}{1+\Omega(\epsilon^2)}}\right)$	Sec. 6.3
		1	$\Omega(n)$	Sec. 6.4
Minimum Vertex Cover	Vertex set	2	2	Sec. 6.2
2-Coloring	Vertex set	—	$\Omega(n)$	Sec. 7

to achieve average sensitivity of  $O(n^{2/3})$  for the minimum  $s$ - $t$  cut problem by allowing for an additive  $O(n^{2/3})$  approximation.

We show that the average sensitivity of every algorithm that outputs the exact maximum matching is  $\Omega(n)$ , implying that some approximation is essential to obtain nontrivial average sensitivity. We also propose two stable-on-average approximation algorithms for maximum matching. Our first algorithm has approximation ratio 1/2 and average sensitivity at most 1. This result immediately implies a 2-approximation algorithm for the minimum vertex cover problem with average sensitivity at most 2. Our second algorithm for maximum matching has approximation ratio  $1 - \epsilon$  and average sensitivity  $\tilde{O}\left(\left(\text{OPT}/\epsilon^3\right)^{1/(1+\Omega(\epsilon^2))}\right)$  for every constant  $\epsilon \in (0, 1)$ .

In the 2-coloring problem, given a bipartite graph, we are to output one part in the bipartition. For this problem, we show a lower bound of  $\Omega(n)$  for the average sensitivity, that is, there is no algorithm with nontrivial average sensitivity.

**Implications on the query complexity of LCAs.** Recall that, by Corollary 1.12, the average sensitivity lower bound for a problem implies an identical lower bound on the query complexity of an LCA (with failure probability  $o(1/n^2)$ ) for the same problem. This implies that every LCA for exact maximum matching, global minimum cut, and minimum  $s$ - $t$  cut has query complexity  $\Omega(n)$ . Additionally, every LCA giving a purely multiplicative approximation guarantee for the global minimum cut has query complexity  $\Omega\left(\frac{n \text{OPT}}{\text{OPT}^2}\right)$ . Additionally, as mentioned earlier, every LCA for 2-coloring has query complexity  $\Omega(n)$ , which answers an open question raised by Czumaj et al. [7].

## 1.4 Discussions on average sensitivity

**Output representation.** Average sensitivity is dependent on the output representation. For example, we can double the average sensitivity by duplicating the output. A natural idea for alleviating this issue is to normalize the average sensitivity by the maximum Hamming weight  $H$  of a solution. However, for minimization problems where the optimal value  $\text{OPT}$  could be much smaller than  $H$ , such a normalization can diminish subtle differences in average sensitivity, e.g.,  $O(\text{OPT}^{1/2})$  vs  $O(\text{OPT})$ . It is an interesting open question whether there is a canonical way to normalize average sensitivity so that the resulting quantity is independent of the output representation.

**Sensitivity against adversarial edge removals.** It is also natural to take the maximum, instead of the average, over edges in definitions (1) and (2), which can be seen as sensitivity against adversarial edge removals. Indeed a similar notion has been proposed to study algorithms for geometric problems [36]. However, in the case of graph algorithms, it is hard to guarantee that the output of an algorithm does not change much after removing an arbitrary edge. Moreover, by a standard averaging argument, one can say that for 99% of arbitrary edge removals, the sensitivity of an algorithm is asymptotically equal to the average sensitivity, which is sufficient in most cases.

**Average sensitivity w.r.t. edge additions.** As another variant of average sensitivity, it is natural to consider incorporating edge additions in definitions (1) and (2). If an algorithm is stable-on-average against edge additions, then in addition to the case of not knowing the true graph as we have discussed earlier, it will be useful for the case that the graph dynamically changes but we want to prevent the output of the algorithm from fluctuating too much. However, in contrast to removing edges, it is not always clear how we should add edges to the graph in definitions (1) and (2). A naive idea is sampling  $k$  pairs of vertices uniformly at random and adding edges between them. This procedure makes the graph close to a graph sampled from the Erdős-Rényi model [10], which does not represent real networks such as social networks and road networks well. To avoid this subtle issue, in this work, we focus on removing edges.

**Alternative notion of average sensitivity for randomized algorithms.** Consider a randomized algorithm  $\mathcal{A}$  that, given a graph  $G$  on  $n$  vertices, generates a random string  $\pi \in \{0, 1\}^{r(n)}$  for some function  $r : \mathbb{N} \rightarrow \mathbb{N}$ , and then runs a deterministic algorithm  $\mathcal{A}_\pi$  on  $G$ , where the algorithm  $\mathcal{A}_\pi$  has  $\pi$  hardwired into it. Assume that  $\mathcal{A}_\pi$  can be applied to any graph. It is also natural to define the average sensitivity of  $\mathcal{A}$  as

$$\mathbb{E}_{e \sim E} \left[ \mathbb{E}_\pi [d_{\text{Ham}}(\mathcal{A}_\pi(G), \mathcal{A}_\pi(G - e))] \right]. \quad (3)$$

In other words, we measure the expected distance between the outputs of  $\mathcal{A}$  on  $G$  and  $G - e$  when we feed the same string  $\pi$  to  $\mathcal{A}$ , over the choice of  $\pi$  and edge  $e$ . Note that (3) upper bounds (2) because, in the definition of the earth mover's distance, we optimally transport probability mass from  $\mathcal{A}(G)$  to  $\mathcal{A}(G - e)$  whereas, in (3), how the probability mass is transported is not necessarily optimal.

We can actually bound (3) for some of our algorithms. In this work, however, we focus on the definition (2) because the assumption that  $\mathcal{A}_\pi$  can be applied to any graph does not hold in general, since the length of  $\pi$  might depend on parameters of the input graph other than the number of vertices. Moreover, bounding (3) is unnecessarily tedious and is not very enlightening.



## 1.5 Overview of our techniques

**Global minimum cut.** For the global minimum cut problem, our algorithm is inspired by a differentially private algorithm<sup>1</sup> for the same problem by Gupta et al. [15]. Our algorithm, given a parameter  $\varepsilon > 0$  and a graph  $G$  as input, first enumerates a list of cuts whose sizes are at most  $(2 + \varepsilon) \cdot \text{OPT}$ ; this enumeration can be done in polynomial time as shown by Karger’s theorem [19]. The algorithm then outputs a cut from the list with probability exponentially small in the product of the size of the cut and  $O(1/\varepsilon \cdot \text{OPT})$ . The main argument in analyzing the average sensitivity of the algorithm is that the aforementioned distribution is very close (in earth mover’s distance) to a related Gibbs distribution on the set of all cuts in the graph. Therefore the average sensitivity of the algorithm is of the same order as that of the average sensitivity of sampling a cut from such a Gibbs distribution, where the latter sampling task requires exponential time. We finally show that the average sensitivity of sampling a cut from this Gibbs distribution is at most  $n^{O(1/\varepsilon \text{OPT})}$ .

**Minimum  $s$ - $t$  cut.** The first stage of our algorithm consists of solving an LP relaxation for the minimum  $s$ - $t$  cut problem in a stable way, for an appropriately defined notion of sensitivity. Given a graph  $G = (V, E)$  and vertices  $s, t \in V$ , the relaxation contains variables  $\mathbf{d}(\{u, v\}) \in [0, 1]$  for each  $\{u, v\} \in \binom{V}{2}$ , where these variables can be thought of as representing a *pseudometric* over the vertices. The constraints include triangle inequalities, and also a special constraint  $\mathbf{d}(\{s, t\}) = 1$ . The objective is to minimize the sum of the variables associated with the edges in  $G$ . Intuitively, if  $\mathbf{d}(\{s, v\})$  is *large* in a solution to the linear program, then the vertex  $v$  falls on the  $t$ -side of the  $s$ - $t$  cut represented by the solution.

Our stable LP solving strategy works by solving a related LP that is identical to the original LP, except for a regularization term added to the objective function. This regularization term is  $n^{-1/3}$  times the  $\ell_2$  norm of the vector of variables  $(\mathbf{d}(\{s, v\}) : v \in V)$ , where we use  $\|\mathbf{d}\|_s$  to denote this norm. It is easy to show that the value of the optimal solution to the regularized LP is within an additive  $O(n^{2/3})$  of the value of the optimal solution to the original LP. We also show that for all  $e \in E$ , the solutions output by our LP solver graphs  $G$  and  $G - e$  are close to each other with respect to  $\|\cdot\|_s$ . Here, we use the fact that the regularized objective function is strongly convex with respect to  $\|\cdot\|_s$ .

Given a solution to the (regularized) LP relaxation, our rounding procedure samples a threshold  $\tau \in [0, 1]$  uniformly at random and outputs the set  $S$  consisting of all vertices  $u \in V$  such that  $d(\{s, u\}) \leq \tau$ . The approximation guarantee of this algorithm follows from the fact that we are rounding based on a near optimal solution to the linear programming relaxation. To analyze the average sensitivity of the algorithm, we first show that the earth mover’s distance (with respect to Hamming distance) between the outputs of the rounding procedure for inputs  $\mathbf{d}, \mathbf{d}' \in [0, 1]^{\binom{V}{2}}$  is bounded by the  $\ell_1$  distance between the vectors  $(\mathbf{d}(\{s, v\}) : v \in V)$  and  $(\mathbf{d}'(\{s, v\}) : v \in V)$ . Combining this with the bound on the average sensitivity of our LP solving strategy, we obtain our final bound on the average sensitivity for our algorithm to approximate the minimum  $s$ - $t$  cut.

**Maximum matching.** Our stable-on-average  $\frac{1}{2}$ -approximation algorithm for the maximum matching problem first considers a uniformly random ordering of the edges in the input graph, and then greedily adds edges to the matching according to that ordering. In the context of dynamic distributed algorithms, Censor-Hillel et al. [6] showed that at most 1 edge changes in the matching

---

<sup>1</sup>We compare and contrast the definitions of average sensitivity and differential privacy in Section 1.6.

(in expectation) due to the removal of a uniformly random edge, where the expectation is taken over the edge removed and the ordering of edges. This result immediately implies that the average sensitivity of this randomized greedy matching algorithm is at most 1. In addition, it implies a 2-approximation algorithm for minimum vertex cover with average sensitivity at most 2.

There are several components to the design and analysis of our stable-on-average  $(1 - \varepsilon)$ -approximation algorithm. Our starting point is the observation (Theorem 1.10) that the existence of a sublinear-time solution oracle  $\mathcal{O}$  (see Definition 1.9) for an algorithm  $\mathcal{A}$  implies that  $\mathcal{A}$  is stable-on-average. We use Theorem 1.10 to bound the average sensitivity of a  $(1 - \varepsilon')$ -approximation algorithm  $\mathcal{A}$  for the maximum matching problem, where  $\varepsilon' = \Omega(\varepsilon)$ . Specifically,  $\mathcal{A}$  constructs a matching by considering augmenting paths of increasing length, and augmenting the (initially empty) matching iteratively, where the paths of each length are considered in a uniformly random order. Yoshida et al. [52] constructed a sublinear-time solution oracle that, given a uniformly random edge  $e \in E$  as input, makes  $O\left(\Delta^{O(1/(\varepsilon')^2)}\right)$  queries to  $G$  in expectation and answers whether  $e$  is in the matching output by  $\mathcal{A}$  on  $G$ , where the expectation is over the choice of input  $e$  and the randomness in  $\mathcal{A}$ , and  $\Delta$  is the maximum degree of  $G$ . Combined with Theorem 1.10, this implies that the average sensitivity of  $\mathcal{A}$  is  $O\left(\Delta^{O(1/(\varepsilon')^2)}\right)$ .

Next, we transform  $\mathcal{A}$  to also work for graphs of unbounded degree as follows. The idea is to remove vertices of degree at least  $\frac{m}{\varepsilon' \text{OPT}}$  from the graph and run  $\mathcal{A}$  on the resulting graph. This transformation affects the approximation guarantee only by an additive  $\varepsilon' \text{OPT}$  term, since the number of such *high degree* vertices is small. However, this thresholding procedure might itself have high average sensitivity, since the thresholds for  $G$  and  $G - e$  can be very different for all  $e \in E$ .

We circumvent this issue by using a Laplace random variable  $L$  as the threshold, where the distribution of  $L$  is tightly concentrated around  $\frac{m}{\varepsilon' \text{OPT}}$ . We use our sequential composition theorem (Theorem 1.6) in order to analyze the average sensitivity of the resulting procedure, where we consider the instantiation of the Laplace random threshold as the first algorithm, and the remaining steps in the procedure as the second algorithm. The first term in the expression given by Theorem 1.6 turns out to be a negligible quantity and is easy to bound. The main task in bounding the second term is to bound, for all  $x \in \mathbb{R}$ , the average sensitivity of a procedure  $\mathcal{A}_x$  that, on the input graph  $G$ , removes all vertices of degree at least  $x$  from  $G$  and runs the augmenting paths-based matching algorithm. The heart of the argument in bounding this average sensitivity is that given a solution oracle  $\mathcal{O}$  with query complexity  $q(\Delta)$  for an algorithm  $\mathcal{A}$ , we can, for all  $x \in \mathbb{R}$ , construct a solution oracle  $\mathcal{O}_x$  for the algorithm  $\mathcal{A}_x$ . Moreover, the query complexity of  $\mathcal{O}_x$  is at most  $O(x^2 q(x))$ . By Theorem 1.10, this is also a bound on the average sensitivity of  $\mathcal{A}_x$ . Using this, we bound the second term in the expression given by Theorem 1.6 as  $\mathbb{E}_L [O(L^2 q(L))] = O\left(\left(\frac{m}{\varepsilon' \text{OPT}}\right)^{O(1/(\varepsilon')^2)}\right)$ .

An issue with the aforementioned matching algorithm is that its average sensitivity is poor for graphs with small values of  $\text{OPT}$ . We observe that, in contrast to this, the algorithm that simply outputs the lexicographically smallest maximum matching has average sensitivity  $O(\text{OPT}^2/m)$ , since the output matching stays the same unless an edge in the matching is removed. We obtain our final stable-on-average  $(1 - \varepsilon)$ -approximation algorithm for the maximum matching problem by running these two algorithms according to a probability distribution determined by the input graph. Using our parallel composition theorem, we bound the average sensitivity of the resultant algorithm as  $\tilde{O}\left(\left(\text{OPT}/\varepsilon^3\right)^{1/(1+\Omega(\varepsilon^2))}\right)$ .

**2-coloring.** To show our  $\Omega(n)$  lower bound on the average sensitivity for 2-coloring, consider the set of all paths on  $n$  vertices and the set of all graphs obtained by removing exactly one edge from these paths (called 2-part-paths). A path has exactly two ways of being 2-colored and a 2-part-path has four ways of being 2-colored. A path and 2-part-path are neighbors if the latter is obtained from the former by removing an edge. A 2-part-path has at most four neighbors. The output distribution of any 2-coloring algorithm  $\mathcal{A}$  on a 2-part-path can be close (in earth mover’s distance) only to those of at most 2 of its neighboring paths. If  $\mathcal{A}$ , however, has low average sensitivity, the output distributions of  $\mathcal{A}$  have to be close on a large fraction of pairs of neighboring graphs, which gives a contradiction.

## 1.6 Related work

**Average sensitivity of network centralities.** (*Network centrality*) is a collective name for indicators that measure importance of vertices or edges in a network. Notable examples are closeness centrality [3, 4, 49], harmonic centrality [34], betweenness centrality [12], and PageRank [43]. To compare these centralities qualitatively, Murai and Yoshida [37] recently introduced the notion of average-case sensitivity for centralities. Fix a vertex centrality measure  $c$ ; let  $c_G(v)$  denote the centrality of a vertex  $v \in V$  in a graph  $G = (V, E)$ . Then, the *average-case sensitivity* of  $c$  on  $G$  is defined as

$$S_c(G) = \mathbb{E}_{e \sim E} \mathbb{E}_{v \sim V} \frac{|c_{G-e}(v) - c_G(v)|}{c_G(v)},$$

where  $e$  and  $v$  are sampled uniformly at random. They showed various upper and lower bounds for centralities. See [37] for details.

Since a centrality measure assigns real values to vertices, they studied the relative change of the centrality values upon removal of random edges. As our focus in this work is on graph algorithms, our notion (2) measures the Hamming distance between solutions when one removes random edges.

**Average sensitivity of spectral clustering.** In a subsequent work to this study, Pan and Yoshida [45] analyzed the average sensitivity of  $k$ -way spectral clustering to assess its reliability and efficiency. They showed that the average sensitivity is proportional to  $\lambda_k/\lambda_{k+1}^2$ , where  $\lambda_i$  is the  $i$ -th smallest eigenvalue of the normalized Laplacian. This result and higher-order Cheeger’s inequality [25] suggest that, if the input graph can be well partitioned into  $k$  components but not more,  $k$ -way spectral clustering is stable-on-average and hence can be safely used. If not,  $k$ -way spectral clustering may not be stable-on-average, but then we should not try to find  $k$ -way clustering in the first place.

**Differential privacy.** *Differential privacy* [8] is a notion closely related to average sensitivity. Assuming the existence of a neighbor relation over inputs, the definition of differential privacy requires that the distributions of outputs on neighboring inputs are similar. The variant of differential privacy closest to our definition of average sensitivity is edge differential privacy introduced by Nissim et al. [42] and further studied by [18, 15, 21, 22, 20, 46]. Here, the neighbors of a graph  $G = (V, E)$  are defined to be  $\{G - e\}_{e \in E}$ . For  $\varepsilon > 0$ , we say that an algorithm is  $\varepsilon$ -*differentially private* if for all  $e \in E$ ,

$$\exp(-\varepsilon) \cdot \Pr[\mathcal{A}(G - e) \in \mathcal{S}] \leq \Pr[\mathcal{A}(G) \in \mathcal{S}] \leq \exp(\varepsilon) \cdot \Pr[\mathcal{A}(G - e) \in \mathcal{S}] \quad (4)$$

for any set of solutions  $\mathcal{S}$ .

Differential privacy has stricter requirements than average sensitivity. Firstly, differential privacy is a *worst-case* sensitivity notion. Moreover, since differential privacy guarantees that the probabilities of outputting a specific solution on  $G$  and  $G - e$  are close to each other, the total variation distance between the two distributions  $\mathcal{A}(G)$  and  $\mathcal{A}(G - e)$  must be small. The earth mover’s distance between two output distributions can be small even if the total variation distance between them is large, and therefore, even if an algorithm is not differentially private, it could still be stable-on-average. Despite these differences, our stable-on-average algorithm for the global minimum cut problem is inspired by a differentially private algorithm for the same problem [15].

**Generalization and stability of learning algorithms.** Generalization [50] is a fundamental concept in statistical learning theory. Given samples  $\mathbf{z}_1, \dots, \mathbf{z}_n$  from an unknown true distribution  $\mathcal{D}$  over a dataset, the goal of a learning algorithm  $\mathcal{L}$  is to output a parameter  $\theta$  that minimizes *expected loss*  $\mathbb{E}_{\mathbf{z} \sim \mathcal{D}}[\ell(\mathbf{z}; \theta)]$ , where  $\ell(\mathbf{z}; \theta)$  is the loss incurred by a sample  $\mathbf{z}$  with respect to a parameter  $\theta$ . As the true distribution  $\mathcal{D}$  is unknown, a frequently used approach in learning is to compute a parameter  $\theta$  that minimizes the *empirical loss*  $\frac{1}{n} \cdot \sum_{i=1}^n \ell(\mathbf{z}_i; \theta)$ , which is an unbiased estimator of the expected loss and is purely a function of the available samples. The *generalization error* of a learner  $\mathcal{L}$  is a measure of how close the empirical loss is to the expected loss as a function of the sample size  $n$ .

One technique to reduce the generalization error is to add a *regularization* term to the loss function being minimized [5]. This also ensures that the learned parameter  $\theta$  does not change much with respect to minor changes in the samples being used for learning. Therefore, in a sense, learning algorithms that use regularization can be considered as being stable according to our definition of sensitivity.

Bousquet and Elisseeff [5] defined a notion of stability for learning algorithms in relation to reducing the generalization error. Their stability notion requires that the empirical loss of the learning algorithm does not change much by removing or replacing *any* sample in the input data. In contrast, in our definition of average sensitivity, we consider removing *random* edges from a graph and measure the change in the output solution rather than that in the objective value.

## 1.7 Organization

We show our stable-on-average algorithms for the minimum spanning forest problem, the global minimum cut problem, the minimum  $s$ - $t$  cut problem, and the maximum matching problem problems in Sections 3, 4, 5, and 6, respectively. Our lower bounds on the average sensitivity of algorithms for the global minimum cut problem and the maximum matching problem can also be found in Sections 4, and 6, respectively. We show a linear lower bound for the 2-coloring problem in Section 7. We discuss general properties of average sensitivity in Section 8.

## 2 Preliminaries

For a positive integer  $n$ , let  $[n] = \{1, 2, \dots, n\}$ . Let  $G = (V, E)$  be a graph whose vertex set is  $V$  and edge set is  $E$ . We denote by  $\mathcal{G}$  the (infinite) set consisting of all graphs. We often use the symbols  $n, m, \Delta$  to denote the number of vertices, the number of edges, and the maximum degree of a vertex, respectively, in the input graph. We use  $\text{OPT}(G)$  to denote the optimal value of a

graph  $G$  in the graph problem we are concerned with. We simply write  $\text{OPT}$  when  $G$  is clear from the context. For an edge  $e \in E$ , we denote by  $G - e$  the graph obtained by removing  $e$  from  $G$ . Similarly, for a subset of edges  $F \subseteq E$ , we denote by  $G - F$  the graph obtained by removing every edge in  $F$  from  $G$ . For a subset of edges  $F \subseteq E$ , let  $V(F)$  denote the set of vertices incident to an edge in  $F$ . For a positive integer  $k \leq |E|$ , we use the notation  $\binom{E}{k}$  to denote the set of all subsets of  $E$  of cardinality  $k$ . For a subset of vertices  $S$ , let  $G[S]$  be the subgraph of  $G$  induced by  $S$ . We denote by  $\mathbb{R}_+$  the set of non-negative real numbers. For vectors  $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ , we use  $\langle \mathbf{x}, \mathbf{y} \rangle$  to denote the inner product of  $\mathbf{x}$  and  $\mathbf{y}$ .

## 2.1 Exponential Mechanism

The *exponential mechanism* [35] is an algorithm that, given a vector  $\mathbf{x} \in \mathbb{R}^n$  and a real number  $\eta > 0$ , returns an index  $i \in [n]$  with probability proportional to  $e^{-\eta \mathbf{x}(i)}$ . Just as the exponential mechanism is useful to design differentially private algorithms, it is also useful to design stable-on-average algorithms. Lemma 2.1 formalizes this statement.

**Lemma 2.1.** *Let  $\eta > 0$  and let  $\mathcal{A}$  be the algorithm that, given a vector  $\mathbf{x} \in \mathbb{R}^n$ , applies the exponential mechanism to  $\mathbf{x}$  and  $\eta$ . Then for any  $t > 0$ , we have*

$$\Pr_{i \sim \mathcal{A}(\mathbf{x})} \left[ \mathbf{x}(i) \geq \text{OPT} + \frac{\log n}{\eta} + \frac{t}{\eta} \right] \leq e^{-t},$$

where  $\text{OPT} = \min_{i \in [n]} \mathbf{x}(i)$ . Moreover, for all  $\mathbf{x}' \in \mathbb{R}^n$ , we have

$$d_{\text{TV}}(\mathcal{A}(\mathbf{x}), \mathcal{A}(\mathbf{x}')) = O(\eta \cdot \|\mathbf{x} - \mathbf{x}'\|_1).$$

The proof of Lemma 2.1 is deferred to Appendix A. By setting  $\eta = \log n / \epsilon$  and replacing  $t$  with  $t \log n$ , we get the following:

**Lemma 2.2.** *Let  $\epsilon > 0$ . There exists an algorithm  $\mathcal{A}_\epsilon$  such that, given a vector  $\mathbf{x} \in \mathbb{R}^n$  outputs  $i \in [n]$  such that*

$$\Pr_{i \sim \mathcal{A}_\epsilon(\mathbf{x})} [\mathbf{x}(i) \geq \text{OPT} + \epsilon(1 + t)] \leq n^{-t},$$

for any  $t > 0$ , where  $\text{OPT} = \min_{i \in [n]} \mathbf{x}(i)$ . Moreover, for all  $\mathbf{x}' \in \mathbb{R}^n$ , we have

$$d_{\text{TV}}(\mathcal{A}_\epsilon(\mathbf{x}), \mathcal{A}_\epsilon(\mathbf{x}')) = O\left(\|\mathbf{x} - \mathbf{x}'\|_1 \cdot \frac{\log n}{\epsilon}\right).$$

## 3 Warm Up: Minimum Spanning Forest

To get intuition about average sensitivity of algorithms, we start with the *minimum spanning forest problem*. In this problem, we are given a weighted graph  $G = (V, E, w)$ , where  $w : E \rightarrow \mathbb{R}$  is a weight function on edges, and we want to find a forest of the minimum total weight including all the vertices.

Recall that Kruskal's algorithm [24] works as follows: Iterate over edges in the order of increasing weights, where we break ties arbitrarily. At each iteration, add the current edge to the solution if it does not form a cycle with the edges already added. The following theorem states that this simple and deterministic algorithm is stable-on-average.

**Theorem 3.1.** *The average sensitivity of Kruskal’s algorithm is  $O(n/m)$ .*

*Proof.* Let  $G = (V, E)$  be the input graph and  $T$  be the spanning forest obtained by running Kruskal’s algorithm on  $G$ , where  $T$  is represented by its set of edges. We consider how the output changes when we remove an edge  $e \in E$  from  $G$ .

If the edge  $e$  does not belong to  $T$ , clearly the output of Kruskal’s algorithm on  $G - e$  is also  $T$ .

Suppose that the edge  $e$  belongs to  $T$ . Let  $T_1$  and  $T_2$  be the two trees rooted at the endpoints of  $e$  obtained by removing  $e$  from  $T$ . If  $G - e$  is not connected, that is,  $e$  is a bridge in  $G$ , then Kruskal’s algorithm outputs  $T_1 \cup T_2$  on  $G - e$ . If  $G - e$  is connected, then let  $e'$  be the first edge considered by Kruskal’s algorithm among all the edges connecting  $G[V(T_1)]$  and  $G[V(T_2)]$ , where  $V(T_i)$  is the vertex set of  $T_i$  for  $i \in [2]$ . Then, Kruskal’s algorithm outputs  $T_1 \cup T_2 \cup \{e'\}$  on  $G - e$ . It follows that the Hamming distance between  $T$  and the output of the algorithm on  $G - e$  is at most 2.

Therefore, the average sensitivity of Kruskal’s algorithm is at most

$$\frac{m - |T|}{m} \cdot 0 + \frac{|T|}{m} \cdot 2 = O\left(\frac{n}{m}\right). \quad \square$$

Indeed, it is not hard to show a matching lower bound.

**Theorem 3.2.** *The average sensitivity of a (possibly randomized) algorithm for the minimum spanning forest problem is  $\Omega(n/m)$ .*

*Proof.* Let  $\mathcal{A}$  be an algorithm for the minimum spanning forest problem, and let  $G = (V, E)$  be a connected graph with  $n$  vertices and  $m$  edges.

For each  $e \in E$ , let  $\mu_e$  be the probability distribution over  $\mathcal{F}(G) \times \mathcal{F}(G - e)$  such that  $d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}(G - e)) = \mathbb{E}_{(F, F_e) \sim \mu_e} |F \Delta F_e|$ , where  $\mathcal{F}(G)$  is the set of spanning forests of  $G$ . Note that the marginal distribution of  $\mu_e$  on the first coordinate is identical for all  $e \in E$ . Let  $\mu_{e, F}$  be the distribution over  $\mathcal{F}(G - e)$  defined as (the second coordinate of) the distribution  $\mu_e$  conditioned on the first coordinate being  $F$ . Then, we have

$$\begin{aligned} \mathbb{E}_{e \sim E} [d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}(G - e))] &= \mathbb{E}_{e \sim E} \mathbb{E}_{(F, F_e) \sim \mu_e} |F \Delta F_e| = \mathbb{E}_{F \sim \mathcal{F}(G)} \mathbb{E}_{e \sim E} \mathbb{E}_{F_e \sim \mu_{e, F}} |F \Delta F_e| \\ &= \mathbb{E}_{F \sim \mathcal{F}(G)} \left[ \frac{1}{m} \sum_{e \in E} \mathbb{E}_{F_e \sim \mu_{e, F}} |F \Delta F_e| \right] \geq \mathbb{E}_{F \sim \mathcal{F}(G)} \left[ \frac{1}{m} \sum_{e \in F} \mathbb{E}_{F_e \sim \mu_{e, F}} |F \Delta F_e| \right] \geq \mathbb{E}_{F \sim \mathcal{F}(G)} \left[ \frac{1}{m} \sum_{e \in F} \mathbb{E}_{F_e \sim \mu_{e, F}} 1 \right] \\ &= \frac{n - 1}{m} = \Omega\left(\frac{n}{m}\right), \end{aligned}$$

where in the second inequality we used the fact that  $e \in F$  and  $e \notin F_e$ . □

In Appendix B, we show that Prim’s algorithm, another classical algorithm for the minimum spanning forest problem, has average sensitivity  $\Omega(m)$  for a certain natural tie breaking rule. We mention that this lower bound holds even for unweighted graphs.

## 4 Global Minimum Cut

For a graph  $G = (V, E)$  and a vertex set  $S \subseteq V$ , we define  $\text{cost}(G, S)$  to be the number of edges in  $E$  that cross the cut  $(S, V \setminus S)$ . Then in the *global minimum cut problem*, given a graph  $G = (V, E)$ , we want to compute a vertex set  $\emptyset \subsetneq S \subsetneq V$  that minimizes  $\text{cost}(G, S)$ . In this section, we discuss upper and lower bounds on the average sensitivity for the global minimum cut problem.

## 4.1 Upper bound

In this section, we show the following.

**Theorem 4.1.** *For  $\varepsilon > 0$ , there exists a polynomial time algorithm for the global minimum cut problem with approximation ratio  $2 + \varepsilon$  and average sensitivity  $n^{O(1/\varepsilon \text{OPT})}$ .*

Let  $\text{OPT}$  be the minimum size of a cut in  $G$ . Our algorithm enumerates cuts of small size and then output a vertex set  $S$  with probability  $\exp(-\alpha \cdot \text{cost}(G, S))$  for a suitable  $\alpha$ . See Algorithm 1 for details.

---

**Algorithm 1:** STABLE ALGORITHM FOR GLOBAL MINIMUM CUT

---

- Input:** undirected graph  $G = (V, E)$ ,  $\varepsilon > 0$
- 1 Compute the value  $\text{OPT}$ ;
  - 2 Let  $\alpha \leftarrow \frac{(2+1/\varepsilon) \log n}{\text{OPT}}$  denote a parameter;
  - 3 Enumerate all cuts of size at most  $(2 + 7\varepsilon)\text{OPT} + 2\varepsilon$ ;
  - 4 Sample a vertex set  $S$  (from among the cuts enumerated) with probability proportional to  $\exp(-\alpha \cdot \text{cost}(G, S))$ ;
  - 5 **return**  $S$ .
- 

The approximation ratio of the Algorithm 1 is  $2 + 9\varepsilon$ : It clearly holds when  $\text{OPT} \geq 1$ , and it also holds when  $\text{OPT} = 0$  because we only output a cut of size zero (for  $\varepsilon < 1/2$ ). The following theorem due to Karger [19] directly implies that it runs in time polynomial in the input size for any constant  $\varepsilon > 0$ .

**Theorem 4.2** ([19]). *Given a graph  $G$  on  $n$  vertices with the minimum cut size  $c$  and a parameter  $\alpha \geq 1$ , the number of cuts of size at most  $\alpha \cdot c$  is at most  $n^{2\alpha}$  and can be enumerated in time polynomial (in  $n$ ) per cut.*

We now show that Algorithm 1 is stable-on-average.

**Lemma 4.3.** *The average sensitivity of Algorithm 1 is at most*

$$\beta(G) = \frac{n}{m} \cdot n^{(2+1/\varepsilon)/\text{OPT}} \cdot ((2 + 7\varepsilon)\text{OPT} + 2\varepsilon) + o(1).$$

As we have  $\text{OPT} \leq 2m/n$ , the average sensitivity can be bounded by  $n^{O(1/\varepsilon \text{OPT})}$ , and Theorem 4.1 follows by replacing  $\varepsilon$  with  $\varepsilon/9$ .

*Proof.* If  $\text{OPT} = 0$ , then the claim trivially holds because the right hand size is infinity. Hence in what follows, we assume  $\text{OPT} \geq 1$ .

Let  $\mathcal{A}$  denote Algorithm 1. Consider an (inefficient) algorithm  $\mathcal{A}'$  that on input  $G$ , outputs a cut  $S \subseteq V$  (from among all the cuts in  $G$ ) with probability proportional to  $\exp(-\alpha \cdot \text{cost}(G, S))$ . For a graph  $G = (V, E)$ , let  $\mathcal{A}(G)$  and  $\mathcal{A}'(G)$  denote the output distribution of algorithms  $\mathcal{A}$  and  $\mathcal{A}'$  on input  $G$ , respectively. For  $G = (V, E)$  and  $S \subseteq V$ , let  $p_G(S)$  and  $p'_G(S)$  be shorthands for the probabilities that  $S$  is output on input  $G$  by algorithms  $\mathcal{A}$  and  $\mathcal{A}'$ , respectively.

We first bound the earth mover's distance between  $\mathcal{A}(G)$  and  $\mathcal{A}'(G)$  for a graph  $G = (V, E)$ . To this end, we define

$$Z = \sum_{S \subseteq V: \text{cost}(G, S) \leq \text{OPT} + b} \exp(-\alpha \cdot \text{cost}(G, S)), \text{ and } Z' = \sum_{S \subseteq V} \exp(-\alpha \cdot \text{cost}(G, S))$$

where  $b = (1 + 7\varepsilon)\text{OPT} + 2\varepsilon$ . Note that  $Z \leq Z'$  and the quantity  $\frac{Z' - Z}{Z'}$  is the total probability mass assigned by algorithm  $\mathcal{A}'$  to cuts  $S \subseteq V$  such that  $\text{cost}(G, S) > \text{OPT} + b$ .

Now, we start with  $\mathcal{A}'(G)$ . For each  $S \subseteq V$  such that  $\text{cost}(G, S) \leq \text{OPT} + b$ , keep at least  $\frac{Z}{Z'} \cdot p'_G(S)$  mass with a cost of 0 and move a mass of at most  $p'_G(S) - \frac{Z}{Z'} \cdot p'_G(S)$  at a cost of  $n \cdot (p'_G(S) - \frac{Z}{Z'} \cdot p'_G(S))$ . For each  $S \subseteq V$  such that  $\text{cost}(G, S) > \text{OPT} + b$ , we move a mass of  $p'_G(S)$  at a cost of  $n \cdot p'_G(S)$ . The total cost of moving masses is then equal to:

$$\begin{aligned} d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}'(G)) &\leq n \cdot \sum_{S \subseteq V: \text{cost}(G, S) \leq \text{OPT} + b} p'_G(S) \left(1 - \frac{Z}{Z'}\right) + n \cdot \sum_{S \subseteq V: \text{cost}(G, S) > \text{OPT} + b} p'_G(S) \\ &= \frac{n(Z' - Z)}{Z'} \left( \sum_{S \subseteq V: \text{cost}(G, S) \leq \text{OPT} + b} p'_G(S) + 1 \right) \\ &\leq \frac{2n(Z' - Z)}{Z'}. \end{aligned}$$

Let  $n_t$  stand for the number of cuts of cost at most  $\text{OPT} + t$  in  $G$ . By Karger's theorem (Theorem 4.2), we have that  $n_t \leq n^{2+2t/\text{OPT}}$ . Then, we have

$$\begin{aligned} \frac{Z' - Z}{Z'} &\leq \sum_{t > b} \exp(-\alpha t) \cdot (n_t - n_{t-1}) \leq (\exp(\alpha) - 1) \cdot \sum_{t > b} \exp(-\alpha t) n_t \\ &\leq (\exp(\alpha) - 1) n^2 \cdot \sum_{t > b} n^{2t/\text{OPT}} \cdot \exp(-\alpha t) \\ &\leq (\exp(\alpha) - 1) n^2 \cdot \sum_{t > b} n^{-t/\varepsilon \text{OPT}} \leq (\exp(\alpha) - 1) n^2 \cdot \frac{n^{-(b+1)/\varepsilon \text{OPT}}}{1 - n^{-1/\varepsilon \text{OPT}}} \\ &= \left( n^{(2+1/\varepsilon)/\text{OPT}} - 1 \right) \cdot \left( 1 + \frac{1}{n^{1/\varepsilon \text{OPT}} - 1} \right) \cdot \frac{n^2}{n^{(b+1)/\varepsilon \text{OPT}}} \\ &\leq n^{(2+1/\varepsilon)/\text{OPT}} \cdot \left( 1 + \frac{\varepsilon n}{\log n} \right) \cdot \frac{n^2}{n^{(b+1)/\varepsilon \text{OPT}}} \\ &= O\left( \frac{\varepsilon n^{3+(2+1/\varepsilon)/\text{OPT}}}{n^{(b+1)/\varepsilon \text{OPT}}} \right) = O\left( \frac{\varepsilon}{n^{4+1/\varepsilon}} \right). \end{aligned}$$

The last inequality above follows from our choice of  $b$ . Therefore, the earth mover's distance between  $\mathcal{A}(G)$  and  $\mathcal{A}'(G)$  is  $d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}'(G)) \leq O\left(\frac{\varepsilon}{n^{3+1/\varepsilon}}\right)$ .

In addition, we can bound the expected size of the cut output by  $\mathcal{A}'$  on  $G$  as follows. The total probability mass assigned by algorithm  $\mathcal{A}'$  to cuts of size larger than  $\text{OPT} + b$  is equal to  $\frac{Z' - Z}{Z'} = O\left(\frac{\varepsilon}{n^{4+1/\varepsilon}}\right)$ . Hence, the expected size of the cut output by  $\mathcal{A}'$  on  $G$  is at most  $\text{OPT} + b + m \cdot O\left(\frac{\varepsilon}{n^{4+1/\varepsilon}}\right) = (2 + 7\varepsilon)\text{OPT} + 2\varepsilon + O\left(\frac{\varepsilon m}{n^{4+1/\varepsilon}}\right)$ .

We now bound the earth mover's distance between  $\mathcal{A}'(G)$  and  $\mathcal{A}'(G - e)$  for an arbitrary edge  $e \in E$ . Let  $Z'_e$  denote the quantity  $\sum_{S \subseteq V} \exp(-\alpha \cdot \text{cost}(G - e, S))$ . Since the cost of every cut in  $G - e$  is at most the cost of the same cut in  $G$ , we have that  $Z' \leq Z'_e$  and therefore,

$$p'_G(S) = \frac{\exp(-\alpha \cdot \text{cost}(G, S))}{Z'} \leq \frac{\exp(\alpha \cdot \text{cost}(G - e, S))}{Z'_e} \cdot \frac{Z'_e}{Z'} = p'_{G-e}(S) \cdot \frac{Z'_e}{Z'}.$$



We transform  $\mathcal{A}'(G)$  into  $\mathcal{A}'(G - e)$  as follows. For each  $S \subseteq V$ , we leave a probability mass of at most  $p'_{G-e}(S)$  at  $S$  with zero cost and move a mass of  $\max\{0, p'_G(S) - p'_{G-e}(S)\}$  to any other point at a cost of at most  $n \cdot \max\{0, p'_G(S) - p'_{G-e}(S)\} \leq n \cdot \left(\frac{Z'_e}{Z'} - 1\right) \cdot p'_G(S)$ . Hence,

$$d_{\text{EM}}(\mathcal{A}'(G), \mathcal{A}'(G - e)) \leq n \cdot \left(\frac{Z'_e}{Z'} - 1\right) \cdot \sum_{S \subseteq V} p'_G(S) = n \cdot \left(\frac{Z'_e}{Z'} - 1\right).$$

By the triangle inequality, the earth mover's distance between  $\mathcal{A}(G)$  and  $\mathcal{A}(G - e)$  can be bounded as

$$\begin{aligned} d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}(G - e)) &\leq d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}'(G)) + d_{\text{EM}}(\mathcal{A}'(G), \mathcal{A}'(G - e)) + d_{\text{EM}}(\mathcal{A}'(G - e), \mathcal{A}(G - e)) \\ &\leq n \cdot \left(\frac{Z'_e}{Z'} - 1\right) + O\left(\frac{2\varepsilon}{n^{2+1/\varepsilon}}\right). \end{aligned}$$

Hence, the average sensitivity of  $\mathcal{A}$  is bounded as:

$$\begin{aligned} \beta(G) &= \mathbb{E}_{e \sim E} d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}(G - e)) \leq O\left(\frac{2\varepsilon}{n^{3+1/\varepsilon}}\right) + n \cdot \mathbb{E}_{e \sim E} \left(\frac{Z'_e}{Z'} - 1\right) \\ &= O\left(\frac{2\varepsilon}{n^{3+1/\varepsilon}}\right) + \frac{n}{mZ'} \sum_{e \in E} (Z'_e - Z') \\ &= O\left(\frac{2\varepsilon}{n^{3+1/\varepsilon}}\right) + \frac{n}{mZ'} \sum_{e \in E} \sum_{S \subseteq V: e \text{ crosses } S} \exp(-\alpha \cdot \text{cost}(G - e, S)) - \exp(-\alpha \cdot \text{cost}(G, S)) \\ &= O\left(\frac{2\varepsilon}{n^{3+1/\varepsilon}}\right) + \frac{n(\exp(\alpha) - 1)}{mZ'} \sum_{e \in E} \sum_{S \subseteq V: e \text{ crosses } S} \exp(-\alpha \cdot \text{cost}(G, S)) \\ &= O\left(\frac{2\varepsilon}{n^{3+1/\varepsilon}}\right) + \frac{n(\exp(\alpha) - 1)}{m} \sum_{S \subseteq V} \text{cost}(G, S) \cdot \frac{\exp(-\alpha \cdot \text{cost}(G, S))}{Z'}. \end{aligned}$$

The summation in the second term above is equal to the expected size of the cut output by algorithm  $\mathcal{A}'$  on input  $G$ . We argued that it is at most  $(2 + 7\varepsilon)\text{OPT} + 2\varepsilon + O\left(\frac{\varepsilon m}{n^{4+1/\varepsilon}}\right)$ . Hence, the average sensitivity of  $\mathcal{A}$  is at most

$$\begin{aligned} &\frac{n}{m} \cdot n^{(2+1/\varepsilon)/\text{OPT}} \cdot ((2 + 7\varepsilon)\text{OPT} + 2\varepsilon) + O\left(\frac{\varepsilon n^{(2+1/\varepsilon)/\text{OPT}} + 2}{n^{3+1/\varepsilon}}\right) \\ &= \frac{n}{m} \cdot n^{(2+1/\varepsilon)/\text{OPT}} \cdot ((2 + 7\varepsilon)\text{OPT} + 2\varepsilon) + o(1) \end{aligned}$$

as  $\text{OPT} \geq 1$ . □

## 4.2 Lower bound

In this section, we show that the average sensitivity of the algorithm given in Section 4.1 is almost tight. Specifically, we show the following.

**Theorem 4.4.** *Any algorithm for the global minimum cut problem with no additive error (and possibly an arbitrary large multiplicative error) has average sensitivity  $\Omega(n^{1/\text{OPT}}/\text{OPT}^2)$  if  $\text{OPT} = o(\sqrt{n})$ .*

*Proof.* We first show a lower bound for the case  $\text{OPT} = 1$ . Let  $\mathcal{A}$  be an arbitrary algorithm for the global minimum cut problem with no additive error and let  $G = ([n+1], E)$  be a path on  $n+1$  vertices, where  $E = \{(i, i+1) : i \in [n]\}$ . Note that for any  $i \in [n]$ , the graph  $G - (i, i+1)$  is disconnected and  $\mathcal{A}$  must output a vertex set  $[i]$  or  $[n+1] \setminus [i]$ . For a vertex set  $S \subseteq [n+1]$ , let  $p_S$  be the probability that  $\mathcal{A}$  on  $G$  outputs  $S$ . Then, the average sensitivity of  $\mathcal{A}$  on  $G$  is

$$\mathbb{E}_{e \sim E} [d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}(G - e))] = \frac{1}{n} \sum_{i \in [n]} \sum_{S \subseteq [n+1]} p_S \cdot \min\{d_{\text{Ham}}(S, [i]), d_{\text{Ham}}(S, [n+1] \setminus [i])\}. \quad (5)$$

Note that if two sets  $S, T \subseteq [n+1]$  satisfy  $|S| \leq |T| - n/10$  or  $|S| \geq |T| + n/10$ , then  $d_{\text{Ham}}(S, T) \geq n/10$  holds. Hence, we have  $d_{\text{Ham}}(S, [i]) \geq n/10$  for at least a  $4/5$ -fraction of  $i \in [n]$ . Similarly, we have  $d_{\text{Ham}}(S, [n+1] \setminus [i]) \geq n/10$  for at least a  $4/5$ -fraction of  $i \in [n]$ . It follows that we have  $\min\{d_{\text{Ham}}(S, [i]), d_{\text{Ham}}(S, [n+1] \setminus [i])\} \geq n/10$  for at least a  $3/5$ -fraction of  $i \in [n]$ . Then, we have

$$(5) \geq \frac{3}{5} \cdot \frac{n}{10} \cdot \sum_{S \subseteq [n]} p_S = \frac{3n}{50} = \Omega(n). \quad (6)$$

We now consider the case  $t := \text{OPT} \geq 2$ . Consider a multigraph  $G_t = ([n+1], E_t)$ , where  $E_t$  contains  $t$  copies of the edge  $(i, i+1)$  for every  $i \in [n]$ . For  $k = (tn)^{1-1/t}$ , the  $k$ -average sensitivity of  $\mathcal{A}$  on  $G$  without replacement is

$$\begin{aligned} & \mathbb{E}_{\{e_1, \dots, e_k\} \sim \binom{E_t}{k}} [d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}(G - \{e_1, \dots, e_k\}))] \\ & \geq \mathbb{E}_{\{e_1, \dots, e_k\} \sim \binom{E_t}{k}} [d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}(G - \{e_1, \dots, e_k\})) \mid \mathcal{A}(G - \{e_1, \dots, e_k\}) \text{ has two components}] \times \\ & \quad \Pr_{\{e_1, \dots, e_k\} \sim \binom{E_t}{k}} [\mathcal{A}(G - \{e_1, \dots, e_k\}) \text{ has two components}]. \end{aligned} \quad (7)$$

The first factor of (7) is exactly equal to (5), which is  $\Omega(n)$  by (6). Now we bound the second factor. For every  $i \in [n]$ , the probability that we cut all the edges between  $i$ -th and  $(i+1)$ -th vertices is  $\binom{k}{t} / \binom{tn}{t}$  from the property of the hypergeometric distribution. For every distinct  $i, j \in [n]$ , the probability that we cut all the edges between  $i$ -th and  $(i+1)$ -th vertices and all the edges between  $j$ -th and  $(j+1)$ -th vertices is  $\binom{k}{2t} / \binom{tn}{2t}$ . By the inclusion-exclusion principle, the probability that  $G - \{e_1, \dots, e_k\}$  has exactly two components is at least

$$\begin{aligned} & n \frac{\binom{k}{t}}{\binom{tn}{t}} - \binom{n}{2} \frac{\binom{k}{2t}}{\binom{tn}{2t}} \geq n \left(1 - \frac{t^2}{k}\right) \left(\frac{k}{tn}\right)^t - \binom{n}{2} \frac{n}{n-4t} \left(\frac{k}{tn}\right)^{2t} \quad (\text{By } \left(1 - \frac{k^2}{n}\right) \frac{n^k}{k!} \leq \binom{n}{k} \leq \frac{n^k}{k!}) \\ & = n \left(1 - \frac{t^2}{k}\right) \frac{1}{tn} - \binom{n}{2} \frac{n}{n-4t} \frac{1}{t^2 n^2} \geq \frac{3}{4t} - \frac{1}{t^2} \geq \frac{1}{4t} = \Omega\left(\frac{1}{t}\right), \end{aligned}$$

where we used the fact that  $2 \leq t = o(\sqrt{n})$ . Hence, we have (7) =  $\Omega(n/t)$ . By Theorem 1.5, the average sensitivity  $\beta$  of  $\mathcal{A}$  on  $G$  must satisfy

$$\beta \cdot (tn)^{1-1/t} \geq \Omega\left(\frac{n}{t}\right),$$

which implies  $\beta \geq \Omega(n^{1/t}/t^2)$ . □

The proof of the following theorem is implicit in the first part of the proof of Theorem 4.4.

**Theorem 4.5.** *Any algorithm that exactly outputs the global minimum cut has average sensitivity  $\Omega(n)$ .*

## 5 Minimum $s$ - $t$ Cut

In this section, we design a stable-on-average algorithm for the minimum  $s$ - $t$  cut problem. We say that a pair  $\{u, v\} \in \binom{V}{2}$  is *cut by*  $S \subseteq V$  if  $u \in S$  and  $v \in V \setminus S$ , or vice versa. The *cut size* of a vertex set  $S \subseteq V$  in a graph  $G = (V, E)$  is the number of edges  $e \in E$  cut by  $S$ . In the minimum  $s$ - $t$  cut problem, given a graph  $G = (V, E)$  and two vertices  $s, t \in V$ , we want to find a *minimum  $s$ - $t$  cut*, that is, a vertex set  $S$  with  $s \in S$  and  $t \notin S$  that has the minimum cut size. We show the following.

**Theorem 5.1.** *There exists a polynomial time algorithm for the minimum  $s$ - $t$  cut problem with additive error  $O(n^{2/3})$  and average sensitivity  $O(n^{2/3})$ .*

In Section 5.1, we describe our LP relaxation for the minimum  $s$ - $t$  cut problem and introduce a notion of average sensitivity for algorithms that solve the LP. We then provide a stable-on-average LP solver in Section 5.2 and discuss a rounding procedure in Section 5.3. We prove Theorem 5.1 in Section 5.4.

### 5.1 LP Relaxation and Average Sensitivity

Our algorithm is based on an LP relaxation for the minimum  $s$ - $t$  cut problem. For each pair of vertices  $\{u, v\} \in \binom{V}{2}$ , we introduce a variable  $\mathbf{d}(\{u, v\})$ , which we regard as a distance between  $u$  and  $v$ . Roughly speaking,  $\mathbf{d}(\{u, v\})$  is 1 if  $u$  and  $v$  are on different sides of an  $s$ - $t$  cut.  $\mathbf{d}(\{u, v\})$  is 0 otherwise. For notational simplicity, we often write  $\mathbf{d}(u, v)$  to denote  $\mathbf{d}(\{u, v\})$ . Intuitively, the distance between  $s$  and  $t$  should be at least one, and the distance  $\mathbf{d}(\cdot, \cdot)$  should satisfy the triangle inequality. Our LP relaxation is the following.

$$\begin{aligned}
 & \text{minimize} && \text{LP}_G(\mathbf{d}) := \sum_{\{u, v\} \in E} \mathbf{d}(u, v), \\
 & \text{subject to} && \mathbf{d}(s, t) = 1, \\
 & && \mathbf{d}(u, v) \leq \mathbf{d}(u, w) + \mathbf{d}(w, v), \quad \forall \{u, v, w\} \in \binom{V}{3} \\
 & && 0 \leq \mathbf{d}(u, v) \leq 1. \quad \forall \{u, v\} \in \binom{V}{2}
 \end{aligned} \tag{8}$$

It is easy to check that LP (8) is indeed a relaxation for the minimum  $s$ - $t$  cut problem. Let  $S \subseteq V$  be an  $s$ - $t$  cut. Then for each  $\{u, v\} \in \binom{V}{2}$ , we set  $\mathbf{d}(u, v) = 1$  if  $\{u, v\}$  is cut by  $S$ , and set  $\mathbf{d}(u, v) = 0$  otherwise. It is clear that  $\sum_{\{u, v\} \in E} \mathbf{d}(u, v) = |\{\{u, v\} \in E : \{u, v\} \text{ is cut by } S\}|$  is the cut size of  $S$  and that  $\mathbf{d}$  satisfies all the constraints.

Now, we introduce a notion of the average sensitivity of an algorithm for solving (8). First, for a vertex  $s \in V$ , we define a norm  $\|\cdot\|_s$  as  $\|\mathbf{d}\|_s := \sqrt{\sum_{v \in V} \mathbf{d}(s, v)^2}$ . Then, we define the *average sensitivity* of a (deterministic) algorithm  $\mathcal{A}$  for solving LP (8) as

$$\mathbb{E}_{e \sim E} \|\mathcal{A}(G) - \mathcal{A}(G - e)\|_s.$$

We use the norm  $\|\cdot\|_s$  instead of the standard  $\ell_2$  norm because our rounding procedure uses only  $\mathbf{d}(s, v)$  ( $v \in V$ ) (see Section 5.3 for the details of the rounding procedure), and the former norm gives a better approximation guarantee than the latter.

## 5.2 Stable-on-Average LP Solver

In this section, we give a stable-on-average solver for LP (8).

**Theorem 5.2.** *For any  $\eta > 0$ , there exists a polynomial-time algorithm for solving LP (8) with average sensitivity  $O(1/\sqrt{\eta})$  such that the output  $\mathbf{d} \in [0, 1]^{\binom{V}{2}}$  satisfies*

$$\sum_{e \in E} \mathbf{d}(e) \leq \text{OPT} + \frac{\eta n}{2},$$

where OPT is the optimal value of LP (8).

Given a parameter  $\eta > 0$ , our algorithm solves the following regularized LP and then returns the optimal solution.

$$\begin{aligned} \text{minimize} \quad & \text{LP}_G^\eta(\mathbf{d}) := \sum_{\{u,v\} \in E} \mathbf{d}(u,v) + \frac{\eta}{2} \|\mathbf{d}\|_s^2 \\ \text{subject to} \quad & \mathbf{d}(s,t) = 1, \\ & \mathbf{d}(u,v) \leq \mathbf{d}(u,w) + \mathbf{d}(w,v) \quad \forall \{u,v,w\} \in \binom{V}{3}, \\ & 0 \leq \mathbf{d}(u,v) \leq 1 \quad \forall \{u,v\} \in \binom{V}{2}, \end{aligned} \tag{9}$$

The only difference from LP (8) is that we have a regularization term  $\eta \|\mathbf{d}\|_s^2/2$  in the objective function. Note that  $\text{LP}_G^\eta$  is  $\eta$ -strongly convex with respect to the norm  $\|\cdot\|_s$ . Theorem 5.2 follows from Lemmas 5.3 and 5.4, given below.

**Lemma 5.3.** *Let  $\mathbf{d}^*, \mathbf{d} \in [0, 1]^{\binom{V}{2}}$  be the optimal solutions to LPs (8) and (9), respectively. Then, we have*

$$\sum_{\{u,v\} \in E} \mathbf{d}(e) \leq \sum_{\{u,v\} \in E} \mathbf{d}^*(e) + \frac{\eta n}{2}.$$

*Proof.* We have

$$\begin{aligned} \sum_{e \in E} \mathbf{d}(e) - \sum_{e \in E} \mathbf{d}^*(e) &= \left( \sum_{e \in E} \mathbf{d}(e) + \frac{\eta}{2} \|\mathbf{d}\|_s^2 \right) - \left( \sum_{e \in E} \mathbf{d}^*(e) + \frac{\eta}{2} \|\mathbf{d}^*\|_s^2 \right) + \frac{\eta}{2} (\|\mathbf{d}^*\|_s^2 - \|\mathbf{d}\|_s^2) \\ &\leq \frac{\eta}{2} (\|\mathbf{d}^*\|_s^2 - \|\mathbf{d}\|_s^2) && \text{(by the optimality of } \mathbf{d} \text{ for LP (9))} \\ &\leq \frac{\eta}{2} \|\mathbf{d}^*\|_s^2 \\ &\leq \frac{\eta n}{2}, && \text{(by } \mathbf{d}^*(s,v) \leq 1 \text{ for any } v \in V) \end{aligned}$$

as desired.  $\square$

**Lemma 5.4.** *Let  $G = (V, E)$  be a graph,  $\{u, v\} \in E$  be an edge, and let  $\mathbf{d}$  and  $\mathbf{d}'$  be the optimal solutions to LPs (9) for  $G$  and  $G' := G - \{u, v\}$ , respectively. Then,  $\|\mathbf{d} - \mathbf{d}'\|_s = O(1/\sqrt{\eta})$ .*

---

**Algorithm 2:** Rounding procedure for the minimum  $s$ - $t$  cut problem

---

**1 Procedure** THRESH( $\mathbf{x}$ )  
     **Input:**  $\mathbf{x} \in [0, 1]^V$   
**2**     Sample  $\tau$  from  $[0, 1]$  uniformly at random;  
**3**     **return** the set  $\{v \in V : \mathbf{x}(v) \geq \tau\}$

---

*Proof.* By the  $\eta$ -strong convexity of  $\text{LP}_G^\eta$  with respect to  $\|\cdot\|_s$ , we have

$$\text{LP}_G^\eta(\mathbf{d}') - \text{LP}_G^\eta(\mathbf{d}) \geq \langle \nabla \text{LP}_G^\eta(\mathbf{d}), \mathbf{d}' - \mathbf{d} \rangle + \eta \|\mathbf{d}' - \mathbf{d}\|_s^2 \geq \eta \|\mathbf{d} - \mathbf{d}'\|_s^2,$$

where the second inequality holds because  $\mathbf{d}$  is an optimal solution. Similarly for  $\text{LP}_{G'}^\eta$ , we have

$$\text{LP}_{G'}^\eta(\mathbf{d}) - \text{LP}_{G'}^\eta(\mathbf{d}') \geq \eta \|\mathbf{d} - \mathbf{d}'\|_s^2.$$

Therefore, we have

$$\begin{aligned} 2\eta \|\mathbf{d} - \mathbf{d}'\|_s^2 &\leq \text{LP}_G^\eta(\mathbf{d}') - \text{LP}_G^\eta(\mathbf{d}) + \text{LP}_{G'}^\eta(\mathbf{d}) - \text{LP}_{G'}^\eta(\mathbf{d}') \\ &= \sum_{e \in E} (\mathbf{d}'(e) - \mathbf{d}(e)) + \sum_{e \in E \setminus \{u, v\}} (\mathbf{d}(e) - \mathbf{d}'(e)) = \mathbf{d}'(u, v) - \mathbf{d}(u, v) \leq 1. \end{aligned}$$

Hence, we have  $\|\mathbf{d} - \mathbf{d}'\|_s \leq O(1/\sqrt{\eta})$ . □

### 5.3 Rounding Procedure

Suppose we have obtained a solution  $\mathbf{d}$  to LP (8). Then, we will round the vector  $\mathbf{z} := (\mathbf{d}(s, v))_{v \in V} \in [0, 1]^V$  using a thresholding procedure THRESH (Algorithm 2), which returns a set of vertices  $v \in V$  such that  $\mathbf{d}(s, v)$  is at least a threshold  $\tau$  sampled from  $[0, 1]$  uniformly at random. We can also interpret THRESH( $\mathbf{z}$ ) as follows. Let  $v_1, \dots, v_n$  be the ordering of  $V$  such that  $\mathbf{z}(v_i) \geq \mathbf{z}(v_{i+1})$  for every  $i \in [n - 1]$ . Then, THRESH( $\mathbf{z}$ ) outputs the set  $\{v_1, \dots, v_i\}$  with probability  $\mathbf{z}(v_i) - \mathbf{z}(v_{i+1})$  for  $i \in [n - 1]$ , where we define  $\mathbf{z}(v_{n+1}) = 0$  for a dummy vertex  $v_{n+1}$ .

First, we analyze the solution quality of THRESH.

**Lemma 5.5.** *We have*

$$\mathbb{E}[\text{The cut size of THRESH}(\mathbf{z})] \leq \sum_{e \in E} \mathbf{d}(e),$$

where  $\mathbf{z} = (\mathbf{d}(s, v))_{v \in V}$ .

*Proof.* For each edge  $e \in E$ , the probability that it is cut by  $S$  is

$$|\mathbf{z}(u) - \mathbf{z}(v)| = |\mathbf{d}(s, u) - \mathbf{d}(s, v)| \leq \mathbf{d}(u, v),$$

where  $u, v \in V$  are the endpoints of  $e$ . The claim follows by the linearity of expectations. □

Now, we bound the average sensitivity of THRESH when only one coordinate differs.

**Lemma 5.6.** Let  $\mathbf{x} \in [0, 1]^V$  and  $\mathbf{x}' \in [0, 1]^V$  be such that

$$\mathbf{x}'(u) = \begin{cases} \mathbf{x}(u) + \Delta & \text{if } u = v, \\ \mathbf{x}(u) & \text{otherwise,} \end{cases}$$

for some  $v \in V$  and  $0 \leq \mathbf{x}(v) \leq 1 - \Delta$ . Then,  $d_{\text{EM}}(\text{THRESH}(\mathbf{x}), \text{THRESH}(\mathbf{x}')) \leq 2\Delta$ .

*Proof.* We can assume  $\Delta \geq 0$  as otherwise we can switch the roles of  $\mathbf{x}$  and  $\mathbf{x}'$ . Starting with the vector  $\mathbf{x}_0 = \mathbf{x}$ , we iteratively construct  $\mathbf{x}_k \in [0, 1]^V$  from  $\mathbf{x}_{k-1}$  as

$$\mathbf{x}_k(u) = \begin{cases} \min \left\{ \mathbf{x}'(v), \min_{w \in V: \mathbf{x}_{k-1}(w) > \mathbf{x}_{k-1}(v)} \{ \mathbf{x}_{k-1}(w) \} \right\} & \text{if } u = v, \\ \mathbf{x}_{k-1}(u) & \text{otherwise.} \end{cases}$$

Let  $\ell$  be the smallest integer such that  $\mathbf{x}_\ell(v) = \mathbf{x}'(v)$ . Note that for every  $k \in [\ell]$ , there is an ordering  $v_1, \dots, v_n$  of  $V$  such that both  $\mathbf{x}_{k-1}(v_i) \geq \mathbf{x}_{k-1}(v_{i+1})$  and  $\mathbf{x}_k(v_i) \geq \mathbf{x}_k(v_{i+1})$  hold for every  $i \in [n-1]$ . Note that  $\sum_{k \in [\ell]} \|\mathbf{x}_k - \mathbf{x}_{k-1}\|_1 = \Delta$ .

Now we show that for each  $k \in [\ell]$  we have  $d_{\text{EM}}(\text{THRESH}(\mathbf{x}_k), \text{THRESH}(\mathbf{x}_{k-1})) \leq 2\|\mathbf{x}_k - \mathbf{x}_{k-1}\|_1$ . Let  $v_1, \dots, v_n$  be an ordering of  $V$  with the property mentioned above, and let  $S = \{v_1, \dots, v_{i-1}\}$  and  $S' = \{v_1, \dots, v_i\}$ , where  $i \in [n]$  is such that  $v_i = v$ . Then, the only difference in the output distributions of  $\text{THRESH}(\mathbf{x}_k)$  and  $\text{THRESH}(\mathbf{x}_{k-1})$  is that the former outputs  $S$  with probability  $\mathbf{x}_k(v_{i-1}) - \mathbf{x}_k(v_i)$  and  $S'$  with probability  $\mathbf{x}_k(v_i) - \mathbf{x}_k(v_{i+1})$  whereas the latter outputs  $S$  with probability  $\mathbf{x}_{k-1}(v_{i-1}) - \mathbf{x}_{k-1}(v_i)$  and  $S'$  with probability  $\mathbf{x}_{k-1}(v_i) - \mathbf{x}_{k-1}(v_{i+1})$ . It follows that

$$\begin{aligned} d_{\text{EM}}(\text{THRESH}(\mathbf{x}_k), \text{THRESH}(\mathbf{x}_{k-1})) &\leq 2(\mathbf{x}_k(v_i) - \mathbf{x}_{k-1}(v_i)) \cdot d_{\text{Ham}}(S, S') \\ &= 2(\mathbf{x}_k(v) - \mathbf{x}_{k-1}(v)) \cdot d_{\text{Ham}}(S, S') = 2\|\mathbf{x}_k - \mathbf{x}_{k-1}\|_1. \end{aligned}$$

Then, we have

$$\begin{aligned} d_{\text{EM}}(\text{THRESH}(\mathbf{x}), \text{THRESH}(\mathbf{x}')) &= d_{\text{EM}}(\text{THRESH}(\mathbf{x}_0), \text{THRESH}(\mathbf{x}_\ell)) \\ &\leq \sum_{k \in [\ell]} d_{\text{EM}}(\text{THRESH}(\mathbf{x}_k), \text{THRESH}(\mathbf{x}_{k-1})) \leq \sum_{k \in [\ell]} 2\|\mathbf{x}_k - \mathbf{x}_{k-1}\|_1 = 2\Delta. \quad \square \end{aligned}$$

**Corollary 5.7.**  $d_{\text{EM}}(\text{THRESH}(\mathbf{x}), \text{THRESH}(\mathbf{x}')) \leq 2\|\mathbf{x} - \mathbf{x}'\|_1$  for  $\mathbf{x}, \mathbf{x}' \in [0, 1]^V$ .

*Proof.* The inequality can be obtained by iteratively applying Lemma 5.6 to each coordinate of the vectors.  $\square$

## 5.4 Putting Things Together

Our algorithm is given in Algorithm 3. It simply computes a solution  $\mathbf{d} \in [0, 1]^{\binom{V}{2}}$  to LP (8) using the algorithm given in Theorem 5.2 with parameter  $\eta = n^{-1/3}$ , and then rounds  $\mathbf{d}$  using the procedure `THRESH`.

---

**Algorithm 3:** Algorithm for the minimum  $s$ - $t$  cut problem

---

1 **Procedure** MINCUT( $G, s, t, \epsilon$ )  
2     **Input:** A graph  $G = (V, E)$ , two vertices  $s, t \in V$ .  
3     Solve LP (8) using the algorithm given by Theorem 5.2 with parameter  $\eta = n^{-1/3}$ , and  
4     let  $\mathbf{d} \in [0, 1]^{\binom{V}{2}}$  be the solution obtained;  
5      $\mathbf{z} \leftarrow (\mathbf{d}(s, v))_{v \in V}$ ;  
6     **return** THRESH( $\mathbf{z}$ ).

---

*Proof of Theorem 5.1.* Let  $\mathcal{A}$  be Algorithm 3. First, we consider the approximation guarantee of  $\mathcal{A}$ . Let  $\mathbf{d}^* \in \mathbb{R}^{\binom{V}{2}}$  be an optimal solution to LP (8). By Theorem 5.2, we get  $\mathbf{d} \in [0, 1]^{\binom{V}{2}}$  satisfying the constraints in LP (8) such that

$$\sum_{e \in E} \mathbf{d}(e) \leq \text{OPT} + \frac{\eta n}{2} = \text{OPT} + O(n^{2/3}).$$

By Lemma 5.5, the expected cut size of the output set is the same.

Now, we consider the average sensitivity of  $\mathcal{A}$ . Let  $G = (V, E)$  be a graph, and let  $\mathbf{d} \in [0, 1]^{\binom{V}{2}}$  be the solution to LP (8) computed by  $\mathcal{A}$  on  $G$ . For each edge  $e \in E$ , let  $\mathbf{d}_e \in [0, 1]^{\binom{V}{2}}$  be the solution to LP (8) computed by  $\mathcal{A}$  on  $G - e$ . Then, we have

$$\begin{aligned} \mathbb{E}_{e \sim E} d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}(G - e)) &\leq \mathbb{E}_{e \sim E} d_{\text{EM}}(\text{THRESH}((\mathbf{d}(s, v))_{v \in V}), \text{THRESH}((\mathbf{d}_e(s, v))_{v \in V})) \\ &\leq 2 \mathbb{E}_{e \sim E} \|(\mathbf{d}(s, v))_{v \in V} - (\mathbf{d}_e(s, v))_{v \in V}\|_1 && \text{(by Corollary 5.7)} \\ &= 2\sqrt{n} \mathbb{E}_{e \sim E} \|\mathbf{d} - \mathbf{d}_e\|_s \\ &= O\left(\sqrt{\frac{n}{\eta}}\right) && \text{(by Theorem 5.2)} \\ &= O(n^{2/3}). \end{aligned} \quad \square$$

## 6 Maximum Matching

A vertex-disjoint set of edges is called a *matching*. In the maximum matching problem, given a graph, we want to find a matching of the maximum size. In this section, we describe several algorithms with low average sensitivity that approximate the maximum matching in a graph.

### 6.1 Lexicographically smallest matching

In this section, we describe an algorithm that computes a maximum matching in a graph with average sensitivity at most  $\text{OPT}^2/m$  and prove Theorem 6.1, where  $\text{OPT}$  is the maximum size of a matching.

First, we define some ordering among vertex pairs. Then, we can naturally define the lexicographical order among matchings by considering a matching as a sorted sequence of vertex pairs. Then, our algorithm simply outputs the lexicographically smallest matching. Note that this can be done in polynomial time using Edmonds' algorithm [9].

---

**Algorithm 4: RANDOMIZED GREEDY ALGORITHM**

---

**Input:** undirected unweighted graph  $G = (V, E)$

- 1 Sample a uniformly random ordering  $\pi$  of edges in  $E$ ;
  - 2 Set  $M \leftarrow \emptyset$ ;
  - 3 Consider edges one by one according to  $\pi$  and add an edge  $(u, v)$  to  $M$  only if both  $u$  and  $v$  are unmatched in  $M$ ;
  - 4 **return**  $M$ .
- 

**Theorem 6.1.** *Let  $\mathcal{A}$  be the algorithm that outputs the lexicographically smallest maximum matching. Then, the average sensitivity of  $\mathcal{A}$  is at most  $\text{OPT}^2/m$ , where  $\text{OPT}$  is the maximum size of a matching.*

*Proof.* For a graph  $G = (V, E)$ , let  $M(G) \subseteq E$  be its lexicographically smallest maximum matching. As long as  $e \notin M$ , we have  $M(G) = M(G - e)$ . Hence, the average sensitivity of the algorithm is at most

$$\frac{\text{OPT}}{m} \cdot \text{OPT} + \left(1 - \frac{\text{OPT}}{m}\right) \cdot 0 = \frac{\text{OPT}^2}{m}. \quad \square$$

**Remark 6.2.** Consider the path graph  $P_n = (\{1, \dots, n\}, E)$ , where  $E = \{(i, i + 1) : i \in [n - 1]\}$ . The average sensitivity of the above algorithm on  $P_n$  is  $\Omega(\frac{\text{OPT}^2}{m})$ . Hence the above analysis of the average sensitivity is tight.

## 6.2 Greedy matching algorithm

In this section, we analyze the average sensitivity of the randomized greedy algorithm (Algorithm 4) that outputs a maximal matching. It is evident that Algorithm 4 runs in polynomial time and that the matching it outputs has size at least  $\frac{1}{2}$  the size of a maximum matching in the input graph.

**Theorem 6.3.** *Algorithm 4 is a  $\frac{1}{2}$ -approximation algorithm for the maximum matching problem and has average sensitivity at most 1.*

*Proof.* For a permutation  $\pi$  of edges in  $E$ , let  $M_\pi(G)$  denote the matching obtained by running Algorithm 4 on a graph  $G$ . Using [6, Theorem 1], we get that for every  $e \in E(G)$ , it holds that  $\mathbb{E}_\pi [\text{Ham}(M_\pi(G), M_\pi(G - e))] \leq 1$ . This implies that the average sensitivity of Algorithm 4 is at most 1.  $\square$

A vertex set  $S \subseteq V$  in a graph  $G = (V, E)$  is called a *vertex cover* if every edge in  $E$  is incident to a vertex in  $S$ . In the *minimum vertex cover problem*, given a graph  $G$ , we want to compute a vertex cover of the minimum size. It is well known that, for any maximal matching  $M$ , the vertex set consisting of all endpoints of edges in  $M$  is a 2-approximate vertex cover. The following theorem is immediate from Theorem 6.3.

**Theorem 6.4.** *There exists a 2-approximation algorithm for the minimum vertex cover problem with average sensitivity at most 2.*



### 6.3 Matching algorithm based on augmenting paths

In this section, we describe a  $(1 - \varepsilon)$ -approximation algorithm for the maximum matching problem with average sensitivity  $\tilde{O}\left(\text{OPT}^{\frac{c}{c+1}} / \varepsilon^{\frac{3c}{c+1}}\right)$  for  $c = O(1/\varepsilon^2)$  in Theorem 6.11. The basic building block is a  $(1 - \varepsilon)$ -approximation algorithm (Algorithm 5) for maximum matching that is based on iteratively augmenting a matching with greedily chosen augmenting paths of increasing lengths. In Theorem 6.5, we show that the average sensitivity of this algorithm is  $\Delta^{O(1/\varepsilon^2)}$ , where  $\Delta$  is the maximum degree of the input graph. We obtain Theorem 6.5 by applying Theorem 1.10 to a result by Yoshida et al. [52].

We then apply Theorem 6.6 to Theorem 6.5 in order to get rid of the dependence of the average sensitivity on the maximum degree and obtain Theorem 6.10. We then combine (using Theorem 1.8, the parallel composition theorem) the algorithm guaranteed by Theorem 6.10 with the algorithm guaranteed by Theorem 6.1 to obtain Theorem 6.11.

#### 6.3.1 Greedy matching algorithm based on augmenting paths

In this section, we present an approximation algorithm that starts with an empty matching and then iteratively improves its size with augmenting paths of increasing lengths. We show that the average sensitivity of this algorithm can be bounded using Theorem 1.10.

---

#### Algorithm 5: GREEDY AUGMENTING PATHS ALGORITHM

---

**Input:** undirected unweighted graph  $G = (V, E)$ , parameter  $\varepsilon \in (0, 1)$

- 1  $M_0 \leftarrow \emptyset$ ;
  - 2 **for**  $i \in \{1, 2, \dots, \lceil \frac{1}{\varepsilon} - 1 \rceil\}$  **do**
  - 3     Let  $A_i$  denote the set of augmenting paths of length  $2i - 1$  for the matching  $M_{i-1}$ ;
  - 4     Let  $A'_i$  denote a maximal set of disjoint paths from  $A_i$ , where  $A'_i$  is made from a random ordering of  $A_i$ ;
  - 5      $M_i \leftarrow M_{i-1} \Delta A'_i$ .
  - 6 **return**  $M_{\lceil \frac{1}{\varepsilon} - 1 \rceil}$ .
- 

**Theorem 6.5.** *Algorithm 5 with parameter  $\varepsilon > 0$  has approximation ratio  $1 - \varepsilon$  and average sensitivity  $\Delta^{O(1/\varepsilon^2)}$ , where  $\Delta$  is the maximum degree of the input graph.*

*Proof.* For all  $k \geq 0$ , it is known that  $|M_k| \geq \frac{k}{k+1} \cdot |M^*|$  [14], where  $M^*$  denotes a maximum matching in  $G$ . Hence, the matching  $M_{\lceil \frac{1}{\varepsilon} - 1 \rceil}$  is a  $(1 - \varepsilon)$ -approximation to  $M^*$ .

Yoshida et al. [52, Theorem 3.7] show that for all  $k \geq 0$ , determining whether a uniformly random edge  $e \sim E$  belongs to  $M_k$  can be done by querying at most  $\Delta^{O(k^2)}$  edges in expectation, where  $\Delta$  is the maximum degree of  $G$ . Applying Theorem 1.10 to this result, we can see that the average sensitivity of Algorithm 5 with parameter  $\varepsilon > 0$  and input  $G$  is  $\Delta^{O(1/\varepsilon^2)}$ , where  $\Delta$  is the maximum degree of  $G$ .  $\square$

#### 6.3.2 Stable-on-average thresholding transformation

In this section, we show a transformation from matching algorithms whose average sensitivity is a function of the maximum degree to matching algorithms whose average sensitivity does not depend

on the maximum degree. This is done by adding to the algorithm, a preprocessing step that removes vertices from the input graph, where the removed vertices have degree at least an appropriate random threshold. Such a transformation helps us to design stable-on-average algorithms for graphs with unbounded degree. Let  $\text{Lap}(\mu, \phi)$  denote the Laplace distribution with a location parameter  $\mu$  and a scale parameter  $\phi$ .

**Theorem 6.6.** *Let  $\mathcal{A}'$  be a randomized algorithm for the maximum matching problem such that the size of the matching output by  $\mathcal{A}'$  on a graph  $G$  is always at least  $a \cdot \text{OPT}$  for some  $a \geq 0$ . In addition, assume that there exists a solution oracle  $\mathcal{O}$  (see Definition 1.9) for  $\mathcal{A}'$  that makes at most  $q(\Delta)$  queries to  $G$  in expectation, where  $\Delta$  is the maximum degree of  $G$ , and the expectation is taken over the random coins of  $\mathcal{A}'$  and edges  $e \in E$ . Let  $\delta > 0$  and  $\tau$  be a non-negative function on graphs. Then, there exists an algorithm  $\mathcal{A}$  for the maximum matching problem with average sensitivity*

$$\beta(G) \leq O\left(\frac{K_G}{\delta(\tau(G) - K_G)} + \exp\left(-\frac{1}{\delta}\right)\right) \cdot \text{OPT} + \mathbb{E}_L \left[ (2L - 2)^2 q(L) \right],$$

where  $L$  is a random variable distributed as  $\text{Lap}(\tau(G), \delta\tau(G))$  and  $K_G = \max_{e \in E(G)} |\tau(G) - \tau(G - e)|$ . Moreover, the expected size of the matching output by  $\mathcal{A}$  is at least

$$a \cdot \text{OPT} - \frac{am}{(1 - \delta \ln(\text{OPT}/2)) \cdot \tau(G)} - a.$$

The following fact will be useful in the proof of Theorem 6.6.

**Proposition 6.7.** *Let  $L$  be a random variable distributed as  $\text{Lap}(\mu, \phi)$ . Then,  $\Pr[L < (1 - \varepsilon)\mu] \leq \exp(-\varepsilon\mu/\phi)/2$ . Similarly,  $\Pr[L > (1 + \varepsilon)\mu] \leq \exp(-\varepsilon\mu/\phi)/2$ .*

*Proof of Theorem 6.6.* The algorithm  $\mathcal{A}$  is given below.

**Algorithm  $\mathcal{A}$ :** On input  $G = (V, E)$ ,

1. Sample a random variable  $L$  according to the distribution  $\text{Lap}(\tau(G), \delta\tau(G))$ .
2. Let  $[G]_L$  be the graph obtained after removing from  $G$  all vertices of degree at least  $L$ .
3. Run  $\mathcal{A}'$  on  $[G]_L$ .

We first bound the average sensitivity of  $\mathcal{A}$ . We can think of  $\mathcal{A}$  as being sequentially composed of two algorithms, where the first algorithm takes in a graph  $G = (V, E)$  and outputs a number  $L \sim \text{Lap}(\tau(G), \delta\tau(G))$ . The second algorithm takes both  $L$  and  $G$  and runs  $\mathcal{A}'$  on  $[G]_L$ .

Let  $L_e$  for  $e \in E$  denote a Laplace random variable distributed as  $\text{Lap}(\tau(G - e), \delta\tau(G - e))$ . Using Theorem 1.6, we get that the average sensitivity of  $\mathcal{A}$  is bounded by

$$\text{OPT} \cdot \mathbb{E}_{e \sim E} [d_{\text{TV}}(L, L_e)] + \mathbb{E}_L \left[ \mathbb{E}_{e \sim E} [d_{\text{EM}}(\mathcal{A}'([G]_L), \mathcal{A}'([G - e]_L))] \right].$$

**Claim 6.8.** *For  $x \in \mathbb{R}$ ,  $\mathbb{E}_{e \sim E} [d_{\text{EM}}(\mathcal{A}'([G]_x), \mathcal{A}'([G - e]_x))] \leq (2x - 2)^2 q(x)$ .*

*Proof.* Fix  $x \in \mathbb{R}$ . In order to bound the term  $\mathbb{E}_{e \sim E} [d_{\text{EM}}(\mathcal{A}'([G]_x), \mathcal{A}'([G - e]_x))]$ , consider the following algorithm  $\mathcal{A}'_x$ . On input  $G = (V, E)$ , the algorithm  $\mathcal{A}'_x$  first removes every vertex of degree at least  $x$  from  $G$  and then runs  $\mathcal{A}'$  on the resulting graph. Hence, the quantity  $\mathbb{E}_{e \sim E} [d_{\text{EM}}(\mathcal{A}'([G]_x), \mathcal{A}'([G - e]_x))]$  denotes the average sensitivity of  $\mathcal{A}'_x$ .

In order to bound the average sensitivity of  $\mathcal{A}'_x$ , construct a solution oracle  $\mathcal{O}_x$  for  $\mathcal{A}'_x$  as follows. The oracle  $\mathcal{O}_x$ , when given access to a graph  $G = (V, E)$  and input  $e$  sampled uniformly at random from  $E$ , does the following. It first checks whether at least one of the endpoints of  $e$  has degree at least  $x$ . If so, it returns that  $e$  does not belong to the solution obtained by running  $\mathcal{A}'_x$  on  $G$ . Otherwise, it runs  $\mathcal{O}$  with access to  $[G]_x$  and  $e$  as input and outputs the answer of  $\mathcal{O}$ .

We can analyze the query complexity of  $\mathcal{O}_x$  as follows. Call an edge  $e \in E$  *alive* if both the endpoints of  $e$  have degree less than  $x$ . Otherwise,  $e$  is *dead*.

The oracle  $\mathcal{O}_x$  can check whether an edge  $e = (u, v)$  is alive or not by querying at most  $2x - 2$  edges incident to  $e$ . In particular  $\mathcal{O}_x$  examines the neighbors of  $u$  and  $v$  one by one, and, as soon  $\mathcal{O}_x$  encounters  $x - 1$  distinct neighbors (excluding  $u$  or  $v$  themselves) for either  $u$  or  $v$ ,  $\mathcal{O}_x$  can declare  $e$  to be a dead edge.

If the edge  $e \in E$  input to  $\mathcal{O}_x$  is a dead edge, therefore,  $\mathcal{O}_x$  queries at most  $2x - 2$  edges and returns that  $e$  cannot be part of a solution to running  $\mathcal{A}'_x$  on  $G$ .

If the input edge  $e \in E$  is alive, then we know that it is a uniformly random alive edge. By the guarantee on  $\mathcal{O}$ , we then know that  $\mathcal{O}$  makes at most  $q(x)$  queries to the alive edges in expectation over the randomness of  $\mathcal{A}'$  and the choice of the input alive edge, since the maximum degree of  $[G]_x$  is at most  $x$ . In order for the oracle  $\mathcal{O}_x$  to simulate oracle access to  $[G]_x$  for the purpose of answering queries made by oracle  $\mathcal{O}$ , for each alive edge  $e$  queried by  $\mathcal{O}$ , the oracle  $\mathcal{O}_x$  has to query each edge incident to  $e$  in  $G$  and determine which among these are alive. Since  $e$  is alive, both endpoints of  $e$  have degrees less than  $x$ . Hence,  $\mathcal{O}_x$  need only check whether at most  $2x - 2$  edges incident to  $e$  are alive or not. This can be done by querying  $(2x - 2)^2$  edges in  $E$  in total.

Combining all of the above, the expected query complexity of  $\mathcal{O}_x$  is at most  $(2x - 2)^2 q(x)$ , where the expectation is taken over the edges of  $e \in E$  and the randomness in  $\mathcal{A}_x$ .

Therefore, by Theorem 1.10, we get that the average sensitivity of algorithm  $\mathcal{A}_x$  is bounded by  $(2x - 2)^2 q(x)$ .  $\square$

We now bound the quantity  $\mathbb{E}_{e \sim E} [d_{\text{TV}}(L, L_e)]$ .

**Claim 6.9.** *For any  $e \in E$ , we have*

$$d_{\text{TV}}(L, L_e) \leq O\left(\frac{K}{\delta(\tau - K)} + \exp\left(-\frac{1}{\delta}\right)\right).$$

*Proof.* Let  $f_L, f_{L_e} : \mathbb{R} \rightarrow \mathbb{R}$  be the probability density functions of the Laplace random variables  $L$  and  $L_e$ , respectively. Let  $\tau = \tau(G)$ ,  $\tau_e = \tau(G_e)$ , and  $K = K_G$ . Then

$$\begin{aligned} \frac{f_L(x)}{f_{L_e}(x)} &= \frac{\frac{1}{2\delta\tau} \exp\left(-\frac{|x-\tau|}{\delta\tau}\right)}{\frac{1}{2\delta\tau_e} \exp\left(-\frac{|x-\tau_e|}{\delta\tau_e}\right)} = \frac{\tau_e}{\tau} \exp\left(\frac{|x-\tau_e|}{\delta\tau_e} - \frac{|x-\tau|}{\delta\tau}\right) \\ &= \left(1 - \frac{\tau - \tau_e}{\tau}\right) \exp\left(\frac{\tau|x-\tau_e| - \tau_e|x-\tau|}{\delta\tau\tau_e}\right). \end{aligned}$$

A direct calculation shows that for  $0 \leq x \leq 2 \max\{\tau, \tau_e\}$ , we have

$$\left(1 - \frac{K}{\tau}\right) \exp\left(\frac{-2K}{\delta(\tau - K)}\right) \leq \frac{f_L(x)}{f_{L_e}(x)} \leq \left(1 + \frac{K}{\tau}\right) \exp\left(\frac{2K}{\delta(\tau - K)}\right).$$

This implies that for all  $S \subseteq [0, 2 \max\{\tau, \tau_e\}]$ ,

$$\left(1 - \frac{K}{\tau}\right) \exp\left(\frac{-2K}{\delta(\tau - K)}\right) - 1 \leq \Pr[L \in S] - \Pr[L_e \in S] \leq \left(1 + \frac{K}{\tau}\right) \exp\left(\frac{2K}{\delta(\tau - K)}\right) - 1.$$

By Proposition 6.7, the probability that  $L$  (and  $L_e$  as well) falls in the range  $[-\infty, 0] \cup [2 \max\{\tau, \tau_e\}, \infty]$  is bounded by  $\exp(-1/\delta)$ . Hence, total variation distance between  $L$  and  $L_e$  is

$$\begin{aligned} & \left(1 + \frac{K}{\tau}\right) \exp\left(\frac{2K}{\delta(\tau - K)}\right) - \left(1 - \frac{K}{\tau}\right) \exp\left(\frac{-2K}{\delta(\tau - K)}\right) + 2 \exp\left(-\frac{1}{\delta}\right) \\ &= \left(1 + \frac{K}{\tau}\right) \left(1 + \frac{2K}{\delta(\tau - K)} + O\left(\frac{K^2}{\delta^2(\tau - K)^2}\right)\right) \\ & \quad - \left(1 - \frac{K}{\tau}\right) \left(1 - \frac{2K}{\delta(\tau - K)} - O\left(\frac{K^2}{\delta^2(\tau - K)^2}\right)\right) + 2 \exp\left(-\frac{1}{\delta}\right) \\ &= \frac{2K}{\tau} + \frac{4K}{\delta(\tau - K)} + O\left(\frac{K^2}{\delta^2(\tau - K)^2}\right) + 2 \exp\left(-\frac{1}{\delta}\right) \\ &\leq \frac{6K}{\delta(\tau - K)} + 2 \exp\left(-\frac{1}{\delta}\right) + O\left(\frac{K^2}{\delta^2(\tau - K)^2}\right). \\ &= O\left(\frac{K}{\delta(\tau - K)} + \exp\left(-\frac{1}{\delta}\right)\right). \end{aligned} \quad \square$$

Therefore, the average sensitivity of  $\mathcal{A}$  is bounded as

$$\begin{aligned} \beta(G) &= \mathbb{E}_{e \sim E} d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}(G - e)) \\ &\leq O\left(\frac{K}{\delta(\tau - K)} + \exp\left(-\frac{1}{\delta}\right)\right) \cdot \text{OPT} + \mathbb{E}_L \left[ (2x - 2)^2 q(x) \right]. \end{aligned}$$

We now bound the approximation guarantee of  $\mathcal{A}$ . By Proposition 6.7,

$$\Pr \left[ L < \left(1 - \delta \ln\left(\frac{\text{OPT}}{2}\right)\right) \cdot \tau(G) \right] \leq \frac{1}{\text{OPT}}.$$

Therefore, with probability at least  $1 - 1/\text{OPT}$ , only those vertices with degree at least  $(1 - \delta \ln(\text{OPT}/2)) \cdot \tau(G)$  are removed from  $G$ . The number of such vertices is at most  $\frac{m}{(1 - \delta \ln(\text{OPT}/2)) \cdot \tau(G)}$ . Therefore, with probability at least  $1 - 1/\text{OPT}$ , the size of a maximum matching in the resulting graph is at most  $\frac{m}{(1 - \delta \ln(\text{OPT}/2)) \cdot \tau(G)}$  smaller than that of  $G$ . With probability at most  $1/\text{OPT}$ , the size of a maximum matching in the resulting instance could be smaller by an additive term of at most  $\text{OPT}$ . Hence, the expected size of a maximum matching in the new instance is at least

$$\text{OPT} - \frac{m}{(1 - \delta \ln(\text{OPT}(G)/2)) \cdot \tau(G)} - 1.$$

The statement on approximation guarantee follows.  $\square$

### 6.3.3 Average sensitivity of the greedy augmenting paths algorithm with thresholding

**Theorem 6.10.** *Let  $\varepsilon \in (0, 1)$  be a parameter. There exists an algorithm with approximation ratio  $1 - \varepsilon$  and average sensitivity*

$$O\left(\frac{\varepsilon}{1-\varepsilon} \log n\right) + \left(\frac{m}{\varepsilon^3 \text{OPT}}\right)^{O(1/\varepsilon^2)}.$$

*Proof.* The algorithm guaranteed by the theorem statement is as follows.

**Algorithm  $\mathcal{A}_\varepsilon$ :** On input  $G = (V, E)$ ,

1. Compute OPT.
2. If  $\text{OPT} \leq \frac{2}{\varepsilon} + 1$  or  $m \leq \frac{1}{3\varepsilon}$ , then output an arbitrary maximum matching.
3. Otherwise, run the algorithm obtained by applying Theorem 6.6 with the setting  $\tau := \tau(G) = \frac{m}{\varepsilon' \text{OPT}}$  and  $\delta := \frac{1}{2 \ln n}$  to Algorithm 5 run with parameter  $\varepsilon'$ , where  $\varepsilon' = \frac{\varepsilon}{3} - \frac{1}{3\text{OPT}}$ .

*Approximation guarantee:* If  $\text{OPT} \leq \frac{1}{\varepsilon} + 1$  or  $m \leq \frac{1}{2\varepsilon}$ , the approximation guarantee is clear. Otherwise, since Algorithm 5 outputs a maximal matching whose size is always at least  $(1 - \varepsilon') \cdot \text{OPT}$ , the size of the matching output by  $\mathcal{A}_\varepsilon$  is at least  $(1 - \varepsilon') \cdot \text{OPT} - \frac{\varepsilon' \cdot (1 - \varepsilon') \cdot \text{OPT}}{1 - \frac{\ln(\text{OPT}/2)}{2 \ln n}} - (1 - \varepsilon')$ , which is at least  $(1 - \varepsilon) \cdot \text{OPT}$  by the setting of  $\varepsilon'$  and the fact that  $\frac{\ln(\text{OPT}/2)}{2 \ln n} \leq \frac{1}{2}$ .

*Average sensitivity:* If  $\text{OPT} \leq \frac{2}{\varepsilon} + 1$  or  $m \leq \frac{1}{3\varepsilon}$ , the average sensitivity of  $\mathcal{A}_\varepsilon$  is bounded by  $O(\frac{1}{\varepsilon})$ , since the size of maximum matching in  $G$  is small and it can decrease only by at most 1 by the removal of an edge.

We now analyze the average sensitivity of  $\mathcal{A}_\varepsilon$  for the case that  $\text{OPT} > \frac{2}{\varepsilon} + 1$  and  $m > \frac{1}{3\varepsilon}$ . Let  $c = O(1/\varepsilon^2)$ . The average sensitivity of the algorithm resulting from applying Theorem 6.6 to Algorithm 5 is bounded as:

$$O\left(\frac{K_G}{\delta(\tau - K_G)} + \exp\left(-\frac{1}{\delta}\right)\right) \cdot \text{OPT} + \int_0^\infty (2x - 2)^2 \cdot x^c \cdot \frac{1}{2\delta\tau} \cdot \exp\left(-\frac{|x - \tau|}{\delta\tau}\right) dx. \quad (10)$$

To obtain the above expression, we used the fact (from [52, Theorem 3.7]) that  $q(x) \leq x^c$  when  $x > 0$  and  $q(x) = 0$  otherwise.

The second term of (10) can be bounded as:

$$\begin{aligned} & \int_0^\infty (2x - 2)^2 x^c \frac{1}{2\delta\tau} \exp\left(-\frac{|x - \tau|}{\delta\tau}\right) dx = 4 \int_\tau^\infty x^{c+2} \frac{1}{\delta\tau} \exp\left(-\frac{x - \tau}{\delta\tau}\right) dx \\ & = \exp\left(\frac{1}{\delta}\right) (\delta\tau)^{c+2} \Gamma\left(c + 3, \frac{1}{\delta}\right) = (\delta\tau)^{c+2} (c + 2)! \sum_{k=0}^{c+2} \frac{(1/\delta)^k}{k!} = \left(\frac{m}{\varepsilon^3 \text{OPT}}\right)^{O(1/\varepsilon^2)} \end{aligned}$$

where  $\Gamma(\cdot, \cdot)$  is the incomplete Gamma function and we have used the fact that  $\Gamma(s + 1, x) = s! \exp(-x) \sum_{k=0}^s x^k / k!$  if  $s$  is a non-negative integer. Moreover, each term in the summation  $\delta^{c+2} \cdot (c + 2)! \sum_{k=0}^{c+2} \frac{(1/\delta)^k}{k!}$  is  $o(1)$ . Hence, the summation is  $O(\frac{1}{\varepsilon^2})$ .

In order to bound the first term of (10), note that

$$K_G = \max_{e \in E} |\tau(G) - \tau(G - e)|$$

$$\begin{aligned}
&= 3 \max_{e \in E} \left| \frac{m}{\varepsilon \text{OPT}(G) - 1} - \frac{m-1}{\varepsilon \text{OPT}(G-e) - 1} \right| \\
&\leq 3 \max \left\{ \frac{m}{\varepsilon \text{OPT}(G) - 1} - \frac{m-1}{\varepsilon \text{OPT}(G) - 1}, \frac{m-1}{\varepsilon(\text{OPT}(G) - 1) - 1} - \frac{m}{\varepsilon \text{OPT}(G) - 1} \right\} \\
&= 3 \max \left\{ \frac{1}{\varepsilon \text{OPT}(G) - 1}, \frac{\varepsilon(m - \text{OPT}(G)) + 1}{(\varepsilon(\text{OPT}(G) - 1) - 1) \cdot (\varepsilon \text{OPT}(G) - 1)} \right\} \\
&= \frac{3}{\varepsilon \text{OPT}(G) - 1} \max \left\{ 1, \frac{\varepsilon(m - \text{OPT}(G)) + 1}{\varepsilon(\text{OPT}(G) - 1) - 1} \right\}.
\end{aligned}$$

The inequality above uses the fact that for numbers  $u, v, w \geq 0$  such that  $u \leq v$  and  $w(v-1) - 1 \geq 0$ , we have that  $\frac{u}{wv-1} \leq \frac{u-1}{w(v-1)-1}$ .

Since  $\frac{\varepsilon(m - \text{OPT}) + 1}{\varepsilon(\text{OPT} - 1) - 1}$  is a nonincreasing function of  $\text{OPT}$  and  $\text{OPT} > \frac{2}{\varepsilon} + 1$ , we have that

$$\frac{\varepsilon(m - \text{OPT}) + 1}{\varepsilon(\text{OPT} - 1) - 1} < \varepsilon m.$$

Hence,  $K_G < \frac{3}{\varepsilon \text{OPT} - 1} \max \{1, \varepsilon m\} = \frac{9\varepsilon m}{\varepsilon \text{OPT} - 1}$ , since  $m > \frac{1}{3\varepsilon}$  and therefore, we have that  $\tau - K_G \geq \tau(1 - 9\varepsilon)$ . Hence, the first term of (10) can be upper bounded by

$$\begin{aligned}
O \left( \frac{K_G}{\delta \tau (1 - 9\varepsilon)} + \exp \left( -\frac{1}{\delta} \right) \right) \cdot \text{OPT} &= O \left( \frac{9\varepsilon m}{\varepsilon \text{OPT} - 1} \cdot \frac{1}{1 - 9\varepsilon} \cdot \frac{2 \ln n}{\frac{m}{\varepsilon \text{OPT} - 1}} + \frac{\text{OPT}}{n^2} \right) \\
&= O \left( \frac{\varepsilon}{1 - \varepsilon} \cdot \log n \right).
\end{aligned}$$

Hence, the average sensitivity of the algorithm obtained can be bounded by:

$$\begin{aligned}
\beta(G) &= \max \left\{ O \left( \frac{1}{\varepsilon} \right), O \left( \frac{\varepsilon}{1 - \varepsilon} \cdot \log n \right) + \left( \frac{m}{\varepsilon^3 \text{OPT}} \right)^{O(1/\varepsilon^2)} \right\} \\
&= O \left( \frac{\varepsilon}{1 - \varepsilon} \log n \right) + \left( \frac{m}{\varepsilon^3 \text{OPT}} \right)^{O(1/\varepsilon^2)}. \quad \square
\end{aligned}$$

### 6.3.4 Average sensitivity of a combined matching algorithm

In this section, we combine the algorithms guaranteed by Theorems 6.1 and 6.10 in order to get a matching algorithm with improved sensitivity.

**Theorem 6.11.** *Let  $\varepsilon \in (0, 1)$  be a parameter. There exists an algorithm with approximation ratio  $1 - \varepsilon$  and average sensitivity*

$$\text{OPT}(G)^{\frac{c}{c+1}} \cdot O \left( \left( \frac{\varepsilon}{1 - \varepsilon} \cdot \log n \right)^{\frac{1}{c+1}} + \frac{1}{\varepsilon^{\frac{3c}{c+1}}} \right)$$

for  $c = O(1/\varepsilon^2)$ .

---

**Algorithm 6:** COMBINED ALGORITHM TO  $(1 - \varepsilon)$ -APPROXIMATE MAXIMUM MATCHING

---

**Input:** undirected unweighted graph  $G = (V, E)$

- 1 Compute OPT.;
- 2 **if**  $\text{OPT} < 2c$  or  $m < 2c$  **then**
- 3     **return** an arbitrary maximum matching in  $G$ .
- 4 **else**
- 5     Let  $f(G) \leftarrow \frac{\text{OPT}^2}{m}$  and  $g(G) \leftarrow \frac{\varepsilon}{(1-\varepsilon)} \cdot \log n + \left(\frac{m}{\varepsilon^3 \text{OPT}}\right)^c$ ;
- 6     Run the algorithm given by Theorem 6.1 with probability  $\frac{g(G)}{f(G)+g(G)}$  and run the algorithm given by Theorem 6.10 with the remaining probability.

---

*Proof.* Let  $c = O(1/\varepsilon^2)$ . The algorithm guaranteed by the theorem is given as Algorithm 6. The bounds on approximation guarantee and average sensitivity are both straightforward when  $\text{OPT} < 2c$  or  $m < 2c$ .

The approximation guarantee in the case when  $\text{OPT} \geq 2c$  and  $m \geq 2c$  is also straightforward since Algorithm 6 is simply a distribution over algorithms guaranteed by Theorem 6.1 and Theorem 6.10.

We now bound the average sensitivity of Algorithm 6 when  $\text{OPT} \geq 2c$  and  $m \geq 2c$ . Let  $\rho(G)$  denote the probability  $\frac{g(G)}{f(G)+g(G)}$ . By Theorem 8.2, the average sensitivity is at most

$$\frac{O(f(G)) \cdot g(G) + O(g(G)) \cdot f(G)}{f(G) + g(G)} + 2\text{OPT} \cdot \mathbb{E}_{e \sim E} [|\rho(G) - \rho(G - e)|]. \quad (11)$$

We first bound the quantity  $\mathbb{E}_{e \sim E} [|\rho(G) - \rho(G - e)|]$ .

**Claim 6.12.** For every graph  $G = (V, E)$  such that  $\text{OPT} \geq c + 1$ , and for every  $e \in E$ ,

$$\left(1 - \frac{c}{m}\right) \cdot g(G) \leq g(G - e) \leq \left(1 + \frac{c}{\text{OPT} - c}\right) \cdot g(G).$$

*Proof.* We first prove the upper bound. We know that

$$\begin{aligned} \frac{g(G - e)}{g(G)} &\leq \left(1 + \frac{\left(\frac{m-1}{\varepsilon^3(\text{OPT}-1)}\right)^c - \left(\frac{m}{\varepsilon^3 \text{OPT}}\right)^c}{\frac{\varepsilon \log n}{(1-\varepsilon)} + \frac{m^c}{\varepsilon^{3c} \text{OPT}^c}}\right) \leq \left(1 + \frac{\left(\frac{m-1}{\varepsilon^3(\text{OPT}-1)}\right)^c - \left(\frac{m}{\varepsilon^3 \text{OPT}}\right)^c}{\frac{m^c}{\varepsilon^{3c} \text{OPT}^c}}\right) \\ &= \left(1 - \frac{1}{m}\right)^c \cdot \left(1 + \frac{1}{\text{OPT} - 1}\right)^c \leq \left(1 + \frac{c}{\text{OPT} - c}\right). \end{aligned}$$

Note that the last inequality holds whenever  $\text{OPT} > c$ , because  $(1 + x)^r \leq 1 + \frac{rx}{1 - (r-1)x}$  for  $x \in [0, \frac{1}{r-1})$  and  $r > 1$ .

For the lower bound,

$$\begin{aligned} \frac{g(G - e)}{g(G)} &\geq \left(1 - \frac{\left(\frac{m}{\varepsilon^3 \text{OPT}}\right)^c - \left(\frac{m-1}{\varepsilon^3 \text{OPT}}\right)^c}{\frac{\varepsilon \log n}{(1-\varepsilon)} + \frac{m^c}{\varepsilon^{3c} \text{OPT}^c}}\right) \geq \left(1 - \frac{\left(\frac{m}{\varepsilon^3 \text{OPT}}\right)^c - \left(\frac{m-1}{\varepsilon^3 \text{OPT}}\right)^c}{\frac{m^c}{\varepsilon^{3c} \text{OPT}^c}}\right) \\ &= \left(1 - \frac{1}{m}\right)^c \geq 1 - \frac{c}{m}. \quad \square \end{aligned}$$

**Claim 6.13.** For every graph  $G = (V, E)$  and every  $e \in E$ ,

$$f(G) \cdot \left(1 - \frac{2}{\text{OPT}}\right) \leq f(G - e) \leq f(G) \cdot \left(1 + \frac{1}{m-1}\right).$$

*Proof.* To prove the upper bound,

$$\frac{f(G - e)}{f(G)} \leq \left(\frac{m}{m-1}\right) = \left(1 + \frac{1}{m-1}\right).$$

For the lower bound,

$$\frac{f(G - e)}{f(G)} \geq \left(\frac{\text{OPT} - 1}{\text{OPT}}\right)^2 \cdot \left(\frac{m}{m-1}\right)^2 \geq \left(\frac{\text{OPT} - 1}{\text{OPT}}\right)^2 \geq 1 - \frac{2}{\text{OPT}}. \quad \square$$

**Claim 6.14.** For every graph  $G = (V, E)$  such that  $\text{OPT} \geq 2c$  and  $m \geq 2c$ , and for every  $e \in E$ ,

$$\rho(G) \cdot \left(1 - \frac{2c}{\text{OPT} - c}\right) \leq \rho(G - e) \leq \rho(G) \cdot \left(1 + \frac{5c}{\text{OPT} - c}\right).$$

*Proof.* Note that  $\left(1 - \frac{2}{\text{OPT}}\right)^{-1} \leq 1 + \frac{4}{\text{OPT}}$  and  $\left(1 - \frac{c}{m}\right)^{-1} \leq 1 + \frac{2c}{m}$  for  $\text{OPT} \geq 4$  and  $m \geq 2c$ . We also have  $\left(1 + \frac{c}{\text{OPT} - c}\right)^{-1} \geq 1 - \frac{c}{\text{OPT} - c}$  and  $\left(1 + \frac{1}{m-1}\right)^{-1} \geq 1 - \frac{1}{m-1}$  for  $\text{OPT} \geq 2c$  and  $m \geq 2$ . Combining all of the above,

$$\begin{aligned} \rho(G - e) &= \frac{g(G - e)}{f(G - e) + g(G - e)} \\ &\leq \frac{g(G) \cdot \left(1 + \frac{c}{\text{OPT} - c}\right)}{(f(G) + g(G)) \cdot \min\left\{1 - \frac{c}{m}, 1 - \frac{2}{\text{OPT}}\right\}} \\ &\leq \rho(G) \cdot \left(1 + \frac{c}{\text{OPT} - c}\right) \cdot \max\left\{1 + \frac{2c}{m}, 1 + \frac{4}{\text{OPT}}\right\} \\ &\leq \rho(G) \cdot \left(1 + \frac{5c}{\text{OPT} - c}\right). \end{aligned}$$

Using similar calculations, we can see that

$$\rho(G - e) \geq \rho(G) \cdot \left(1 - \frac{c}{m}\right) \cdot \min\left\{1 - \frac{c}{\text{OPT} - c}, 1 - \frac{1}{m-1}\right\} \geq \rho(G) \cdot \left(1 - \frac{2c}{\text{OPT} - c}\right). \quad \square$$

Thus, for all  $e \in E$ , we have that  $|\rho(G) - \rho(G - e)| \leq \max\left\{\frac{2c}{\text{OPT} - c}, \frac{5c}{\text{OPT} - c}\right\} \cdot \rho(G) = \frac{5c\rho(G)}{\text{OPT} - c}$ . Hence,  $\mathbb{E}_{e \sim E} [|\rho(G) - \rho(G - e)|] \leq \frac{5c\rho(G)}{\text{OPT} - c}$ .

Therefore, the average sensitivity of Algorithm 6 is at most

$$\begin{aligned} &\frac{O(f(G)) \cdot g(G) + O(g(G)) \cdot f(G)}{f(G) + g(G)} + 2\text{OPT} \cdot \mathbb{E}_{e \sim E} [|\rho(G) - \rho(G - e)|] \\ &= O\left(\frac{f(G)^{c/(c+1)}g(G)^{1/(c+1)}}{\frac{g(G)^{1/(c+1)}}{f(G)^{1/(c+1)}} + \frac{f(G)^{c/(c+1)}}{g(G)^{c/(c+1)}}}\right) + O\left(\frac{\text{OPT}c\rho(G)}{\text{OPT}}\right) \end{aligned}$$



$$\begin{aligned}
&= O\left(f(G)^{c/(c+1)}g(G)^{1/(c+1)}\right) + O(1/\varepsilon^2) \\
&= O\left(\left(\frac{\text{OPT}^2}{m}\right)^{c/(c+1)} \cdot \left(\left(\frac{\varepsilon}{1-\varepsilon}\log n\right)^{1/(c+1)} + \left(\frac{m^c}{\varepsilon^{3c}\text{OPT}^c}\right)^{1/(c+1)}\right)\right) + O(1/\varepsilon^2) \\
&= O\left(\left(\frac{\text{OPT}^{2c/(c+1)}}{m^{c/(c+1)}}\left(\frac{\varepsilon}{1-\varepsilon}\right)^{1/(c+1)}\log^{1/(c+1)}n + \frac{\text{OPT}^{2c/(c+1)}}{m^{c/(c+1)}}\frac{m^{c/(c+1)}}{\varepsilon^{3c/(c+1)}\text{OPT}^{c/(c+1)}}\right)\right) + O(1/\varepsilon^2) \\
&= O\left(\text{OPT}^{c/(c+1)}\left(\frac{\varepsilon}{1-\varepsilon}\right)^{1/(c+1)}\log^{1/(c+1)}n + \frac{\text{OPT}^{c/(c+1)}}{\varepsilon^{3c/(c+1)}}\right) + O(1/\varepsilon^2) \\
&= O\left(\text{OPT}^{c/(c+1)}\left(\left(\frac{\varepsilon}{1-\varepsilon}\right)^{1/(c+1)}\log^{1/(c+1)}n + \frac{1}{\varepsilon^{3c/(c+1)}}\right)\right). \quad \square
\end{aligned}$$

To obtain the first term of the expression resulting from the first equality, we divide both the numerator and denominator by  $f(G)^{\frac{1}{c+1}} \cdot g(G)^{\frac{c}{c+1}}$ . The second term of the first equality above follows since  $\frac{\text{OPT}}{\text{OPT}-c} \leq 2$  as  $\text{OPT} \geq 2c$ .

## 6.4 Lower bound

In this section, we show a lower bound of  $\Omega(n)$  for the problem of exactly computing the maximum matching in a graph.

**Theorem 6.15.** *Every algorithm that exactly computes the maximum matching in a graph has average sensitivity  $\Omega(n)$ .*

*Proof.* Let  $n \in \mathbb{N}$  be even. Consider the cycle  $C_n$  on  $n$  vertices.  $C_n$  has exactly two maximum matchings  $M_1$  and  $M_2$  of size  $n/2$  each. Both  $M_1$  and  $M_2$  consist of alternating edges of the cycle. Let  $A$  be an algorithm that outputs  $M_1$  with probability  $p$  and  $M_2$  with probability  $1-p$ . Assume, without loss of generality, that  $p \geq \frac{1}{2}$ . For every edge  $e \in M_1$ , the unique maximum matching in the odd-length path  $G-e$  has Hamming distance  $n-1$  from  $M_1$ . Thus, for each  $e \in M_1$ , the earth mover's distance between  $A(G)$  and  $A(G-e)$  is at least  $\frac{n-1}{2}$ . Hence, the average sensitivity of  $A$  is at least  $\frac{1}{n} \sum_{e \in M_1} \frac{n-1}{2} = \Omega(n)$ .  $\square$

## 7 2-Coloring

In the *2-coloring problem*, given a bipartite graph  $G = (V, E)$ , we are to output a (proper) 2-coloring on  $G$ , that is, an assignment  $f : V \rightarrow \{0, 1\}$  such that  $f(u) \neq f(v)$  for every edge  $(u, v) \in E$ . Clearly this problem can be solved in linear time. In this section, however, we show that there is no stable-on-average algorithm for the 2-coloring problem.

**Theorem 7.1.** *Any (randomized) algorithm for the 2-coloring problem has average sensitivity  $\Omega(n)$ .*

*Proof.* Suppose that there is a (randomized) algorithm  $\mathcal{A}$  whose average sensitivity is at most  $\beta n$  for  $\beta < 1/256$ . In what follows, we assume that  $n$ , that is, the number of vertices in the input graph, is a multiple of 16.

Let  $\mathcal{P}_n$  be the family of all possible paths on  $n$  vertices, and let  $\mathcal{Q}_n$  be the family of all possible graphs on  $n$  vertices consisting of two paths. Note that  $|\mathcal{P}_n| = n!/2$  and  $|\mathcal{Q}_n| = (n-1)n!/4$ .

Consider a bipartite graph  $H = (\mathcal{P}_n, \mathcal{Q}_n; E)$ , where a pair  $(P, Q)$  is in  $E$  if and only if  $Q$  can be obtained by removing an edge in  $P$ . Note that each  $P \in \mathcal{P}_n$  has  $n - 1$  neighbors in  $H$  and each  $Q \in \mathcal{Q}_n$  has four neighbors in  $H$ .

We say that an edge  $(P, Q) \in E$  is *intimate* if  $d_{\text{EM}}(\mathcal{A}(P), \mathcal{A}(Q)) \leq 8\beta n$ . We observe that for every  $P \in \mathcal{P}_n$ , at least a  $7/8$ -fraction of the edges incident to  $P$  are intimate; otherwise

$$\mathbb{E}_{e \sim E(P)} [d_{\text{EM}}(\mathcal{A}(P), \mathcal{A}(P - e))] > \frac{1}{8} \cdot 8\beta n = \beta n,$$

which is a contradiction, where  $E(P)$  denotes the set of edges in  $P$ .

We say that a graph  $Q \in \mathcal{Q}_n$  is *heavy* if both components of  $Q$  have at least  $n/16$  vertices, and say that an edge  $(P, Q) \in E$  is *heavy* if  $Q$  is heavy. We observe that for every  $P \in \mathcal{P}_n$ , at least a  $7/8$ -fraction of the edges incident to  $P$  are heavy.

We say that an edge  $(P, Q) \in E$  is *good* if it is intimate and heavy. Observe that for every  $P \in \mathcal{P}_n$ , by the union bound, at least a  $3/4$ -fraction of the edges incident to  $P$  are good. In particular, this means that the fraction of good edges in  $H$  is at least  $3/4$ . Hence, there exists  $Q^* \in \mathcal{Q}_n$  that has at least three good incident edges; otherwise the fraction of good edges in  $H$  is at most  $2/4 = 1/2$ , which is a contradiction.

Let  $f_1, \dots, f_4$  be the four 2-colorings of  $Q^*$ . As  $Q^*$  has three good incident edges, without loss of generality, there are adjacent paths  $P_1, P_2 \in \mathcal{P}_n$  such that both  $(P_1, Q^*)$  and  $(P_2, Q^*)$  are good, and there is no assignment that is a 2-coloring for both  $P_1$  and  $P_2$ . Without loss of generality, we assume that  $f_1, f_2$  are 2-colorings of  $P_1$ , and  $f_3, f_4$  are 2-colorings of  $P_2$ . Note that  $d_{\text{Ham}}(f_i, f_j) \geq n/16$  for  $i \neq j$  because  $Q$  is heavy. Let  $q_i = \Pr[\mathcal{A}(Q^*) = f_i]$  for  $i \in [4]$ . As the edge  $(P_1, Q^*)$  is intimate, we have

$$\begin{aligned} 8\beta n &\geq d_{\text{EM}}(\mathcal{A}(P_1), \mathcal{A}(Q^*)) \geq \frac{n}{16} (|\Pr[\mathcal{A}(P_1) = f_1] - q_1| + |\Pr[\mathcal{A}(P_1) = f_2] - q_2| + q_3 + q_4) \\ &= \frac{n}{16} (|\Pr[\mathcal{A}(P_1) = f_1] - q_1| + |\Pr[\mathcal{A}(P_1) = f_2] - q_2| + 1 - q_1 - q_2) \end{aligned}$$

and hence we must have  $q_1 + q_2 \geq 1 - 128\beta$ . Considering  $d_{\text{EM}}(\mathcal{A}(P_2), \mathcal{A}(Q^*))$ , we also have  $q_3 + q_4 \geq 1 - 128\beta$ . However,

$$1 = q_1 + q_2 + q_3 + q_4 \geq (1 - 128\beta) + (1 - 128\beta) = 2 - 256\beta > 1$$

as  $\beta < 1/256$ , which is a contradiction.  $\square$

## 8 General Results on Average Sensitivity

In this section, we state and prove some basic properties of average sensitivity and show that locality guarantees of solutions output by an algorithm imply low average sensitivity for that algorithm.

### 8.1 Bounds on $k$ -average sensitivity from bounds on average sensitivity

In this section, we prove Theorem 1.5, which says that, if an algorithm is stable-on-average against deleting a single edge, it is also stable-on-average against deleting multiple edges. We restate the theorem here.

**Theorem 1.5.** *Let  $\mathcal{A}$  be an algorithm for a graph problem with average sensitivity at most  $f(n, m)$ . Then, for any integer  $k \geq 1$ , the algorithm  $\mathcal{A}$  has  $k$ -average sensitivity at most  $\sum_{i=1}^k f(n, m - i + 1)$ .*

*Proof.* We have

$$\begin{aligned}
& \mathbb{E}_{\{e_1, \dots, e_k\} \sim \binom{E}{k}} \left[ d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}(G - \{e_1, \dots, e_k\})) \right] \\
& \leq \mathbb{E}_{\{e_1, \dots, e_k\} \sim \binom{E}{k}} \left[ \sum_{i=1}^k d_{\text{EM}}(\mathcal{A}(G - \{e_1, \dots, e_{i-1}\}), \mathcal{A}(G - \{e_1, \dots, e_i\})) \right] \\
& = \mathbb{E}_{e_1 \sim E} \left[ d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}(G - \{e_1\})) \right] + \mathbb{E}_{e_2 \sim E \setminus \{e_1\}} \left[ d_{\text{EM}}(\mathcal{A}(G - \{e_1\}), \mathcal{A}(G - \{e_1, e_2\})) \right] + \dots \\
& \quad + \mathbb{E}_{e_k \sim E \setminus \{e_1, \dots, e_{k-1}\}} \left[ d_{\text{EM}}(\mathcal{A}(G - \{e_1, \dots, e_{k-1}\}), \mathcal{A}(G - \{e_1, \dots, e_k\})) \dots \right] \\
& \leq f(n, m) + \mathbb{E}_{e_1 \sim E} \left[ \beta(G - \{e_1\}) \right] + \mathbb{E}_{e_2 \sim E \setminus \{e_1\}} \left[ \beta(G - \{e_1, e_2\}) \right] + \dots \\
& \quad + \mathbb{E}_{e_{k-1} \sim E \setminus \{e_1, \dots, e_{k-2}\}} \left[ \beta(G - \{e_1, \dots, e_{k-1}\}) \dots \right] \\
& \leq \sum_{i=1}^k f(n, m - i + 1).
\end{aligned}$$

Here, the first inequality is due to the triangle inequality.  $\square$

## 8.2 Sequential composition

In this section, we state and prove our two sequential composition theorems Theorem 1.6 and Theorem 1.7.

**Theorem 1.6** (Sequential composition). *Consider two randomized algorithms  $\mathcal{A}_1 : \mathcal{G} \rightarrow \mathcal{S}_1$ ,  $\mathcal{A}_2 : \mathcal{G} \times \mathcal{S}_1 \rightarrow \mathcal{S}_2$ . Suppose that the average sensitivity of  $\mathcal{A}_1$  with respect to the total variation distance is  $\gamma_1(G)$  and the average sensitivity of  $\mathcal{A}_2(\cdot, S_1)$  is  $\beta_2^{(S_1)}(G)$  for any  $G \in \mathcal{G}$ ,  $S_1 \in \mathcal{S}_1$ . Let  $\mathcal{A} : \mathcal{G} \rightarrow \mathcal{S}_2$  be a randomized algorithm obtained by composing  $\mathcal{A}_1$  and  $\mathcal{A}_2$ , that is,  $\mathcal{A}(G) = \mathcal{A}_2(G, \mathcal{A}_1(G))$  for  $G \in \mathcal{G}$ . Then, the average sensitivity of  $\mathcal{A}$  on a graph  $G \in \mathcal{G}$  is  $H \cdot \gamma_1(G) + \mathbb{E}_{S_1 \sim \mathcal{A}_1(G)} \left[ \beta_2^{(S_1)}(G) \right]$ , where  $H$  denotes the maximum over Hamming weights of all the solutions output (with nonzero probability) by running  $\mathcal{A}$  on  $G$  and all of the graphs in  $\{G - e : e \in E\}$ .*

*Proof.* Consider  $G = (V, E)$  and let  $e \in E$ . We bound the earth mover's distance between  $\mathcal{A}(G)$  and  $\mathcal{A}(G - e)$  as follows. For a distribution  $\mathcal{D}$ , we use  $f_{\mathcal{D}}$  to denote its probability mass function. We know that for all  $S_1 \in \mathcal{S}_1$  and  $S_2 \in \mathcal{S}_2$

$$f_{(\mathcal{A}_1(G), \mathcal{A}_2(G, S_1))}(S_1, S_2) = f_{\mathcal{A}_1(G)}(S_1) \cdot f_{\mathcal{A}_2(G, S_1)}(S_2),$$

where  $(\mathcal{A}_1(G), \mathcal{A}_2(G, S_1))$  denotes the joint distribution of  $\mathcal{A}_1(G)$  and  $\mathcal{A}_2(G, S_1)$ . Fix  $S_1 \in \mathcal{S}_1$ . For each  $S_2 \in \mathcal{S}_2$ , we transform probabilities of the form  $f_{(\mathcal{A}_1(G), \mathcal{A}_2(G, S_1))}(S_1, S_2)$  to  $f_{\mathcal{A}_1(G)}(S_1) \cdot f_{\mathcal{A}_2(G-e, S_1)}(S_2)$ . This incurs a total cost of  $f_{\mathcal{A}_1(G)}(S_1) \cdot d_{\text{EM}}(\mathcal{A}_2(G, S_1), \mathcal{A}_2(G - e, S_1))$ . We can now, for each  $S_1 \in \mathcal{S}_1$  and  $S_2 \in \mathcal{S}_2$ , transform the probability  $f_{\mathcal{A}_1(G)}(S_1) \cdot f_{\mathcal{A}_2(G-e, S_1)}(S_2)$  into  $f_{\mathcal{A}_1(G-e)}(S_1) \cdot f_{\mathcal{A}_2(G-e, S_1)}(S_2)$  at a cost of at most  $d_{\text{TV}}(\mathcal{A}_1(G), \mathcal{A}_1(G - e)) \cdot H$ , where  $H$  denotes the

maximum Hamming weight among those of solutions obtained by running  $\mathcal{A}$  on  $G$  and  $\{G - e\}_{e \in E}$ . Thus, the earth mover's distance between  $\mathcal{A}(G)$  and  $\mathcal{A}(G - e)$  is at most

$$d_{\text{TV}}(\mathcal{A}_1(G), \mathcal{A}_1(G - e)) \cdot H + \int_{\mathcal{S}_1} f_{\mathcal{A}_1(G)}(S_1) \cdot d_{\text{EM}}(\mathcal{A}_2(G, S_1), \mathcal{A}_2(G - e, S_1)) \, dS_1.$$

Hence, the average sensitivity of  $\mathcal{A}$  can be bounded as:

$$\begin{aligned} \mathbb{E}_{e \sim E} [d_{\text{EM}}(\mathcal{A}(G), \mathcal{A}(G - e))] &\leq H \cdot \mathbb{E}_{e \sim E} [d_{\text{TV}}(\mathcal{A}_1(G), \mathcal{A}_1(G - e))] \\ &\quad + \mathbb{E}_{e \sim E} \left[ \int_{\mathcal{S}_1 \in \mathcal{S}_1} f_{\mathcal{A}_1(G)}(S_1) \cdot d_{\text{EM}}(\mathcal{A}_2(G, S_1), \mathcal{A}_2(G - e, S_1)) \, dS_1 \right] \\ &\leq H\gamma_1(G) + \mathbb{E}_{S_1 \sim \mathcal{A}_1(G)} [d_{\text{EM}}(\mathcal{A}_2(G, S_1), \mathcal{A}_2(G - e, S_1))] \\ &= H\gamma_1(G) + \mathbb{E}_{S_1 \sim \mathcal{A}_1(G)} \left[ \mathbb{E}_{e \sim E} d_{\text{EM}}(\mathcal{A}_2(G, S_1), \mathcal{A}_2(G - e, S_1)) \right] \\ &= H\gamma_1(G) + \mathbb{E}_{S_1 \sim \mathcal{A}_1(G)} \left[ \beta_2^{(S_1)}(G) \right]. \end{aligned}$$

We are able to interchange the order of expectations because of Fubini's theorem [13].  $\square$

The following theorem states the composition of average sensitivity with respect to the total variation distance.

**Theorem 1.7** (Sequential composition w.r.t. the TV distance). *Consider  $\ell$  randomized algorithms  $\mathcal{A}_i : \mathcal{G} \times \prod_{j=1}^{i-1} \mathcal{S}_j \rightarrow \mathcal{S}_i$  for  $i \in \{1, \dots, \ell\}$ . Suppose that, for each  $i \in \{1, \dots, \ell\}$ , the average sensitivity of  $\mathcal{A}_i(\cdot, S_1, \dots, S_{i-1})$  is  $\gamma_i(G)$  with respect to the total variation distance for every  $G \in \mathcal{G}$  and  $S_1 \in \mathcal{S}_1, \dots, S_{i-1} \in \mathcal{S}_{i-1}$ . Consider a sequence of computations  $S_1 = \mathcal{A}_1(G), S_2 = \mathcal{A}_2(G, S_1), \dots, S_\ell = \mathcal{A}_\ell(G, S_1, \dots, S_{\ell-1})$  for  $G \in \mathcal{G}$ . Let  $\mathcal{A} : \mathcal{G} \rightarrow \mathcal{S}_\ell$  be a randomized algorithm that performs this sequence of computations on input  $G$  and outputs  $S_\ell$ . Then, the average sensitivity of  $\mathcal{A}$  on a graph  $G \in \mathcal{G}$  with respect to the total variation distance is at most  $\sum_{i=1}^{\ell} \gamma_i(G)$ .*

Theorem 1.7 can be immediately obtained by iteratively applying Lemma 8.1.

**Lemma 8.1.** *Consider two randomized algorithms  $\mathcal{A}_1 : \mathcal{G} \rightarrow \mathcal{S}_1, \mathcal{A}_2 : \mathcal{G} \times \mathcal{S}_1 \rightarrow \mathcal{S}_2$  for a graph problem. Suppose that the average sensitivity of  $\mathcal{A}_1$  is  $\gamma_1(G)$  and the average sensitivity of  $\mathcal{A}_2(\cdot, S_1)$  is  $\gamma_2(G)$  for any  $S_1 \in \mathcal{S}_1$ , both with respect to the total variation distance. Let  $\mathcal{A} : \mathcal{G} \rightarrow \mathcal{S}_2$  be a randomized algorithm obtained by composing  $\mathcal{A}_1$  and  $\mathcal{A}_2$ , that is,  $\mathcal{A}(G) = \mathcal{A}_2(G, \mathcal{A}_1(G))$ . Then, the average sensitivity of  $\mathcal{A}$  is  $\gamma_1(G) + \gamma_2(G)$  with respect to the total variation distance.*

*Proof.* For a distribution  $\mathcal{D}$ , we use  $f_{\mathcal{D}}$  to denote its probability mass function. Consider a graph  $G = (V, E)$ . Note that

$$f_{\mathcal{A}(G)}(S_2) = \int_{\mathcal{S}_1} f_{\mathcal{A}_2(G, S_1)}(S_2) f_{\mathcal{A}_1(G)}(S_1) \, dS_1.$$

Then we have that, for  $e \in E$ ,

$$d_{\text{TV}}(\mathcal{A}(G), \mathcal{A}(G - e))$$

$$\begin{aligned}
&= \frac{1}{2} \int_{\mathcal{S}_2} \left| \int_{\mathcal{S}_1} f_{\mathcal{A}_2(G, S_1)}(S_2) f_{\mathcal{A}_1(G)}(S_1) \, dS_1 - \int_{\mathcal{S}_1} f_{\mathcal{A}_2(G-e, S_1)}(S_2) f_{\mathcal{A}_1(G-e)}(S_1) \, dS_1 \right| \, dS_2 \\
&= \frac{1}{2} \int_{\mathcal{S}_2} \left| \int_{\mathcal{S}_1} f_{\mathcal{A}_2(G, S_1)}(S_2) \left( f_{\mathcal{A}_1(G)}(S_1) - f_{\mathcal{A}_1(G-e)}(S_1) \right) \, dS_1 - \right. \\
&\quad \left. \int_{\mathcal{S}_1} \left( f_{\mathcal{A}_2(G-e, S_1)}(S_2) - f_{\mathcal{A}_2(G, S_1)}(S_2) \right) f_{\mathcal{A}_1(G-e)}(S_1) \, dS_1 \right| \, dS_2 \\
&\leq \frac{1}{2} \int_{\mathcal{S}_1} \left| f_{\mathcal{A}_1(G)}(S_1) - f_{\mathcal{A}_1(G-e)}(S_1) \right| \, dS_1 \cdot \int_{\mathcal{S}_2} f_{\mathcal{A}_2(G, S_1)}(S_2) \, dS_2 + \\
&\quad \int_{\mathcal{S}_1} f_{\mathcal{A}_1(G-e)}(S_1) \, dS_1 \cdot \frac{1}{2} \int_{\mathcal{S}_2} \left| f_{\mathcal{A}_2(G-e, S_1)}(S_2) - f_{\mathcal{A}_2(G, S_1)}(S_2) \right| \, dS_2 \\
&= \frac{1}{2} \int_{\mathcal{S}_1} \left| f_{\mathcal{A}_1(G)}(S_1) - f_{\mathcal{A}_1(G-e)}(S_1) \right| \, dS_1 + \\
&\quad \int_{\mathcal{S}_1} f_{\mathcal{A}_1(G-e)}(S_1) \, dS_1 \cdot \frac{1}{2} \int_{\mathcal{S}_2} \left| f_{\mathcal{A}_2(G-e, S_1)}(S_2) - f_{\mathcal{A}_2(G, S_1)}(S_2) \right| \, dS_2 \\
&= d_{\text{TV}}(\mathcal{A}_1(G), \mathcal{A}_1(G-e)) + \int_{\mathcal{S}_1} f_{\mathcal{A}_1(G-e)}(S_1) \cdot d_{\text{TV}}(\mathcal{A}_2(G, S_1), \mathcal{A}_2(G-e, S_1)) \, dS_1.
\end{aligned}$$

Hence, the average sensitivity of  $\mathcal{A}$  with respect to the total variation distance can be bounded as,

$$\begin{aligned}
\mathbb{E}_{e \sim E} [d_{\text{TV}}(\mathcal{A}(G), \mathcal{A}(G-e))] &\leq \mathbb{E}_{e \sim E} [d_{\text{TV}}(\mathcal{A}_1(G), \mathcal{A}_1(G-e))] + \\
&\quad \mathbb{E}_{e \sim E} \left[ \int_{\mathcal{S}_1} f_{\mathcal{A}_1(G-e)}(S_1) \cdot d_{\text{TV}}(\mathcal{A}_2(G, S_1), \mathcal{A}_2(G-e, S_1)) \, dS_1 \right] \\
&\leq \gamma_1(G) + \int_{\mathcal{S}_1} f_{\mathcal{A}_1(G-e)}(S_1) \, dS_1 \cdot \gamma_2(G) = \gamma_1(G) + \gamma_2(G). \quad \square
\end{aligned}$$

### 8.3 Parallel composition

In this section, we prove Theorem 1.8, which bounds the average sensitivity of an algorithm obtained by running different algorithms according to a distribution in terms of the average sensitivities of the component algorithms. We restate the theorem here.

**Theorem 1.8** (Parallel composition). *Let  $\mathcal{A}_1, \dots, \mathcal{A}_\ell$  be algorithms for a graph problem with average sensitivities  $\beta_1(G), \dots, \beta_\ell(G)$ , respectively. Let  $\mathcal{A}$  be an algorithm that, given a graph  $G$ , runs  $\mathcal{A}_i$  with probability  $\rho_i(G)$  for  $i \in \{1, \dots, \ell\}$ , where  $\sum_{i \in \{1, \dots, \ell\}} \rho_i(G) = 1$ . Let  $\mathsf{H}$  denote the maximum over Hamming weights of all solutions output (with nonzero probability) by running  $\mathcal{A}$  on  $G$  and on all the graphs in  $\{G-e : e \in E\}$ . Then the average sensitivity of  $\mathcal{A}$  is at most  $\sum_{i \in \{1, \dots, \ell\}} \rho_i(G) \cdot \beta_i(G) + \mathsf{H} \cdot \mathbb{E}_{e \sim E} \left[ \sum_{i \in \{1, \dots, \ell\}} |\rho_i(G) - \rho_i(G-e)| \right]$ .*

*Proof.* Consider a graph  $G = (V, E)$ . For a solution  $S$ , let  $p^G(S)$  denote the probability that  $S$  is output on input  $G$  by  $\mathcal{A}$ . Let  $p_i^G(S)$  denote the probability that  $S$  is output on input  $G$  by  $\mathcal{A}_i$ . For every solution  $S$ , we know that  $p^G(S) = \sum_{i \in [\ell]} \rho_i(G) \cdot p_i^G(S)$ .

Let  $\mathcal{A}(G)$  denote the output distribution of  $\mathcal{A}$  on  $G$ . Fix  $e \in E$ . We first bound the earth mover's distance between  $\mathcal{A}(G)$  and  $\mathcal{A}(G-e)$ . In order to transform  $\mathcal{A}(G)$  into  $\mathcal{A}(G-e)$ , we first

transform  $p^G(S)$ , for each solution  $S$ , into  $\sum_{i \in [\ell]} \rho_i(G) \cdot p_i^{G-e}(S)$ . This can be done at a cost of at most  $\sum_{i \in [\ell]} \rho_i(G) \cdot d_{\text{EM}}(\mathcal{A}_i(G), \mathcal{A}_i(G-e))$ .

We now convert  $\sum_{i \in [\ell]} \rho_i(G) \cdot p_i^{G-e}(S)$ , for each solution  $S$ , into  $\sum_{i \in [\ell]} \rho_i(G-e) \cdot p_i^{G-e}(S)$  at a cost of at most  $2\mathsf{H} \cdot \frac{1}{2} \sum_{i \in [\ell]} |\rho_i(G) - \rho_i(G-e)|$ , where  $\frac{1}{2} \sum_{i \in [\ell]} |\rho_i(G) - \rho_i(G-e)|$  is the total variation distance between the probability distributions with which  $\mathcal{A}$  selects the algorithms on inputs  $G$  and  $G-e$ . Hence, the average sensitivity of  $\mathcal{A}$  is at most

$$\sum_{i \in [\ell]} \rho_i(G) \cdot \beta_i(G) + \mathsf{H} \cdot \mathbb{E}_{e \sim E} \left[ \sum_{i \in [\ell]} |\rho_i(G) - \rho_i(G-e)| \right]. \quad \square$$

We separately state the special case of Theorem 1.8 for  $\ell = 2$ .

**Theorem 8.2.** *Let  $\mathcal{A}_1$  and  $\mathcal{A}_2$  be two algorithms for a graph problem with average sensitivities  $\beta_1(G)$  and  $\beta_2(G)$ , respectively. Let  $\mathcal{A}$  be an algorithm that, given a graph  $G$ , runs  $\mathcal{A}_1$  with probability  $\rho(G)$  and runs  $\mathcal{A}_2$  with the remaining probability. Let  $\mathsf{H}$  denote the maximum Hamming weight among those of solutions obtained by running  $\mathcal{A}$  on  $G$  and  $\{G-e\}_{e \in E}$ . Then the average sensitivity of  $\mathcal{A}$  is at most  $\rho(G) \cdot \beta_1(G) + (1 - \rho(G)) \cdot \beta_2(G) + 2\mathsf{H} \cdot \mathbb{E}_{e \sim E} [|\rho(G) - \rho(G-e)|]$ .*

## 8.4 Sublinearity implies low average sensitivity

In this section, we prove Theorem 1.10, which show that the existence of a sublinear-time solution oracle (Definition 1.9) for an algorithm  $\mathcal{A}$  implies that the average sensitivity of  $\mathcal{A}$  is bounded by the query complexity of that oracle.

**Theorem 1.10** (Sublinearity implies low average sensitivity). *Consider a randomized algorithm  $\mathcal{A} : \mathcal{G} \rightarrow \mathcal{S}$  for a graph problem, where each solution output by  $\mathcal{A}$  is a subset of the set of edges in the input graph. Assume that there exists a solution oracle  $\mathcal{O}$  for  $\mathcal{A}$  such that  $\mathcal{O}$  makes at most  $q(G)$  queries to  $G$  in expectation, where this expectation is taken over the random coins of  $\mathcal{O}$  and over input edges  $e \in E$ . Then,  $\mathcal{A}$  has average sensitivity at most  $q(G)$ . Moreover, given the promise that the input graphs satisfy  $|E| \geq |V|$ , the statement applies also to algorithms for which each solution is a subset of the set of vertices in the input graph.*

*Proof.* We prove the theorem for the case that solutions output by  $\mathcal{A}$  are subsets of edges of the input graph. It can be easily modified to work for the case that the solutions output by  $\mathcal{A}$  are subsets of vertices of the input graph in which case, we will use the technical condition that  $n \leq m$ .

Without loss of generality, assume that  $\mathcal{A}$  uses  $r(n)$  random bits when run on graphs of  $n$  vertices<sup>2</sup>. Consider a graph  $G = (V, E)$  that  $\mathcal{O}$  gets access to. For  $e \in E$  and a string  $\pi \in \{0, 1\}^{r(n)}$ , let  $Q_{e, \pi}$  denote the set of edges in  $E$  queried by  $\mathcal{O}$  on input  $e$ , while simulating the run of  $\mathcal{A}$  with  $\pi$  as the random string. The set  $Q_{e, \pi}$  denotes the set of edges  $e'$  such that the status of  $e$  in the solutions output by  $\mathcal{A}$  with randomness  $\pi$  on inputs  $G$  and  $G-e'$  could be different. For each edge  $e' \in E$  and string  $\pi \in \{0, 1\}^{r(n)}$ , define  $R_{e', \pi}$  as the set of edges  $e \in E$  such that  $e' \in Q_{e, \pi}$ .

By definition, for each  $\pi \in \{0, 1\}^{r(n)}$ , we have  $\sum_{e \in E} |R_{e, \pi}| = \sum_{e \in E} |Q_{e, \pi}|$ . Hence we have:

$$\sum_{\pi \in \{0, 1\}^{r(n)}} \sum_{e \in E} |R_{e, \pi}| = \sum_{\pi \in \{0, 1\}^{r(n)}} \sum_{e \in E} |Q_{e, \pi}|,$$

<sup>2</sup>If  $r(G)$  is the length of the random string used for  $G$ , we can simply set  $r(n) = \max\{r(G) : G = (V, E), |V| = n\}$ . If we do not need  $r(n)$  bits for some particular graph  $G$  on  $n$  vertices, we can just throw away the unused bits.

and

$$\mathbb{E}_{\pi \sim \{0,1\}^{r(n)}} \mathbb{E}_{e \sim E} |R_{e,\pi}| = \mathbb{E}_{\pi \sim \{0,1\}^{r(n)}} \mathbb{E}_{e \sim E} |Q_{e,\pi}| \leq q(G),$$

where the last inequality follows from our assumption on  $\mathcal{O}$ .

For  $\pi \in \{0,1\}^{r(n)}$  and  $e \in E$ , the set  $R_{e,\pi}$  contains the set of edges whose presence in the solution could be affected by the removal of  $e$  from  $G$ . Therefore, it is a superset of the set of edges contained in the symmetric difference between the outputs of  $\mathcal{A}$  on inputs  $G$  and  $G - e$  when run with  $\pi$  as the random string.

Let  $\mathcal{H}_{\mathcal{A},\pi}(G, G')$  denote the Hamming distance between the outputs of the algorithm  $\mathcal{A}$  on inputs  $G$  and  $G'$  when run with  $\pi$  as the random string. As per this notation, for each  $e \in E$ ,

$$\mathbb{E}_{\pi \in \{0,1\}^{r(n)}} \mathcal{H}_{\mathcal{A},\pi}(G, G - e) \leq \mathbb{E}_{\pi \in \{0,1\}^{r(n)}} |R_{e,\pi}|.$$

The following claim relates the quantity on the left hand side of the above inequality with the average sensitivity of  $\mathcal{A}$ .

**Claim 8.3.** *The average sensitivity of  $\mathcal{A}$  is bounded as*

$$\beta(G) \leq \mathbb{E}_{e \in E(G)} \mathbb{E}_{\pi \in \{0,1\}^{r(n)}} \mathcal{H}_{\mathcal{A},\pi}(G, G - e).$$

*Proof.* Fix  $G \in \mathcal{G}$  and  $e \in E(G)$ . We first bound the earth mover's distance between  $\mathcal{A}(G)$  and  $\mathcal{A}(G - e)$ , where  $\mathcal{A}(G)$  and  $\mathcal{A}(G - e)$  are the output distributions of  $\mathcal{A}$  on inputs  $G$  and  $G - e$ , respectively. For  $S \in \mathcal{S}$ , let  $p_G(S)$  and  $p_{G-e}(S)$  denote the probabilities that  $\mathcal{A}$  outputs  $S$  on  $G$  and  $G - e$ , respectively. We start with  $\mathcal{A}(G)$ . Consider a string  $\pi \in \{0,1\}^{r(n)}$ . Let  $S \in \mathcal{S}$  denote the output of  $\mathcal{A}$  on input  $G$  when using the string  $\pi$  as its random string. Let  $S'$  denote the output that is generated when running  $\mathcal{A}$  on input  $G - e$  with  $\pi$  as the random string. We move a mass of  $\frac{1}{2^{r(n)}}$  (corresponding to the string  $\pi$ ) from  $p_G(S)$  to  $p_G(S')$  at a cost of  $\frac{d_{\text{Ham}}(S, S')}{2^{r(n)}}$ . Moving masses corresponding to every string  $\pi \in \{0,1\}^{r(n)}$  this way, we can transform  $\mathcal{A}(G)$  to  $\mathcal{A}(G - e)$ . The total cost incurred during this transformation is  $\mathbb{E}_{\pi \in \{0,1\}^{r(n)}} \mathcal{H}_{\mathcal{A},\pi}(G, G - e)$ . Therefore the earth mover's distance between  $\mathcal{A}(G)$  and  $\mathcal{A}(G - e)$  is at most  $\mathbb{E}_{\pi \in \{0,1\}^{r(n)}} \mathcal{H}_{\mathcal{A},\pi}(G, G - e)$ . Therefore the average sensitivity of  $\mathcal{A}$  is  $\beta(G) \leq \mathbb{E}_{e \in E(G)} \mathbb{E}_{\pi \in \{0,1\}^{r(n)}} \mathcal{H}_{\mathcal{A},\pi}(G, G - e)$ .  $\square$

Therefore, the average sensitivity of  $\mathcal{A}$  is:

$$\beta(G) \leq \mathbb{E}_{e \sim E} \mathbb{E}_{\pi \in \{0,1\}^{r(n)}} \mathcal{H}_{\mathcal{A},\pi}(G, G - e) \leq \mathbb{E}_{e \sim E} \mathbb{E}_{\pi \in \{0,1\}^{r(n)}} |R_{e,\pi}| \leq q(G). \quad \square$$

We now prove Corollary 1.12 which says that the existence of an LCA (Definition 1.11) for a graph problem implies the existence of a stable-on-average algorithm for the same problem.

**Corollary 1.12** (LCAs imply stable-on-average algorithms). *Consider a graph problem  $\mathcal{P} : \mathcal{G} \rightarrow \mathcal{S}$ . Let  $\delta : \mathbb{N} \rightarrow [0, 1]$  and  $q, r : \mathbb{N} \rightarrow \mathbb{N}$ . If  $\mathcal{P}$  has a  $(q, r, \delta)$ -LCA  $\mathcal{L}$ , then, there exists an algorithm  $\mathcal{A}$  for  $\mathcal{P}$ , that on input  $G = (V, E)$ , has an average sensitivity of at most  $q(|V|) + |E| \cdot \delta(|V|)$ .*

*Proof.* Assume without loss of generality that each solution in  $\mathcal{S}$  is a subset of edges of its preimage with respect to  $\mathcal{P}$ . Consider the algorithm  $\mathcal{A}$  that, on input  $G = (V, E)$ , constructs a solution to  $\mathcal{P}$  by running  $\mathcal{L}$  on each edge  $e \in E$  and combining the outputs of  $\mathcal{L}$ . It is clear that  $\mathcal{L}$  is a solution oracle (Definition 1.9) for the algorithm  $\mathcal{A}$ . Hence, the average sensitivity of  $\mathcal{A}$  is upper bounded by the expected number of queries made by  $\mathcal{L}$ , which is at most  $q(|V|) + |E| \cdot \delta(|V|)$ .  $\square$

**Acknowledgments.** We are grateful to anonymous reviewers for suggesting a major improvement to the average sensitivity analysis of Algorithm 4. We thank Tasuku Soma and Samson Zhou for several helpful discussions. We extend our gratitude to Sofya Raskhodnikova for helpful comments that improved the presentation of this article.

## References

- [1] N. K. Ahmed, J. Neville, and R. Kompella. Network sampling. *ACM Transactions on Knowledge Discovery from Data*, 8(2):1–56, 2014.
- [2] N. Alon, R. Rubinfeld, S. Vardi, and N. Xie. Space-efficient local computation algorithms. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1132–1139, 2012.
- [3] A. Bavelas. Communication patterns in task-oriented groups. *The Journal of the Acoustical Society of America*, 22(6):725–730, 1950.
- [4] M. A. Beauchamp. An improved index of centrality. *Behavioral Science*, 10(2):161–163, 1965.
- [5] O. Bousquet and A. Elisseeff. Stability and generalization. *Journal of Machine Learning Research*, pages 499–526, 2002.
- [6] K. Censor-Hillel, E. Haramaty, and Z. S. Karnin. Optimal dynamic distributed MIS. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 217–226, 2016.
- [7] A. Czumaj, Y. Mansour, and S. Vardi. Sublinear graph augmentation for fast query implementation. In *Proceedings of the 16th International Workshop on Approximation and Online Algorithms (WAOA)*, pages 181–203, 2018.
- [8] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Theory of Cryptography Conference (TCC)*, pages 265–284, 2006.
- [9] J. Edmonds. Paths, trees, and flowers. *Canadian Journal of mathematics*, pages 449–467, 1965.
- [10] P. Erdős and A. Rényi. On random graphs. *Publicationes Mathematicae*, 6:290–297, 1959.
- [11] G. Even, M. Medina, and D. Ron. Best of two local models: Centralized local and distributed local algorithms. *Inf. Comput.*, 262(Part):69–89, 2018.
- [12] L. C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, 40(1):35–41, 1977.
- [13] G. Fubini. Sugli integrali multipli. *Accademia dei Lincei, Rendiconti*, 16(1):608–614, 1907.
- [14] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1979.



- [15] A. Gupta, K. Ligett, F. McSherry, A. Roth, and K. Talwar. Differentially private combinatorial optimization. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1106–1125, 2010.
- [16] A. Hassidim, J. A. Kelner, H. N. Nguyen, and K. Onak. Local graph partitions for approximation and testing. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 22–31, 2009.
- [17] A. Hassidim, Y. Mansour, and S. Vardi. Local computation mechanism design. *ACM Trans. Economics and Comput.*, 4(4):21:1–21:24, 2016.
- [18] M. Hay, C. Li, G. Miklau, and D. D. Jensen. Accurate estimation of the degree distribution of private networks. In *Proceedings of the 9th IEEE International Conference on Data Mining (ICDM)*, pages 169–178, 2009.
- [19] D. R. Karger. Global min-cuts in rnc, and other ramifications of a simple min-cut algorithm. In *Proceedings of the 4th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 21–30, 1993.
- [20] V. Karwa, S. Raskhodnikova, A. D. Smith, and G. Yaroslavtsev. Private analysis of graph structure. *ACM Transactions on Database Systems*, 39(3):22:1–22:33, 2014.
- [21] V. Karwa and A. B. Slavkovic. Differentially private graphical degree sequences and synthetic graphs. In *Proceedings of the International Conference on Privacy in Statistical Databases (PSD)*, pages 273–285, 2012.
- [22] S. P. Kasiviswanathan, K. Nissim, S. Raskhodnikova, and A. D. Smith. Analyzing graphs with node differential privacy. In *Proceedings of the 10th Theory of Cryptography (TCC)*, pages 457–476, 2013.
- [23] D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 137–146, 2003.
- [24] J. B. Kruskal. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical Society*, 7(1):48–50, 1956.
- [25] J. R. Lee, S. Oveis Gharan, and L. Trevisan. Multiway spectral partitioning and higher-order cheeger inequalities. *Journal of the ACM*, 61(6):37, 2014.
- [26] S. H. Lee, P.-J. Kim, and H. Jeong. Statistical properties of sampled networks. *Physical Review E*, 73(1), 2006.
- [27] C. Lenzen and R. Levi. A centralized local algorithm for the sparse spanning graph problem. In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 87:1–87:14, 2018.
- [28] R. Levi and M. Medina. A (centralized) local guide. *Bulletin of the EATCS*, 122, 2017.
- [29] R. Levi, D. Ron, and R. Rubinfeld. Local algorithms for sparse spanning graphs. *Algorithmica*, 82(4):747–786, 2020.

- [30] R. Levi, R. Rubinfeld, and A. Yodpinyanee. Local computation algorithms for graphs of non-constant degrees. *Algorithmica*, 77(4):971–994, 2017.
- [31] Y. Mansour, B. Patt-Shamir, and S. Vardi. Constant-time local computation algorithms. *Theory Comput. Syst.*, 62(2):249–267, 2018.
- [32] Y. Mansour, A. Rubinfeld, S. Vardi, and N. Xie. Converting online algorithms to local computation algorithms. In *Proceedings of the 39th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 653–664, 2012.
- [33] Y. Mansour and S. Vardi. A local computation approximation scheme to maximum matching. In *Proceedings of 16th International Workshop on Approximation Algorithms for Combinatorial Optimization (APPROX)*, pages 260–273, 2013.
- [34] M. Marchiori and V. Latora. Harmony in the small-world. *Physica A: Statistical Mechanics and its Applications*, 285(3-4):539–546, 2000.
- [35] F. McSherry and K. Talwar. Mechanism design via differential privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 94–103, 2007.
- [36] W. Meulemans, B. Speckmann, K. Verbeek, and J. Wulms. A framework for algorithm stability and its application to kinetic euclidean MSTs. In *Proceedings of the 13th Latin American Symposium on Theoretical Informatics (LATIN)*, pages 805–819, 2018.
- [37] S. Murai and Y. Yoshida. Sensitivity analysis of centralities on unweighted networks. In *Proceedings of the 2019 World Wide Web Conference (WWW)*, pages 1332–1342, 2019.
- [38] M. E. J. Newman. Fast algorithm for detecting community structure in networks. *Physical Review E*, 69(6):066133, 2004.
- [39] M. E. J. Newman. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23):8577–8582, 2006.
- [40] M. E. J. Newman. Network structure from rich but noisy data. *Nature Physics*, 14(6):542–545, 2018.
- [41] H. N. Nguyen and K. Onak. Constant-time approximation algorithms via local improvements. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 327–336, 2008.
- [42] K. Nissim, S. Raskhodnikova, and A. D. Smith. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC)*, pages 75–84, 2007.
- [43] L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999.
- [44] M. Parter, R. Rubinfeld, A. Vakilian, and A. Yodpinyanee. Local computation algorithms for spanners. In *Proceedings of 10th Innovations in Theoretical Computer Science Conference (ITCS)*, pages 58:1–58:21, 2019.

- [45] P. Peng and Y. Yoshida. Average sensitivity of spectral clustering. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2020. to appear.
- [46] S. Raskhodnikova and A. D. Smith. Lipschitz extensions for node-private graph statistics and the generalized exponential mechanism. In *Proceedings of the IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 495–504, 2016.
- [47] O. Reingold and S. Vardi. New techniques and tighter bounds for local computation algorithms. *J. Comput. Syst. Sci.*, 82(7):1180–1200, 2016.
- [48] R. Rubinfeld, G. Tamir, S. Vardi, and N. Xie. Fast local computation algorithms. In *Proceedings of the 1st Symposium on Innovations in Computer Science (ICS)*, pages 223–238, 2011.
- [49] G. Sabidussi. The centrality index of a graph. *Psychometrika*, 31(4):581–603, 1966.
- [50] S. Shalev-Shwartz and S. Ben-David. *Understanding Machine Learning. From Theory to Algorithms*. Cambridge University Press, Cambridge, 2009.
- [51] N. Varma and Y. Yoshida. Average sensitivity of graph algorithms. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 684–703. SIAM, 2021.
- [52] Y. Yoshida, M. Yamamoto, and H. Ito. Improved constant-time approximation algorithms for maximum matchings and other optimization problems. *SIAM Journal on Computing*, 41(4):1074–1093, 2012.

## A Average Sensitivity of Exponential Mechanism

In this section, we prove Lemma 2.1.

*Proof of Lemma 2.1.* Let  $t > 0$  be a parameter. Note that any index  $i \in [n]$  with  $\mathbf{x}(i) > \text{OPT} + \log n/\eta + t/\eta$  has probability at most  $e^{-t}/n$  of being sampled by  $\mathcal{A}$ . Hence, by a union bound, for every  $t > 0$

$$\Pr_{i \sim \mathcal{A}(\mathbf{x})} \left[ \mathbf{x}(i) \geq \text{OPT} + \frac{\log n}{\eta} + \frac{t}{\eta} \right] \leq e^{-t}.$$

Next, we analyze the distance between the output distributions. Let  $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^n$  be vectors, and let  $Z = \sum_{i \in [n]} e^{-\eta \mathbf{x}(i)}$  and  $Z' = \sum_{i \in [n]} e^{-\eta \mathbf{x}'(i)}$ . Without loss of generality, we assume that  $Z \geq Z'$ . First, note that for all  $i \in [n]$  such that  $\mathbf{x}(i) \geq \mathbf{x}'(i)$ , we have

$$0 \leq e^{-\eta \mathbf{x}'(i)} - e^{-\eta \mathbf{x}(i)} = e^{-\eta \mathbf{x}'(i)} \left( 1 - e^{-\eta(\mathbf{x}(i) - \mathbf{x}'(i))} \right) \leq \eta e^{-\eta \mathbf{x}'(i)} (\mathbf{x}(i) - \mathbf{x}'(i)).$$

Hence for any  $i \in [n]$ , we have

$$\begin{aligned} |e^{-\eta \mathbf{x}(i)} - e^{-\eta \mathbf{x}'(i)}| &\leq \max \left\{ \eta e^{-\eta \mathbf{x}(i)} (\mathbf{x}'(i) - \mathbf{x}(i)), \eta e^{-\eta \mathbf{x}'(i)} (\mathbf{x}(i) - \mathbf{x}'(i)) \right\} \\ &\leq \eta |\mathbf{x}(i) - \mathbf{x}'(i)| \max \{ e^{-\eta \mathbf{x}(i)}, e^{-\eta \mathbf{x}'(i)} \} \leq \eta |\mathbf{x}(i) - \mathbf{x}'(i)| \left( e^{-\eta \mathbf{x}(i)} + e^{-\eta \mathbf{x}'(i)} \right). \end{aligned}$$

Then, we have

$$\begin{aligned}
& \frac{1}{Z} \sum_{i \in [n]} |e^{-\eta \mathbf{x}(i)} - e^{-\eta \mathbf{x}'(i)}| \leq \frac{\eta}{Z} \sum_{i \in [n]} |\mathbf{x}(i) - \mathbf{x}'(i)| \left( e^{-\eta \mathbf{x}(i)} + e^{-\eta \mathbf{x}'(i)} \right) \\
& \leq \frac{\eta}{Z} \max_{i \in [n]} |\mathbf{x}(i) - \mathbf{x}'(i)| \sum_{i \in [n]} \left( e^{-\eta \mathbf{x}(i)} + e^{-\eta \mathbf{x}'(i)} \right) = \frac{\eta(Z + Z')}{Z} \max_{i \in [n]} |\mathbf{x}(i) - \mathbf{x}'(i)| \\
& \leq 2\eta \|\mathbf{x} - \mathbf{x}'\|_1. \tag{12}
\end{aligned}$$

Then, the total variation distance between  $\mathcal{A}(\mathbf{x})$  and  $\mathcal{A}(\mathbf{x}')$  is at most

$$\begin{aligned}
& \sum_{i \in [n]} \left| \frac{\exp(-\eta \mathbf{x}(i))}{Z} - \frac{\exp(-\eta \mathbf{x}'(i))}{Z'} \right| = \sum_{i \in [n]} \left| \frac{\exp(-\eta \mathbf{x}(i))}{Z} - \frac{\exp(-\eta \mathbf{x}'(i))}{Z} \left( \frac{Z - Z' + Z'}{Z'} \right) \right| \\
& = \sum_{i \in [n]} \left| \frac{\exp(-\eta \mathbf{x}(i))}{Z} - \frac{\exp(-\eta \mathbf{x}'(i))}{Z} - \frac{\exp(-\eta \mathbf{x}'(i))}{Z} \left( \frac{Z - Z'}{Z'} \right) \right| \\
& \leq \frac{1}{Z} \sum_{i \in [n]} |e^{-\eta \mathbf{x}(i)} - e^{-\eta \mathbf{x}'(i)}| + \frac{Z - Z'}{Z} \frac{1}{Z'} \sum_{i \in [n]} \exp(-\eta \mathbf{x}'(i)) \\
& \leq \frac{2}{Z} \sum_{i \in [n]} |e^{-\eta \mathbf{x}(i)} - e^{-\eta \mathbf{x}'(i)}| \leq 4\eta \|\mathbf{x} - \mathbf{x}'\|_1. \quad \square
\end{aligned}$$

## B Average Sensitivity of Prim's algorithm

In this section, we show that Prim's algorithm (with a simple tie-breaking rule, as described in Algorithm 7) has high average sensitivity even on unweighted graphs. This is in contrast to the low average sensitivity of Kruskal's algorithm that we discussed in Section 3.

---

### Algorithm 7: PRIM'S ALGORITHM

---

**Input:** undirected graph  $G = ([n], E)$

- 1 Let  $T \leftarrow \{1\}$ ;
  - 2 **while** there exists a vertex not spanned by  $T$  **do**
  - 3     Let  $E'$  be the set of edges with the smallest weight among all the edges in  $E$  that have exactly one endpoint in  $T$ ;
  - 4     Add to  $T$ , an edge from  $E'$  that has lexicographically smallest  $T$ -endpoint among all edges in  $E'$ , breaking further ties arbitrarily.
  - 5 **return** Output  $T$ .
- 

**Lemma B.1.** *The average sensitivity of Prim's algorithm is  $\Omega(m)$ .*

*Proof.* Consider the graph family  $\{G_n\}_{n \in 2\mathbb{N}}$  in Figure 1. For a large enough  $n \in 2\mathbb{N}$ , consider running Algorithm 7 on  $G_n$ . The tree  $T$  output will consist of the edges  $(i, i+1)$  for all  $i \in [n/2 - 2]$ , the edges  $(n/2 - 1, j)$  for all  $j \in \{n/2 + 1, \dots, n\}$ , and the edge  $(n/2, 1)$ .

If we remove an edge  $(i', i' + 1)$  for  $i' \in [n/2 - 2]$  from  $G_n$  and run Algorithm 7 on the resulting graph, the tree, say  $T_{i'}$ , output will consist of all edges of the form  $(i, i + 1)$  for  $i \in [n/2 - 1] \setminus \{i'\}$ ,




Figure 1: The graph family  $\{G_n\}_{n \in 2\mathbb{N}}$ .

all edges of the form  $(n/2, j)$  for all  $j \in \{n/2 + 1, \dots, n\}$ , and the edges  $(n/2 + 1, n/2 - 1)$  and  $(n/2, 1)$ . The Hamming distance of  $T_{i'}$  from  $T$  is equal to  $n/2$ .

Since a uniformly random edge removed from  $G_n$  is of the form  $(i, i + 1)$  for  $i \in [n/2 - 2]$  with probability  $\frac{n/2-2}{3n/2-1}$ , the average sensitivity of Algorithm 7 is at least  $\frac{n}{2} \cdot \frac{n/2-2}{3n/2-1}$ , which is at least  $\frac{n}{6} - 1 = \Omega(m)$  for the family  $\{G_n\}_{n \in 2\mathbb{N}}$ .  $\square$