


Sensitivity Analysis of the Maximum Matching Problem

Yuichi Yoshida 

National Institute of Informatics, Tokyo, Japan
yyoshida@nii.ac.jp

Samson Zhou 

Carnegie Mellon University, Pittsburgh, PA, USA
samsonzhou@gmail.com

Abstract

We consider the *sensitivity* of algorithms for the maximum matching problem against edge and vertex modifications. When an algorithm A for the maximum matching problem is deterministic, the sensitivity of A on G is defined as $\max_{e \in E(G)} |A(G) \Delta A(G - e)|$, where $G - e$ is the graph obtained from G by removing an edge $e \in E(G)$ and Δ denotes the symmetric difference. When A is randomized, the sensitivity is defined as $\max_{e \in E(G)} d_{EM}(A(G), A(G - e))$, where $d_{EM}(\cdot, \cdot)$ denotes the earth mover's distance between two distributions. Thus the sensitivity measures the difference between the output of an algorithm after the input is slightly perturbed. Algorithms with low sensitivity, or *stable* algorithms are desirable because they are robust to edge failure or attack.

In this work, we show a randomized $(1 - \epsilon)$ -approximation algorithm with *worst-case* sensitivity $O_\epsilon(1)$, which substantially improves upon the $(1 - \epsilon)$ -approximation algorithm of Varma and Yoshida (SODA'21) that obtains *average* sensitivity $n^{O(1/(1+\epsilon^2))}$ sensitivity algorithm, and show a deterministic $1/2$ -approximation algorithm with sensitivity $\exp(O(\log^* n))$ for bounded-degree graphs. We then show that any deterministic constant-factor approximation algorithm must have sensitivity $\Omega(\log^* n)$. Our results imply that randomized algorithms are strictly more powerful than deterministic ones in that the former can achieve sensitivity independent of n whereas the latter cannot. We also show analogous results for vertex sensitivity, where we remove a vertex instead of an edge.

Finally, we introduce the notion of normalized weighted sensitivity, a natural generalization of sensitivity that accounts for the weights of deleted edges. For a graph with weight function w , the normalized weighted sensitivity is defined to be the sum of the weighted edges in the symmetric difference of the algorithm normalized by the altered edge, i.e., $\max_{e \in E(G)} \frac{1}{w(e)} w(A(G) \Delta A(G - e))$. Hence the normalized weighted sensitivity measures the weighted difference between the output of an algorithm after the input is slightly perturbed, normalized by the weight of the perturbation. We show that if all edges in a graph have polynomially bounded weight, then given a trade-off parameter $\alpha > 2$, there exists an algorithm that outputs a $\frac{1}{4\alpha}$ -approximation to the maximum weighted matching in $O(m \log_\alpha n)$ time, with normalized weighted sensitivity $O(1)$.

2012 ACM Subject Classification Theory of computation \rightarrow Graph algorithms analysis; Theory of computation \rightarrow Approximation algorithms analysis; Theory of computation \rightarrow Models of computation

Keywords and phrases Sensitivity analysis, maximum matching, graph algorithms

Digital Object Identifier 10.4230/LIPIcs.ITCS.2021.58

Related Version <https://arxiv.org/pdf/2009.04556.pdf>

Funding *Yuichi Yoshida*: Supported by JST, PRESTO Grant Number JPMJPR192B, Japan.
Samson Zhou: Supported by a Simons Investigator Award of David P. Woodruff.

Acknowledgements We would like to thank anonymous ITCS reviewers for helpful comments on an earlier version of the paper.



© Yuichi Yoshida and Samson Zhou;
licensed under Creative Commons License CC-BY
12th Innovations in Theoretical Computer Science Conference (ITCS 2021).
Editor: James R. Lee; Article No. 58; pp. 58:1–58:20



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

The problem of finding the maximum matching in a graph is a fundamental problem in graph theory with a wide range of applications in computer science. For example, the maximum matching problem on a bipartite graph G captures a typical example where a number of possible clients want to access content distributed across multiple providers. Each client can download their specific content from a specific subset of the possible providers, but each provider can only connect to a limited number of clients. A maximum matching between clients and providers would ensure that the largest possible number of clients receive their content.

However in many modern applications, the underlying graph G represents some large dataset that is often dynamic or incomplete. In the above example, the content preference of clients may change, which alters the set of suppliers that provide their desired content. Connections between specific providers and clients may become online or offline, effectively adding or removing edges in the underlying graph. Providers and clients may themselves join or leave the network, adding or removing entire vertices from the graph. Thus, it is reasonable to assume that our knowledge of important properties of G may also change or be incomplete. Nevertheless, we must extract information from our current knowledge of G either for pre-processing or to perform tasks on the current infrastructure. At the same time, we would like to maintain as much consistency as possible when updates to G are revealed.

Motivated by a formal definition of consistency of algorithms across graph updates, Varma and Yoshida [12] first defined the *average sensitivity* of a deterministic algorithm A to be the Hamming distance¹ between the output of A on graphs G and $G - e$, where G' is the graph formed by deleting a random edge of G . Then, they defined *average sensitivity* for randomized algorithms as

$$\mathbb{E}_{e \sim E(G)} [d_{EM}(A(G), A(G - e))],$$

where $d_{EM}(\cdot, \cdot)$ denotes the earth mover's distance and $G - e$ is the graph obtained from G by deleting an edge $e \in E(G)$. For the maximum matching problem, they showed a randomized $1/2$ -approximation algorithm with average sensitivity $O(1)$ and a randomized $(1 - \epsilon)$ -approximation algorithm with average sensitivity $O(n^{1/(1+\epsilon^2)})$.

Worst case sensitivity. In this work, we continue the study of sensitivity for the maximum matching problem. Instead of average sensitivity as in [12], we consider a stronger notion of (worst-case) sensitivity. Specifically, the *sensitivity* of a deterministic algorithm A is the maximum Hamming distance between the output of A on graphs G and G' , where G' is the graph formed by deleting an edge of G . Then, the *sensitivity* of a randomized algorithm A is

$$\max_{e \in E(G)} d_{EM}(A(G), A(G - e)).$$

Clearly, the sensitivity of an algorithm is no smaller than its average sensitivity. As a natural variant, we also consider *vertex sensitivity*, where we delete a vertex instead of an edge. To avoid confusion, sensitivity with respect edge deletion will be sometimes called *edge sensitivity*.

¹ Here we regard the output as a binary string so we can think of the Hamming distance between outputs.

1.1 Our Contributions

We first show that, for any $\epsilon > 0$, there exists a randomized $(1 - \epsilon)$ -approximation algorithm whose sensitivity solely depends on ϵ (Section 2).

► **Theorem 1.** *For any $\epsilon > 0$, there exists an algorithm that outputs a $(1 - \epsilon)$ -approximation to the maximum matching problem with probability at least 0.99, using time complexity $O((n + m) \cdot K)$ and edge/vertex sensitivity $O(3^K)$, where $K = (1/\epsilon)^{2^{O(1/\epsilon)}}$.*

This result improves upon the previous $(1 - \epsilon)$ -approximation algorithm [12] in that (1) the sensitivity is constant instead of $O(n^{1/(1+\epsilon^2)})$ and (2) it bounds worst-case sensitivity instead of average sensitivity.

We observe that approximation is necessary to achieve a small sensitivity. For example, consider an n -cycle for an even n , and let M_1 and M_2 be the two maximum matchings of size $n/2$ in the graph. Consider a deterministic algorithm that always outputs a maximum matching, say, M_1 for the n -cycle. Then, it must output M_2 after removing an edge in M_1 , and hence the sensitivity is $\Omega(n)$. With a similar reasoning, we can show a lower bound of $\Omega(n)$ for randomized algorithms. Also as we show in Section 4.3, the dependency on ϵ in Theorem 1 is necessary.

Next, we show a deterministic algorithm for finding a maximal matching on bounded degree graphs that has low sensitivity (Section 3). Note that it has approximation ratio $1/2$ because the size of any maximal matching is a $1/2$ -approximation to the maximum matching.

► **Theorem 2.** *There exists a deterministic algorithm that finds a maximal matching with edge/vertex sensitivity $\Delta^{O(6^\Delta + \log^* n)}$, where Δ is the maximum degree of a vertex in the graph.*

Then, we show that randomness is necessary to achieve sensitivity independent of n (Section 4):

► **Theorem 3.** *Any deterministic constant-factor approximation algorithm for the maximum matching problem has edge sensitivity $\Omega(\log^* n)$.*

Namely, we show in Section 4.2 that we cannot obtain sublinear sensitivity just by derandomizing the randomized greedy algorithm. Theorems 1 and 3 imply that randomized algorithms are strictly more powerful than deterministic ones in that the former can achieve sensitivity independent of n whereas the latter cannot for the maximum matching problem.

We then introduce the idea of *weighted sensitivity*, which is a natural generalization of sensitivity for both the average and worst cases. For the problems that we consider, the sensitivity of a deterministic graph algorithm is the number of edges that changes in the output induced by the alteration of a single vertex/edge in the input. Thus for a weighted graph, the weighted sensitivity is the total weight of the edges that are changed in the output, following the deletion of a vertex/edge. For randomized algorithms, the definition extends naturally to the earth mover's distance between the distributions with the corresponding weighted loss function. Finally, we can also normalize by the weight of the edge that is deleted.

The motivation for studying weighted sensitivity is natural; in many applications with evolving data, the notion of sensitivity arises in the context of recourse, a quantity that measures the change in the underlying topology of the optimal solution. For example in the facility location problem, the goal is to construct a set of facilities to minimize the sum of

the costs of construction and service to a set of consumers. As the information about the set of consumers evolves, it would be ideal to minimize the number of relocations for the facilities, due to the construction costs, which is measured by the sensitivity of the algorithm. However, as construction costs may not be uniform, a more appropriate quantity to minimize would be the total cost of the relocations for the facilities, which is measured by the weighted sensitivity.

Similarly, matchings are often used to maximize flow across a bipartite graph, but the physical structures that support the flow may incur varying costs to construct or demolish, corresponding to the amount of flow that the structures support. In this case, we note that it may not be possible for the worst case weighted sensitivity to be small. For example, if a single edge has weight n^C for some large constant C and the remaining edges have weight 1, any constant factor approximation to the maximum weighted matching must include the heavy edge. But if the heavy edge is then removed from the graph, the weighted sensitivity of any constant factor approximation algorithm is $\Omega(n^C)$. This issue is circumvented by the normalized weighted sensitivity, which scales the sensitivity by the weight of the deleted edge. We give approximation algorithms for maximum weighted matching with low normalized weighted worst-case sensitivity.

► **Theorem 4.** *Let $G = (V, E)$ be a weighted graph with $\frac{1}{n^c} \leq w(e) \leq n^c$ for some constant $c > 0$ and all $e \in E$. For a trade-off parameter $\alpha > 2$, there exists an algorithm that outputs a $\frac{1}{4\alpha}$ -approximation to the maximum weighted matching in $O(m \log_\alpha n)$ time and has normalized weighted sensitivity $O(1)$.*

Our results also extend to $\alpha = 2$ and general worst-case weighted sensitivity, i.e., weighted sensitivity that is not normalized. We detail these algorithms in Section 5.

1.2 Proof Sketch

We explain the idea behind the algorithm of Theorem 1. For simplicity, we focus on edge sensitivity. We note that if we only sought a $1/2$ -approximation to the maximum matching, then it would suffice to find any maximal matching. Although the well-known greedy algorithm produces a maximal matching, the output of the algorithm is highly sensitive to the ordering of the edges in the input. One may hope that, if we choose an ordering of the edges uniformly at random, then the resulting output will be stable against edge deletions to the underlying graph. This is not immediately obvious because the deleted edge will appear about halfway through the ordering (of the edges in the original graph) in expectation, so it seems possible that it can impact about the remaining half of the edges. Luckily, we show that the edges at the beginning of the ordering are significantly more important, so that even if the deleted edge appears about halfway through the ordering, the sensitivity of the maximal matching is $O(1)$ (Section 2.3.1). Our analysis is similar to [2], who show that the vertices at the beginning of an ordering are significantly more important in maintaining a maximal independent set in the dynamic distributed model.

Adapting this idea to a $(1 - \epsilon)$ -approximation is more challenging. The natural approach is to take a maximal matching and repeatedly find a large number of augmenting paths, but the change of even a single edge in a maximal matching can potentially impact a large number of edges if the augmenting paths are found in a sequential manner. We instead adapt a layered graph of [9] that is used to randomly find a large number of augmenting paths in a small number of passes in the streaming model. Crucially, we instead find a large number of disjoint augmenting paths in a small number of parallel rounds, which results in low sensitivity.

Now we turn to explaining the idea behind Theorem 2. Again we focus on edge sensitivity. Our algorithm first uses a deterministic local computation algorithm (LCA) of [4] for 6^Δ -coloring a graph G with maximum degree Δ , using $O(\Delta \log^* n)$ probes to an adjacency list oracle. Here we want to design an algorithm that answer queries about the colors of vertices by making a series of probes to the oracle. The answers of the algorithm must be consistent so that there exists at least one proper coloring that is consistent with the answers. In our case, each probe to the oracle is a query (v, i) with $v \in V$ and a positive integer i . If the degree of v is at least i , the oracle responds with the i -th neighbor of v to the probe. Otherwise, the oracle outputs a special symbol \perp . In particular, the deterministic 6^Δ -coloring LCA only probes vertices that are within a “small” neighborhood of the query.

Given a coloring for G , we then give a local distributed algorithm that takes a coloring of a graph and outputs a maximal matching. It follows from a framework of [10] that our local distributed algorithm can actually be simulated by a deterministic LCA that again only probes a “small” neighborhood of the query. Thus to bound the sensitivity of the algorithm, we bound the number of queries for which a deleted edge would be probed. Since only a small number of queries probes the deleted edge, then the output of the algorithm only has a small number of changes and thus low worst-case sensitivity.

Our lower bound of Theorem 3 considers the set of length- t cycles on a graph with n vertices. Any matching on length- t cycles can be represented as a series of indicator variables denoting whether edge $i \in [t]$ is in the matching. We can then interpret the indicator variables as an integer encoding from 0 to $2^t - 1$ through the natural binary representation. Ramsey theory claims that for $t = O(\log^* n)$, there exists a set S of $t + 1$ nodes of n so that any subset of t nodes has the same encoding. We then choose G and G' to be the cycle graphs consisting of the first t nodes of S and the last t nodes of S , respectively. Since the encodings of the matchings of G and G' are the same, but the edge indices are shifted by one, it follows that $\Omega(t)$ edges must be in the symmetric difference between G and G' , which implies from $t = O(\log^* n)$ that the worst-case sensitivity of the algorithm must be $\Omega(\log^* n)$.

1.3 Related Work

Varma and Yoshida [12] introduced the notion of sensitivity and performed a systematic study of average sensitivity on many graph problems. Namely, they gave efficient approximation algorithms with low average sensitivities for the minimum spanning forest problem, the global minimum cut problem, the minimum s - t cut problem, and the maximum matching problem. They also introduced a low-sensitivity algorithm for linear programming, and proved many fundamental properties of average sensitivity, such as sequential or parallel composition. Peng and Yoshida [11] gave an algorithm for the problem of spectral clustering with average sensitivity $\frac{\lambda_2}{\lambda_3}$, where λ_i is the i -th smallest eigenvalue of the normalized Laplacian, which is small when there are exactly two clusters in the graph.

The effects of graph updates have also been studied significantly in the dynamic/online model, where updates to the graph arrive in a stream, and the goal is to maintain some data structure to answer queries on the underlying graph so that both the update time and query time are efficient. Consequently, most of the literature for dynamic algorithms focuses on optimizing these quantities, rather than the changes in the output as the data evolves. Sensitivity analysis is more relevant when the goal of the dynamic/offline model is to minimize the number of changes between successive outputs of the algorithm over the stream.

Lattanzi and Vassilvitski [8] studied the problem of consistent k -clustering, where the goal is to maintain a constant-factor approximation to some underlying k -clustering problem, such as k -center, k -median, or k -means, while minimizing the total number of changes to the

set of centers as the stream evolves. In this setting, each change to the set of center is known as a *recourse*. Whereas the model of [8] allows only insertions of new points, algorithms with low sensitivity are robust against both insertions and deletions. Cohen-Addad *et. al.* [3] further considered the facility location problem in this model of maintaining a constant-factor approximation while minimizing the total recourse. Although the algorithm of [3] addresses both the insertions and deletions of points, their total recourse across the stream is $O(n)$, where n is the length of the stream; this is inherent to the difficulty of their problem in the model. Whereas their work already provides an amortized $O(1)$ recourse per update, we also study the worst-case sensitivity in our work.

1.4 Preliminaries

For a positive integer n , let $[n]$ denote the set $\{1, 2, \dots, n\}$. For a positive integer n and $p \in [0, 1]$, let $\mathcal{B}(n, p)$ be the binomial distribution with n trials and success probability p . We use the notation $O_\epsilon(\cdot)$ to omit dependencies on ϵ .

Let $G = (V, E)$ be a graph. For an edge $e \in E$, let $N_G(e)$ be the “neighboring” edges of e in G , that is, $N_G(e) = \{e' \in E \mid e' \neq e, |e' \cap e| \geq 1\}$. We omit the subscript if it is clear from the context.

For two (vertex or edge) sets S and S' , let $d_H(S, S') = |S \Delta S'|$, where Δ denotes the symmetric difference. Abusing the notation, for set of paths \mathcal{P} and \mathcal{P}' , we write $d_H(\mathcal{P}, \mathcal{P}')$ to denote $d_H(\cup_{P \in \mathcal{P}} V(P), \cup_{P \in \mathcal{P}'} V(P))$. For two random sets X and X' , let $d_{EM}(X, X')$ be the earth mover’s distance between X and X' , where the distance between two sets is measured by d_H , that is,

$$d_{EM}(X, X') = \min_{\mathcal{D}} \mathbf{E}_{(S, S') \sim \mathcal{D}} d_H(S, S'),$$

where \mathcal{D} is a distribution such that its marginal on the first and second coordinates are X and X' , respectively. For a real-valued function β on graphs, we say that the *sensitivity* of a (randomized) algorithm A that outputs a set of edges is at most β if

$$d_{EM}(A(G), A(G - e)) \leq \beta(G)$$

holds for every $e \in E(G)$.

Given a matching M in a graph $G = (V, E)$, we call a vertex *free* if it does not appear as the endpoint of any edge in M . A path $(v_1, v_2, \dots, v_{2\ell+2})$ of length $2\ell + 1$ is an *augmenting path* if v_1 and $v_{2\ell+2}$ are free vertices and $(v_i, v_{i+1}) \in M$ for even i and $(v_i, v_{i+1}) \in E \setminus M$ for odd i .

2 Randomized $(1 - \epsilon)$ -Approximation

In this section, we prove Theorem 1. Our algorithm, which we describe in Section 2.1, is a slight modification of the multi-pass streaming algorithm due to McGregor [9]. We discuss its approximation guarantee and sensitivity in Sections 2.2 and 2.3, respectively.

2.1 Algorithm Description


A key step of McGregor’s algorithm is to find a large set of augmenting paths of a specified length in a batch manner using the *layered graph*, given below. Given a graph $G = (V, E)$, a matching $M \subseteq E$, and a positive integer ℓ , the layered graph $H = H(G)$ consists of $\ell + 2$ layers $L_0, L_1, \dots, L_{\ell+1}$, where $L_0 = L_{\ell+1} = V$ and $L_1 = L_2 = \dots = L_\ell = V \times V$.

For each vertex $v \in V$, we sample $i_v \in \{0, \ell + 1\}$ uniformly at random independently from others. We say that the copy of v in the L_{i_v} -th layer is *active* and that the other copy is *inactive*. For each edge $\{u, v\} \in M$, with probability half, we sample a value $i_{(u,v)} \in \{1, \dots, \ell\}$ uniformly at random and set $i_{(v,u)} = \perp$, where \perp is a special symbol, and with the remaining probability half, we sample a value $i_{(v,u)} \in \{1, \dots, \ell\}$ uniformly at random and set $i_{(u,v)} = \perp$. For each edge $\{u, v\} \in E \setminus M$, we set $i_{(u,v)} = i_{(v,u)} = \perp$. We say that the copy of (u, v) in the $L_{i_{(u,v)}}$ -th is *active* if $i_{(u,v)} \neq \perp$ and is *inactive* otherwise. Intuitively, some orientation of each edge $\{u, v\}$ in the matching M is assigned to a random internal layer in H and edges of G that are not in the matching are not initially assigned to any layer in H . For $i = 0, \dots, \ell + 1$, we denote by \tilde{L}_i the set of active vertices in L_i . Let $L = \bigcup_{i=0}^{\ell+1} L_i$ be the vertex set of H , and let $\tilde{L} = \bigcup_{i=0}^{\ell+1} \tilde{L}_i$ be the set of active vertices in H .

The edges in the layered graph H are those between active vertices that can be a part of an augmenting path in G . More specifically,

- We add an edge between $t \in \tilde{L}_0$ and $(u, v) \in \tilde{L}_1$ if t is free in M and t is adjacent to v .
- We add an edge between $(u, v) \in \tilde{L}_\ell$ and $s \in \tilde{L}_{\ell+1}$ if s is free in M and s is adjacent to u .
- We add an edge between $(u, v) \in \tilde{L}_i$ and $(u', v') \in \tilde{L}_{i+1}$ for $i \in [\ell - 1]$ if v is adjacent to u' .

Note that inactive vertices are isolated in H .



■ **Figure 1** Example of an active layered graph with respect to a matching M , in solid lines. The free vertex 5 appears in L_0 and the free vertex 6 appears in L_3 . The augmenting path found by the layered graph is represented by a dashed purple line.

We introduce the following definition to handle augmenting paths for a matching M in a graph G via paths in its corresponding layered graph.

► **Definition 5.** We say that a path v_i, v_{i-1}, \dots, v_0 with $v_j \in \tilde{L}_j$ ($j \in \{0, 1, \dots, i\}$) is an i -path. Note that an $(\ell + 1)$ -path in H corresponds to an augmenting path of length $2\ell + 1$ in G .

The layered graph defined above is slightly different from the original one due to McGregor [9] in that he did not include inactive vertices in H , as they are irrelevant to find augmenting paths. However as we consider sensitivity of algorithms, it is convenient to fix the vertex set so that it is independent of the current matching M .

We briefly define the randomized greedy subroutine RANDOMIZEDGREEDY on a graph $G = (V, E)$ as follows. The subroutine first chooses a random ordering π over edges and then starting with an empty matching M , the procedure iteratively adds the i -th edge in the ordering π to M if the edge is not adjacent to any edge in M , until it has processed all edges. See Algorithm 1 for the full details.

■ **Algorithm 1** Randomized Greedy Algorithm.

```

1 Procedure RANDOMIZEDGREEDY(Graph  $G = (V, E)$ )
2   Generate a permutation  $\pi$  of  $E$  uniformly at random;
3   Greedily add edges to a maximal matching  $M$ , in the order of  $\pi$ ;
4   Output  $M$ ;

```

Algorithm 2 shows our algorithm for finding a large set of augmenting paths of length $2\ell + 1$ given a matching M in a graph G . For a matching M and a vertex v belonging to an edge e in M , let $\Gamma_M(v)$ denote the other endpoint of e . Similarly, for a vertex set S such that each edge in M uses at most one vertex in S , let $\Gamma_M(S)$ denote the set of other endpoints. For subsets L and R of adjacent layers in H , let RANDOMIZEDGREEDY(L, R) denote the randomized greedy on the induced bipartite graph $H[L \cup R]$. FINDPATHS tries to find a large set of vertex disjoint i -paths from $S \subseteq L_i$ to L_0 . The result is stored as a tag function $t : V(H) \rightarrow V(H) \cup \{\text{untagged}, \text{dead end}\}$. Here, $t(v)$ is initialized to *untagged*, and it will represent the next vertex in the i -path found. If we could not find any i -path starting from v , $t(v)$ is set to *deadend*.

The difference from McGregor’s algorithm is that we run the loop in FINDPATHS $1/\delta$ times instead of running it until $|M'| \leq \delta|M|$. This makes sure that we compute a maximal matching the same number of times no matter what G and M are, and it is more convenient when analyzing the sensitivity.

Our algorithm for the maximum matching problem (Algorithm 3) simply runs AUGMENTINGPATHS sufficiently many times for various choice of ℓ and then keep applying the obtained augmenting paths. Before analyzing its approximation ratio and sensitivity, we analyze its running time.

► **Lemma 6.** *The total running time of Algorithm 3 is $O(n + m)K$, where $K = (1/\epsilon)^{2^{O(1/\epsilon)}}$.*

Proof. Observe that the outer loop of Algorithm 3 runs for r iterations and the inner loop runs for k iterations, where $k = \lceil \epsilon^{-1} + 1 \rceil$ and $r = 4k^2(16k + 20)(k - 1)(2k)^k$. Each inner loop runs an instance of AUGMENTINGPATHS with parameters $\ell \leq k$ and $\delta = \frac{1}{r(2k+2)}$, which creates a layered graph with $O(\ell)$ layers in $O((n + m)k)$ time, and then calls FINDPATHS. For each time that FINDPATHS is called, the value of δ is squared and the value of ℓ is decremented, starting at $\ell = k$ until $\ell = 1$. Thus, the loop in FINDPATHS is run at most $\frac{1}{\delta} = O\left((r(2k + 2))^{2^k}\right)$ times and each loop uses time $O(n + m)$. Hence, the total runtime is $O(n + m)K$, where $K = (1/\epsilon)^{2^{O(1/\epsilon)}}$. ◀

2.2 Approximation Ratio

In this section, we analyze the approximation ratio of Algorithm 3.

► **Lemma 7** ([9]). *Suppose FINDPATHS(\cdot, \cdot, j, \cdot) is called $\delta^{-2^{i-j+1}+1}$ times in the recursion for FINDPATHS(H, S, i, δ). Then at most $2\delta|L_i|$ paths are removed from consideration as being $(i + 1)$ -paths.*

Algorithm 2 Augmentation Algorithm.

```

1 Procedure AUGMENTINGPATHS( $G, M, \ell, \delta$ )
2   Construct a layered graph  $H$  using  $G, M$ , and  $\ell$ ;
3    $t(v) \leftarrow$  untagged for every vertex  $v$  in  $H$ ;
4   FINDPATHS( $H, L_{\ell+1}, \ell, \delta, t$ );
5   Convert  $t$  to a set of augmenting paths in  $G$ ;
6   Apply the augmenting paths to  $M$ .
7 Procedure FINDPATHS( $H, S, i, \delta, t$ )
8    $M' \leftarrow$  RANDOMIZEDGREEDY( $S, L_{i-1} \cap t^{-1}(\text{untagged})$ );
9    $S' \leftarrow \Gamma_{M'}(S)$ ;
10  if  $i = 1$  then
11    for  $u \in S$  do
12      if  $u \in \Gamma_{M'}(L_0)$  then
13         $t(u) \leftarrow \Gamma_{M'}(u)$ 
14      if  $u \in S \setminus \Gamma_{M'}(L_0)$  then
15         $t(u) \leftarrow$  dead end.
16    return  $t$ .
17  for  $\lceil 1/\delta \rceil$  times do
18    FINDPATHS( $H, S', i - 1, \delta^2, t$ );
19    for  $v \in S' \setminus t^{-1}(\text{dead end})$  do
20       $t(\Gamma_{M'}(v)) \leftarrow v$ .
21     $M' \leftarrow$  RANDOMIZEDGREEDY( $S \cap t^{-1}(\text{untagged}), L_{i-1} \cap t^{-1}(\text{untagged})$ );
22     $S' \leftarrow \Gamma_{M'}(S \cap t^{-1}(\text{untagged}))$ .
23  for  $v \in S \cap t^{-1}(\text{untagged})$  do
24     $t(v) \leftarrow$  dead end.
25  return

```

Algorithm 3 Algorithm for maximum matching.

```

1 Procedure MATCHING( $G, \epsilon$ )
2   Let  $\pi$  be a random ordering of the edges of  $G$ ;
3   Let  $M$  be the greedy maximal matching of  $G$  induced by  $\pi$ ;
4    $k \leftarrow \lceil \epsilon^{-1} + 1 \rceil$ ;
5    $r \leftarrow 4k^2(16k + 20)(k - 1)(2k)^k$ ;
6   for  $\ell = 1$  to  $k$  do
7     for  $i = 1$  to  $r$  do
8        $M'_{\ell,i} \leftarrow$  AUGMENTINGPATHS( $G, M, \ell, \frac{1}{r(2k+2)}$ );
9        $M \leftarrow M \oplus M'_{\ell,i}$ .
10  return  $M$ .

```

Let \mathcal{L}_i be the set of graphs whose vertices are partitioned into $i + 2$ layers, L_0, \dots, L_{i+1} and whose edges are a subset of $\bigcup_{j=1}^{i+1} (L_j \times L_{j-1})$. Then we immediately have the following lemma, analogous to Lemma 2 in [9]:

► **Lemma 8.** For a graph $G \in \mathcal{L}_i$, $\text{FINDPATHS}(H, S, i, \delta)$ finds at least $(\gamma - \delta)|M|$ of the $(i + 1)$ -paths among some maximal set of $(i + 1)$ -paths of size $\gamma|M|$.

We require the following structural property relating maximal and maximum matchings through the set of connected components in the symmetric difference.

► **Lemma 9** (Lemma 1 in [9]). Let M be a maximal matching and M^* be a maximum matching. Let \mathcal{C} be the set of connected components in $M^* \Delta M$. Let α_ℓ be the constant so that $\alpha_\ell|M|$ is the number of connected components in \mathcal{C} with ℓ edges from M , excluding those with ℓ edges from M^* . If $\max_{\ell \in [k]} \alpha_\ell \leq \frac{1}{2k^2(k+1)}$, then $|M| \geq \frac{|M^*|}{1+1/k}$.

We also require the following result by [9] bounding the number of augmenting paths found by AUGMENTINGPATHS .

► **Lemma 10** (Theorem 1 in [9]). If G has $\alpha_\ell|M|$ augmenting paths of length $2\ell + 1$, then the number of augmenting paths of length $2\ell + 1$ found by AUGMENTINGPATHS is at least $(b_\ell\beta_\ell - \delta)|M|$, where $b_\ell = \frac{1}{2\ell+1}$ and $\beta_\ell \sim \mathcal{B}\left(\alpha_\ell|M|, \frac{1}{2(2\ell)^\ell}\right)$.

We now show that Algorithm 3 outputs a $(1 - \epsilon)$ -approximation to the maximum matching.

► **Theorem 11.** Algorithm 3 finds a $(1 - \epsilon)$ -approximation to the maximum matching with probability at least 0.99.

Proof. We say the algorithm enters *phase* ℓ when the number of layers in the layered graph has been incremented to ℓ , i.e., each invocation of the outer for loop corresponds to a separate phase. We say the algorithm enters *round* i in phase ℓ after the subroutine AUGMENTINGPATHS has completed $i - 1$ iterations within phase ℓ . Let $M_{\ell,i}$ be the matching M prior to the call to AUGMENTINGPATHS in round i of phase ℓ . Let $\alpha_{\ell,i}|M_{\ell,i}|$ be the number of length $2\ell + 1$ augmenting paths of $M_{\ell,i}$. Thus by Lemma 10, the subroutine AUGMENTINGPATHS augments $M_{\ell,i}$ by at least $(b_\ell\beta_{\ell,i} - \delta)|M_{\ell,i}|$ edges in round i of phase ℓ , where $b_\ell = \frac{1}{2\ell+2}$, δ is a parameter that we choose later, and $\beta_{\ell,i}$ is a random variable distributed according to $\mathcal{B}\left(\alpha_{\ell,i}|M_{\ell,i}|, \frac{1}{2(2\ell)^\ell}\right)$. Let M be the output matching. Then by Bernoulli's inequality, we have

$$\begin{aligned} \Pr[|M| \geq 2|M_{1,1}|] &\geq \Pr\left[|M_{1,1}| \prod_{\ell \in [k], i \in [r]} (1 + \max(0, b_\ell\beta_{\ell,i} - \delta)) \geq 2|M_{1,1}| \right] \\ &\geq \Pr\left[\sum_{\ell \in [k]} \max_{i \in [r]} b_\ell\beta_{\ell,i} \geq 2 + r\delta \right]. \end{aligned}$$

We would like to analyze $\sum_{\ell \in [k]} \max_{i \in [r]} b_\ell\beta_{\ell,i}$, but the analysis is challenging due to dependencies between multiple rounds and phases. We thus define independent variables X_1, \dots, X_k and use a coupling argument.

We define $A_\ell = \max_{i \in [r]} b_\ell\beta_{\ell,i}|M_{\ell,i}|$ to be an upper bound on the maximum number of augmented edges during phase ℓ of the algorithm. Suppose by way of contradiction that $\max_{i \in [r]} \alpha_{\ell,i} < \alpha_0 := \frac{1}{2k^2(k-1)}$ for each of the phases $1 \leq \ell \leq k$. Then Lemma 9 would imply that at some point $M_{\ell,i}$ is sufficiently large. Thus we have $\max_{i \in [r]} \alpha_{\ell,i} \geq \alpha_0$.

We have $A_\ell = \max_{i \in [r]} b_\ell\beta_{\ell,i}|M_{\ell,i}|$, $b_\ell = \frac{1}{2\ell+2}$, and $\beta_{\ell,i}|M_{\ell,i}| \sim \mathcal{B}\left(\alpha_{\ell,i}|M_{\ell,i}|, \frac{1}{2(2\ell)^\ell}\right)$. Now for $\ell \leq k$, we have that $\frac{1}{2(2\ell)^\ell} \geq \frac{1}{2(2k)^k}$. Thus $\max_{i \in [r]} \alpha_{\ell,i} \geq \alpha_0$ and $|M_{\ell,i}| \geq |M_{\ell,1}|$ implies that the distribution of A_ℓ statistically dominates the distribution of $b_\ell \cdot \mathcal{B}\left(\alpha_0|M_{\ell,1}|, \frac{1}{2(2k)^k}\right)$.

Hence if we define X_ℓ to be independent random variables distributed as $\mathcal{B}\left(\alpha_0|M_{\ell,1}|, \frac{1}{2(2k)^k}\right)$ for each $\ell \in [r]$, then the distribution of A_ℓ statistically dominates the distribution of $b_\ell \cdot X_\ell$. Thus,

$$\begin{aligned} \Pr\left[\sum_{\ell \in [k]} \max_{i \in [r]} b_\ell \beta_{\ell,i} \geq 2 + r\delta\right] &\geq \Pr\left[\sum_{\ell \in [k]} \max_{i \in [r]} b_\ell \beta_{\ell,i} |M_{\ell,i}| \geq (2 + r\delta)|M_{k,r}|\right] \\ &= \Pr\left[\sum_{\ell \in [k]} A_\ell \geq (2 + r\delta)|M_{k,r}|\right] \\ &\geq \Pr\left[\sum_{\ell \in [k]} X_\ell \geq \frac{4 + 2r\delta}{b_k} \cdot |M_{1,1}|\right], \end{aligned}$$

where the final inequality results from $M_{1,1}$ being a maximal matching and $b_\ell \geq b_k$ for $\ell \in [k]$.

Now the variables X_ℓ are independent but not identically distributed. Nevertheless, we can write $Y = \sum_{\ell \in [k]} X_\ell$ and note that the distribution of Y statistically dominates the distribution of $Z \sim \mathcal{B}\left(\alpha_0|M_{1,1}|r, \frac{1}{2(2k)^k}\right)$ since $|M_{\ell,i}| \geq |M_{1,1}|$ for all $\ell \in [k]$ and $i \in [r]$. Thus for $b_k = \frac{1}{2k+2}$ and $\delta = \frac{b_k}{r}$,

$$\begin{aligned} \Pr\left[\sum_{\ell \in [k]} X_\ell \geq \frac{4 + 2r\delta}{b_k} \cdot |M_{1,1}|\right] &\geq \Pr\left[Z \geq \frac{4 + 2b_k}{b_k} |M_{1,1}|\right] \\ &= \Pr[Z \geq |M_{1,1}|(8k + 10)]. \end{aligned}$$

For $r = \frac{2(2k)^k(16k+20)}{\alpha_0}$, we have that $\mathbf{E}[Z] = (16k + 20)|M_{1,1}|$. Thus from a simple Chernoff bound,

$$\Pr[Z \geq |M_{1,1}|(8k + 10)] = 1 - \Pr\left[Z < \frac{\mathbf{E}[Z]}{2}\right] > 1 - e^{-2(16k+20)|M_{1,1}|} \geq 0.99.$$

Putting things together, we have $\Pr[|M| \geq 2|M_{1,1}|] \geq 0.99$, which implies that there exists a maximum matching with more than double the number of edges of a maximal matching and is a contradiction. Therefore, our assumption that $\max_{i \in [r]} \alpha_{\ell,i} \geq \alpha_0 := \frac{1}{2k^2(k-1)}$ for each of the phases $1 \leq \ell \leq k$ must have been invalid. However, if $\alpha_{\ell,i} < \frac{1}{2k^2(k-1)}$ for some $i \in [r]$ and $\ell \in [k]$, then by Lemma 9, we have for $k = \lceil \epsilon^{-1} + 1 \rceil$ and sufficiently small $\epsilon > 0$ that

$$|M_{\ell,i}| \geq \frac{|M^*|}{1 + 1/k} \geq \frac{|M^*|}{1 + \epsilon} \geq (1 - \epsilon)|M^*|,$$

with probability at least 0.99. Thus, Algorithm 3 outputs a $(1 - \epsilon)$ approximation the maximum matching with probability at least 0.99. \blacktriangleleft

Boosting the Success Probability. To increase the probability of success to $1 - p$ for any $p \in (0, 1)$, a naïve approach would be to run $O\left(\log \frac{1}{p}\right)$ iterations of Algorithm 3 in parallel. However, the sensitivity analysis becomes considerably more challenging. Instead, we note that the $\frac{1}{e}$ probability of failure is actually a significant weakening of the $e^{-2(16k+20)|M_{1,1}|}$ probability of failure. Thus, increasing k by a factor of $O\left(\log \frac{1}{p}\right)$ increases the probability of success to $1 - p$. However, for subconstant p , it also substantially increases the asymptotic sensitivity of Algorithm 3.

2.3 Sensitivity of the Randomized Greedy and Algorithm 3

To analyze the sensitivity of Algorithm 3, we first analyze the sensitivity of the randomized greedy algorithm.

2.3.1 Sensitivity of the Randomized Greedy

In this section, we study the sensitivity of the randomized greedy with respect to vertex deletions. Recall that given a graph $G = (V, E)$, the randomized greedy works as follows. First, it chooses a random ordering π over edges. Then starting with an empty matching M , it iteratively adds the i -th edge in the ordering π to M if it is not adjacent to any edge in M . The main result of this section is the following.

► **Theorem 12.** *Let A be the randomized greedy for the maximum matching problem. Then, for any graph $G = (V, E)$ and a vertex $v \in V$, we have $d_{\text{EM}}(A(G), A(G - v)) \leq 1$.*

We need to consider deleting vertices to analyze the sensitivity of our randomized $(1 - \epsilon)$ -approximation algorithm for the maximum matching problem in Section 2. Our analysis is a slight modification of a similar result for the maximal independent set problem [2]. Hence, we defer the proof to the full version of the paper.

2.3.2 Sensitivity of Algorithm 3

We first analyze the sensitivity of AUGMENTINGPATHS. Let us fix the graphs $G = (V, E)$ and $G' = (V, E')$, and matchings $M \subseteq E$ and $M' \subseteq E'$, a positive integer ℓ , and $\delta > 0$. Let H and H' be the layered graphs constructed using G, M, ℓ and G', M', ℓ , respectively, and let \tilde{L} and \tilde{L}' be the set of active vertices in H and H' , respectively.

► **Lemma 13.** *We have $d_{\text{H}}(\tilde{L}, \tilde{L}') \leq 3d_{\text{H}}(E, E') + 3d_{\text{H}}(M, M')$.*

Proof. Each edge modification in the graph or the matching may cause activate/inactivate at most three vertices in the layered graph (two of them are in the first and last layers, and the remaining one is in one of the middle layers), and hence the lemma follows. ◀

For two tag functions $t, t' : V(H) \rightarrow V(H) \cup \{\text{untagged}, \text{deadend}\}$, we define $d_{\text{H}}(t, t') = |\{v \in V(H) \mid t(v) \neq t'(v)\}|$. We will use symbols t and t' to denote tag functions for H and H' , respectively. Note that the supposed domain of t' is $V(H')$, but it is equal to $V(H)$.

► **Lemma 14.** *Let $\mathcal{P} = \text{AUGMENTINGPATHS}(G, M, \ell, \delta)$ and $\mathcal{P}' = \text{AUGMENTINGPATHS}(G', M', \ell, \delta)$. Then, we have*

$$d_{\text{EM}}(\mathcal{P}, \mathcal{P}') \leq (d_{\text{H}}(E, E') + d_{\text{H}}(M, M')) \cdot 3^K,$$

where $K = (1/\epsilon)^{2^{O(1/\epsilon)}}$.

Proof. Let A and A' denote $\text{AUGMENTINGPATHS}(G, M, \delta, \ell)$ and $\text{AUGMENTINGPATHS}(G', M', \delta, \ell)$, respectively. Let M_1, M_2, \dots, M_K and M'_1, M'_2, \dots, M'_K be the sequences of matchings constructed during the process of A and A' , respectively. Note that A and A' construct the same number of matchings, and that $K \leq (1/\epsilon)^{2^{O(1/\epsilon)}}$, which follows by a similar argument to that in the proof of Lemma 6. For $i \in [K]$, let S_i and S'_i be the vertex sets on which M_i and M'_i , respectively, are constructed, that is, the vertex set passed on to RANDOMIZEDGREEDY, and let t_i and t'_i be the tag functions right before constructing M_i and M'_i , respectively. By Lemma 13, we have $d_{\text{EM}}(S_1, S'_1) \leq c$, where $c = 3d_{\text{H}}(M, M') + 3d_{\text{H}}(E, E')$. First, because each difference between M_{i-1} and M'_{i-1} increases the Hamming distance between t_i and t'_i by one, we have

$$\begin{aligned} d_{\text{EM}}(t_i, t'_i) &\leq d_{\text{EM}}(M_{i-1}, M'_{i-1}) + d_{\text{EM}}(t_{i-1}, t'_{i-1}) \\ &\leq \sum_{j=1}^{i-1} d_{\text{EM}}(M_j, M'_j) + d_{\text{EM}}(t_1, t'_1) = \sum_{j=1}^{i-1} d_{\text{EM}}(M_j, M'_j). \end{aligned}$$

Then we have

$$d_{\text{EM}}(S_i, S'_i) \leq d_{\text{EM}}(M_{i-1}, M'_{i-1}) + d_{\text{EM}}(t_i, t'_i) \leq 2 \sum_{j=1}^{i-1} d_{\text{EM}}(M_j, M'_j) \leq 2 \sum_{j=1}^{i-1} d_{\text{EM}}(S_j, S'_j),$$

where the last inequality is due to Theorem 12. Solving this recursion, we get

$$\sum_{j=1}^i d_{\text{EM}}(S_j, S'_j) \leq c \cdot 3^{i-1},$$

and hence we have $d_{\text{EM}}(S_K, S'_K) = 2 \cdot 3^{K-2}c$, and the claim follows. \blacktriangleleft

We now show that the sensitivity of Algorithm 3 is $O_\epsilon(1)$.

► **Theorem 15.** *The sensitivity of Algorithm 3 is at most 3^K , where $K = (1/\epsilon)^{2^{O(1/\epsilon)}}$.*

Proof. Let $G = (V, E)$ be a graph and $G' = (V, E') = G - e$ for some $e \in E$. Let M_0, M_1, \dots, M_{kr} be the sequence of matchings we construct in Algorithm 3 on G , where M_0 is the matching constructed at Line 3, and M_j is the matching constructed at Line 3 in the round i of the phase ℓ such that $j = (\ell - 1)r + i$. We define $M'_0, M'_1, \dots, M'_{kr}$ similarly using G' . Then, we have by Theorem 12

$$d_{\text{EM}}(M_0, M'_0) \leq 1,$$

and we have by Lemma 14

$$\begin{aligned} d_{\text{EM}}(M_i, M'_i) &\leq d_{\text{EM}}(M_{i-1}, M'_{i-1}) + (d_{\text{EM}}(E, E') + d_{\text{EM}}(M_{i-1}, M'_{i-1})) \cdot 3^K \\ &= d_{\text{EM}}(M_{i-1}, M'_{i-1})(3^K + 1) + 3^K \end{aligned}$$

for $i \in [kr]$, where $K = (1/\epsilon)^{2^{O(1/\epsilon)}}$. Solving the recursion, we get

$$d_{\text{EM}}(M_{kr}, M'_{kr}) \leq 2(1 + 3^K)^{kr} - 1,$$

and we have the desired bound. \blacktriangleleft

The proof of Theorem 1 then follows from Theorem 11 and Theorem 15.

Sensitivity to Vertex Deletions

We remark that Algorithm 3 also has sensitivity $O(3^K)$, for $K = (1/\epsilon)^{2^{O(1/\epsilon)}}$, to vertex deletions. Recall that Lemma 13 crucially relies on each edge deletion changing at most three vertices in the layered graph. That is, due to the construction of the layered graph, each edge deletion changes at most two altered vertices in the first and last layers, and at most one altered vertex in one of the middle layers. This is because the first layer and the last layer encode the vertex set V , while each matched edge is assigned to one of the middle layers.

Observe that when we delete a vertex v , at most one vertex in the vertex set V is altered, so that the first and last layer of the layered graph each have one change. Moreover, at most one matched edge is incident to v , so at most one vertex in one of the middle layers is altered as well. Thus, at most three vertices in the layered graph are changed as a result of the vertex deletion, so the sensitivity of Algorithm 3 to vertex deletions is again $O(3^K)$, for $K = (1/\epsilon)^{2^{O(1/\epsilon)}}$.

3 Deterministic Maximal Matching for Bounded-Degree Graphs

In this section, we give a deterministic algorithm for computing a maximal matching that has low sensitivity on bounded-degree graphs. The main idea is to use deterministic local computation algorithms with a small number of probes to find a maximal matching. Our algorithm uses two main ingredients. The first ingredient is a deterministic LCA of [4] for 6^Δ -coloring a graph with maximum degree Δ , using $O(\Delta \log^* n)$ probes. The second ingredient is a framework of [10] that simulates local distributed algorithms using a deterministic LCA. In particular, we use the framework to simulate an algorithm that takes a coloring of a graph and outputs a maximal matching. We give the details for the local distributed algorithm in Algorithm 5.

To bound the sensitivity of the algorithm, it suffices to analyze the number of queries for which a deleted edge would be probed. Crucially, both the deterministic LCA of [4] and the framework of [10] only probe edges (incident to vertices) within a small radius of the query. Thus, only a small number of queries will probe the edge that is altered, so that the output of the algorithm only has a small number of changes.

We first require a deterministic LCA of [4] for 6^Δ -coloring a graph with degree Δ , using a small number of probes within distance $O(\Delta \log^* n)$ of the query. We give the full details in Algorithm 4, which has the following guarantee.

► **Lemma 16** ([4]). *There exists a deterministic LCA `COLORINGLCA` for 6^Δ -coloring a graph with degree Δ , using $O(\Delta \log^* n)$ probes.*

We now describe a local distributed algorithm that takes a coloring of a graph and outputs a maximal matching. The algorithm iterates over all colors and adds any edge adjacent to a vertex of a particular color to the greedy matching if there is no other adjacent edge already present in the matching. We give the algorithm in full in Algorithm 5.

Putting things together, we obtain a deterministic maximal matching algorithm in Algorithm 6.

We next require the following framework of [10] that simulates local distributed algorithms using a deterministic LCA. In particular, we will implement Algorithm 5.

► **Lemma 17** ([7, 10]). *Given access to an oracle that takes vertices of an underlying graph as queries and outputs a color for the queried vertex, there exists a deterministic LCA that can implement `COLORING-TO-MM` using $\Delta^{O(c)}$ probes.*

Given a `COLORINGLCA` for 6^Δ -coloring, the LCA for maximal matching in Lemma 17 uses the following idea. For a query edge e , we first call `COLORINGLCA` for every vertex with distance roughly 6^Δ from e . Parnas and Ron [10] then shows it suffices to run Algorithm 5 locally on the graph of radius roughly 6^Δ from e .

We now show that our deterministic LCA based algorithm outputs a maximal matching with low worst case sensitivity for low-degree graphs.

■ **Algorithm 4** LCA Algorithm COLORINGLCA for 6^Δ -coloring with $O(\Delta \log^* n)$ probes.

```

1 Procedure FORMFORESTS( $G, \Delta$ )
2   //Decompose graph into  $\Delta$  oriented forests;
3   for  $i = 1$  to  $\Delta$  do
4     Let  $N_u(i)$  denote the  $i$ -th neighbor of  $u$  according to the IDs of the vertices;
5     Let  $E_i = \{(u, v) : \text{id}(u) < \text{id}(v), v = N_u(i)\}$ ;
6     Let  $G_i = (V, E_i)$  be the oriented tree, where the root has no out-going edges;
7 Procedure COLORFORESTS( $G_i, \Delta$ )
8   for  $\Theta(\log^* n)$  rounds do
9     for each nodes  $u$  do
10      if  $u$  is a root node then
11        Set  $\phi_u$  to 0;
12      else
13        Let  $v$  be the parent of  $u$  in  $G_i$ ;
14        Let  $a_u$  be the index of the least significant bit with  $\phi_u \neq \phi_v$ ;
15        Let  $b_u$  be the value of the  $a_u$ -th bit of  $u$ ;
16         $\phi_u \leftarrow a_u \circ b_u$ .
```

■ **Algorithm 5** Maximal Matching Algorithm COLORING-TO-MM.

```

1 Procedure COLORING-TO-MM( $G$  colored with  $c$  colors)
2    $M \leftarrow \emptyset$ ;
3   for color  $i = 1$  to  $i = c$  do
4     if edge  $(u, v)$  has either  $\phi_u = i$  or  $\phi_v = i$  then
5       Add  $(u, v)$  to  $M$  if no adjacent edge is in  $M$ .
```

■ **Algorithm 6** Maximal Matching Algorithm.

```

1 Procedure COLORING-TO-MM(Graph  $G$ )
2   Coloring  $G_C = \text{COLORINGLCA}(G, \Delta)$ ;
3   Output COLORING-TO-MM( $G_C$ );
```

Proof of Theorem 2. Consider running the deterministic LCA from Lemma 17 that simulates COLORING-TO-MM on a graph G and a graph $G' := G - e$, for some $e \in E$. Let S be the set of vertices that are assigned different colors in G and G' by COLORINGLCA. First observe that e is within distance $\Theta(\log^* n)$ from at most $\Delta^{\Theta(\log^* n)}$ other vertices. Hence, from Lemma 16 we have $|S| \leq \Delta^{\Theta(\log^* n)}$. Moreover, each vertex $u \in S$ is within distance $O(6^\Delta)$ from at most $\Delta^{O(6^\Delta)}$ other vertices. Thus the total number of edges that differ between the matchings M and M' output by COLORING-TO-MM for G and G' respectively is at most

$$\Delta^{\Theta(\log^* n)} \cdot \Delta^{O(6^\Delta)} = \Delta^{O(6^\Delta + \log^* n)}. \quad \blacktriangleleft$$

It is clear that an almost identical analysis goes through for vertex sensitivity.

4 Lower Bounds for Maximum Matching

In this section, we show lower bounds for deterministic and randomized algorithms for the maximum matching problem.

4.1 Deterministic Lower Bound

In this section, we prove Theorem 3, which claims that *any* deterministic algorithm for the maximum matching problem has edge sensitivity $\Omega(\log^* n)$. Our proof relies on Ramsey's theorem. First we introduce some definitions. Let Y be a finite set. We say that X is a k -subset of Y if $X \subseteq Y$ and $|X| = k$. Let $Y^{(k)} = \{X \subseteq Y \mid |X| = k\}$ be the collection of all k -subsets of Y . A c -labeling of $Y^{(k)}$ is an arbitrary function $f : Y^{(k)} \rightarrow [c]$. Then we say that $X \subseteq Y$ is *monochromatic* in f if $f(A) = f(B)$ for all $A, B \in X^{(k)}$. Let $R_c(n; k)$ be the smallest integer N such that the following holds: for any set Y with at least N elements, and for any c -labeling f of $Y^{(k)}$, there is an n -subset of Y that is monochromatic in f . If no such N exists, $R_c(n; k) = \infty$. Define $\text{twr}(k)$ as the tower of twos of height k , that is, $\text{twr}(1) = 2$ and $\text{twr}(k + 1) = 2^{\text{twr}(k)}$. We will use the following formulation of Ramsey's theorem.

► **Theorem 18** (Special case of Theorem 1 in [6]). *For any positive integer t , $R_{2^t}(t + 1; t) \leq \text{twr}(O(t))$.*

We now show that any deterministic constant-factor approximation algorithm for the maximum matching problem has edge sensitivity $\Omega(\log^* n)$.

Proof of Theorem 3. Let A be an arbitrary deterministic algorithm that outputs a maximal matching, let t be a positive integer, which will be determined later. Let \mathcal{G} be a class of graphs on the vertex set $[n]$ consisting of a cycle v_1, \dots, v_t with $v_1 < v_2 < \dots < v_t$ and $n - t$ isolated vertices. Given a matching M on the cycle v_1, \dots, v_t , we encode it to an integer $0 \leq k \leq 2^t - 1$ so that the i -th bit of k is 1 if and only if the edge $\{v_i, v_{i+1}\}$ belongs to M , where we regard $v_{t+1} = v_1$. Then, we can regard the algorithm A as a function $f : \binom{[n]}{t} \rightarrow \{0, 1, \dots, 2^t - 1\}$, that is, given a set $\{v_1, \dots, v_t\} \subseteq [n]$ with $v_1 < v_2 < \dots < v_t$, we compute a matching on the cycle v_1, \dots, v_t , and encode it to an integer. Then if $n \geq R_{2^t}(t + 1; t)$, which holds when $t = O(\log^* n)$ by Theorem 18, there exists a set $S = \{s_0, s_1, \dots, s_t\} \subseteq [n]$ with $s_0 < s_1 < \dots < s_t$ such that $f(T)$ is constant whenever $T \subseteq S$ with $|T| = t$. Let $G, G' \in \mathcal{G}$ be the graph with cycles s_0, \dots, s_{t-1} and s_1, \dots, s_t , respectively, and let M and M' be the matching output by A on G and G' , respectively. As M and M' have the same encoding, $\{s_i, s_{i+1 \bmod t}\} \in M$ if and only if $\{s_{i+1}, s_{i+2}\} \in M'$, where we regard $s_{t+2} = s_1$. Note that, however, if $\{s_i, s_{i+1 \bmod t}\} \in M$ then $\{s_{i+1 \bmod t}, s_{i+2 \bmod t}\} \notin M'$. It follows that $d_H(M, M') = \Omega(|M|) = \Omega(t)$, where the last equality holds because A has a constant approximation ratio. ◀

4.2 Lower Bounds for Deterministic Greedy Algorithm

As we have seen in Section 2.3.1, the randomized greedy algorithm has $O(1)$ sensitivity even for vertex deletion. Can we derandomize it without increasing the sensitivity? To make the question more precise, let V be a set of n vertices and π be a permutation over $\binom{V}{2}$. Then, let A_π denote the greedy algorithm such that, starting with an empty matching M , it iteratively adds the i -th edge with respect to π to M if and only if the edge does not share an endpoint with any edge in M . We now show that the answer to the question is negative.

► **Theorem 19.** *For any permutation π over $\binom{V}{2}$, the algorithm A_π has sensitivity $\Omega(n)$.*

Proof. We say that an element e of a poset *covers* another element e' if $e > e'$, where $>$ is the order relation of the poset, and there is no other element e'' such that $e > e'' > e'$. Then, we construct a poset P on the element set $\binom{V}{2}$ in which a pair $e \in \binom{V}{2}$ covers another pair $e' \in \binom{V}{2}$ if $\pi(e) > \pi(e')$ and $|e \cap e'| \geq 1$. Note that the size of any antichain in P is at most $n/2$: A set of elements of size more than $n/2$ must have two elements e, e' with $|e \cap e'| \geq 1$, which form a chain of length two. Hence, we need at least $\binom{n}{2}/(n/2) = n - 1$ antichains to cover all the elements in P . Then by Mirsky's theorem, there exists a chain, say, e_1, \dots, e_{n-1} , of size $n - 1$ in P . From the construction of P , e_1, \dots, e_{n-1} forms a path of length $n - 1$. Then, A_π on the path e_1, \dots, e_{n-1} outputs edges with odd indices, whereas A_π on the path e_2, \dots, e_{n-1} outputs edges with even indices, and hence the sensitivity of A_π is $\Omega(n)$. ◀

4.3 Lower Bounds for Randomized Algorithms

The following shows that sensitivity must increase as approximation ratio goes to one.

▶ **Theorem 20.** *Let $\epsilon > 0$. Any (possibly randomized) $(1 - \epsilon)$ -approximation algorithm for the maximum matching problem has sensitivity $\Omega(1/\epsilon)$.*

Proof. For simplicity, we assume $1/10\epsilon$ is an even integer. Let A be an arbitrary $(1 - \epsilon)$ -approximation algorithm for the maximum matching problem, and let G be a graph consisting of a cycle of length $1/10\epsilon$ and $n - 1/10\epsilon$ isolated vertices. Clearly G has two disjoint maximum matchings, say, M_1, M_2 , of size $1/20\epsilon$. Let p_1 and p_2 be the probability that A on G outputs M_1 and M_2 , respectively. Then as A has approximation ratio $1 - \epsilon$, we have

$$p_1 \cdot \frac{1}{20\epsilon} + p_2 \cdot \frac{1}{20\epsilon} + (1 - p_1 - p_2) \cdot \left(\frac{1}{20\epsilon} - 1 \right) \geq \frac{1 - \epsilon}{20\epsilon}.$$

Hence, we have $p_1 + p_2 \geq 19/20$, and it follows that at least one of $p_1 \geq 19/40$ and $p_2 \geq 19/40$ hold. Without loss of generality, we assume $p_1 \geq 19/40$.

Let G' be the graph obtained from G by removing one edge in M_1 . Then, G' has a unique maximum matching M_2 . Let p'_2 be the probability that A on G' outputs M_2 . As A has approximation ratio $1 - \epsilon$, we have

$$p_2 \cdot \frac{1}{20\epsilon} + (1 - p_2) \cdot \left(\frac{1}{20\epsilon} - 1 \right) \geq \frac{1 - \epsilon}{20\epsilon},$$

which implies $p'_2 \geq 19/20$. Hence, the sensitivity of A is at least

$$\max\left(\Pr[A(G) = M_1] - \Pr[A(G') \neq M_2], 0\right) \cdot d_H(M_1, M_2) \geq \left(\frac{19}{40} - \frac{1}{20}\right) \cdot \frac{1}{10\epsilon} = \Omega\left(\frac{1}{\epsilon}\right). \blacktriangleleft$$

5 Weighted Sensitivity and Maximum Weighted Matching

In this section, we consider a generalization of sensitivity to weighted graphs, and show an approximation algorithm with low sensitivity for the maximum weighted matching problem.

5.1 Weighted Sensitivity

Given a weight function over the edges $w : E \rightarrow \mathbb{R}$ of a graph $G = (V, E)$ and two edge sets S and S' , let

$$d_H^w(S, S') = \sum_{e \in S \Delta S'} w(e),$$

58:18 Sensitivity Analysis of the Maximum Matching Problem

where Δ again denotes the symmetric difference. For random edge sets X and X' , we use $d_{\text{EM}}^w(X, X')$ to denote the weighted earth mover's distance between X and X' with respect to w , so that

$$d_{\text{EM}}^w(X, X') = \min_{\mathcal{D}} \mathbf{E}_{(S, S') \sim \mathcal{D}} d_{\text{H}}^w(S, S'),$$

be the weighted Hamming distance between S and S' with respect to w , where \mathcal{D} is a distribution such that its marginal distributions on the first and second coordinates are X and X' , respectively. For a real-valued function β on graphs, we say that the *weighted sensitivity* of an algorithm A that outputs a set of edges is at most β if for every graph $G = (V, E)$, a weight function $w : E \rightarrow \mathbb{R}$, and an edge $e \in E$,

$$d_{\text{EM}}^w(A(G), A(G - e)) \leq \beta(G).$$

A priori, it is not clear whether the weighted sensitivity of an algorithm should correlate with the weight of removed edges. Thus we say that the *normalized weighted sensitivity* is at most β if

$$\frac{d_{\text{EM}}^w(A(G), A(G - e))}{w(e)} \leq \beta(G).$$

5.2 Algorithm Description

We use a simple approach of partitioning the input by weight, finding a maximal matching on each partition, and finally forming a weighted matching by greedily adding edges from the maximal matchings, beginning with the matchings in the largest weight classes. The approach is known to give a $(4 + \epsilon)$ -approximation [1, 5] to the maximum weighted matching. However to bound the weighted sensitivity of our algorithm, we must choose ϵ carefully.

Formal description of our algorithm is given in Algorithm 7. It first defines subsets of edges E_i , where we assume the weight of each edge is polynomially bounded in n , so that $1 \leq w(e) \leq n^c$ for some constant c . For a parameter $\alpha > 1$, we define E_i to be the subset of edges in E with weight at least α^i . Algorithm 7 first draws a random permutation π of the edges and greedily forms a maximal matching M_i on each set E_i induced by π . It then greedily adds edges to a maximal matching, starting from the matching of the heaviest weight class and moving downward. That is, we initialize M to be the empty set and greedily add edges of M_i to M , starting with $i = O(\log_{\alpha} n^c)$ and decrementing i after each iteration.

► **Theorem 21** ([1, 5]). *Algorithm 7 gives an $\frac{1}{4\alpha}$ -approximation to the maximum weighted matching and uses runtime $O(m \log_{\alpha} n)$ on a graph with m edges and n vertices.*

5.3 Sensitivity Analysis

We first require the following key structural lemma that we use to prove Theorem 12 in the full version of the paper.

► **Lemma 22.** *In expectation, the deletion of an edge e alters at most one edge in M_i , i.e., at most one edge in M_i is inserted or deleted in expectation.*

The sensitivity analysis follows from the observation that the deletion of an edge e can only affect the matchings E_i for which $\alpha^i \leq w(e)$. Moreover by Lemma 22, the deletion of edge e affects at most one edge in E_i , in expectation. Thus in expectation, the deletion of e affects at most two edges in E_{i-1} in expectation and inductively, the deletion of e affects at most 2^j

■ **Algorithm 7** Algorithm for maximum weighted matching.

```

1 Procedure MATCHING( $G, \epsilon, w$ )
2   Let  $C$  be a sufficiently large constant and  $\alpha > 1$  be a trade-off parameter.;
3   Let  $\pi$  be a random ordering of the edges of  $G$ ;
4   For each  $i = 0$  to  $i = C \log n$ , let  $E_i$  be the set of edges with weight at least  $\alpha^i$ ;
5   Let  $M_i$  be the greedy maximal matching of  $E_i$  induced by  $\pi$ ;
6    $M \leftarrow \emptyset$ ;
7   for  $i = C \log n$  to 0 do
8     for  $e \in M_i$  do
9       if  $e$  is not adjacent to  $M$  then
10         $M \leftarrow M \cup \{e\}$ ;
11  return  $M$ .

```

edges in E_{i-j} in expectation. On the other hand, the weight of each edge in E_{i-j} is at most α^{i-j} so the weighted sensitivity is $\sum_j \alpha^{i-j} 2^j$. Hence for $\alpha = 2$, the weighted sensitivity is $O(2^i \log n)$ and for $\alpha > 2$, the weighted sensitivity is $O(2^i)$. Similarly for normalized weighted sensitivity, we rescale by $\frac{1}{2^i}$ so that the normalized weighted sensitivity is $O(\log n)$ for $\alpha = 2$ and $O(1)$ for $\alpha > 2$. We now formalize this intuition.

► **Theorem 23.** *Suppose $1 \leq w(e) \leq W \leq n^c$ for some absolute constants $W, c > 0$ for all $e \in E$. The weighted sensitivity of Algorithm 7 is $O(W \log n)$ for $\alpha = 2$ and $O(W)$ for $\alpha > 2$. The normalized weighted sensitivity of Algorithm 7 is $O(\log n)$ for $\alpha = 2$ and $O(1)$ for $\alpha > 2$.*

Proof. Let e be an edge of weight $w(e) \in [2^i, 2^{i+1}]$ for some integer $i \geq 0$ and suppose e is removed from G . For $j \leq i$, let S_j be the set of edges in E_j affected by the deletion of edge e , so that by Lemma 22, $\mathbb{E}[|S_i|] \leq 1$. Then we have $\mathbb{E}[|S_i| \cup \{e\}] \leq 2$ so that Lemma 22 implies that $\mathbb{E}[|S_{i-1}|] \leq 2$. Now suppose that for a fixed $j \leq i$, we have $\mathbb{E}[|\{e\} \cup \bigcup_{k=j}^i S_k|] \leq 2^{i-j}$. Then Lemma 22 implies that $\mathbb{E}[|S_{j-1}|] \leq 2^{i-j}$ so that

$$\mathbb{E} \left[\left| \{e\} \cup \bigcup_{k=j-1}^i S_k \right| \right] \leq 2^{i-j+1}.$$

Hence by induction, we have $\mathbb{E}[|S_j|] \leq 2^{i-j}$.

Since each edge of E_i has weight at most α^i , then we have

$$d_{\text{EM}}(A(G), A(G - e)) = \sum_{j=0}^i 2^{i-j} \alpha^j,$$

where $A(G)$ represents the output of Algorithm 7 on G . Under the assumption that $1 \leq w(e) \leq n^c$ for some absolute constant $c > 0$ for all $e \in E$, then $i = O(\log n)$. Hence for $\alpha = 2$, we have $\sum_{j=0}^i 2^{i-j} \alpha^j = O(2^i \log n) = O(W \log n)$ and for $\alpha > 2$, we have $\sum_{j=0}^i 2^{i-j} \alpha^j = O(2^i) = O(W)$. Moreover, we have

$$\overline{d_{\text{EM}}}(A(G), A(G - e)) \leq \frac{1}{2^{i+1}} \sum_{j=0}^i 2^{i-j} \alpha^j,$$

so that $\overline{d_{\text{EM}}}(A(G), A(G - e)) \leq O(\log n)$ for $\alpha = 2$ and $\overline{d_{\text{EM}}}(A(G), A(G - e)) = O(1)$ for $\alpha > 2$. ◀

Together, Theorem 21 and Theorem 23 give the full guarantees of Algorithm 7.

► **Theorem 24.** *Let $G = (V, E)$ be a weighted graph with $w(e) \leq W \leq n^c$ for some constant $c > 0$ and all $e \in E$. For a trade-off parameter α , there exists an algorithm that outputs a $\frac{1}{4\alpha}$ -approximation to the maximum weighted matching in $O(m \log_\alpha n)$ time. For $\alpha = 2$, the algorithm has weighted sensitivity $O(W \log n)$ and normalized weighted sensitivity $O(\log n)$. For $\alpha > 2$, the algorithm has weighted sensitivity $O(W)$ and normalized weighted sensitivity $O(1)$.*

We again emphasize that the worst case weighted sensitivity of *any* constant factor approximation algorithm to the maximum weighted matching problem is at least $\Omega(W)$. Recall that if an edge of weight $W = n^c$ is altered in a graph whose remaining edges have weight 1, then any constant factor approximation to the maximum weighted matching must include the heavy edge for sufficiently large $c > 0$, which incurs cost $\Omega(W)$ in the weighted sensitivity. Thus for $\alpha > 2$, Algorithm 7 performs well with respect to both weighted sensitivity and normalized weighted sensitivity.

References

- 1 Marc Bury, Elena Grigorescu, Andrew McGregor, Morteza Monemizadeh, Chris Schwiegelshohn, Sofya Vorotnikova, and Samson Zhou. Structural results on matching estimation with applications to streaming. *Algorithmica*, 81(1):367–392, 2019.
- 2 Keren Censor-Hillel, Elad Haramaty, and Zohar Karnin. Optimal dynamic distributed mis. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 217–226, 2016.
- 3 Vincent Cohen-Addad, Niklas Hjuler, Nikos Parotsidis, David Saulpic, and Chris Schwiegelshohn. Fully dynamic consistent facility location. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3250–3260, 2019.
- 4 Richard Cole and Uzi Vishkin. Deterministic coin tossing with applications to optimal parallel list ranking. *Information and Control*, 70(1):32–53, 1986.
- 5 Michael Crouch and Daniel S. Stubbs. Improved streaming algorithms for weighted matching, via unweighted matching. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, volume 28, pages 96–104, 2014.
- 6 P. Erdős and R. Rado. Combinatorial theorems on classifications of subsets of a given set. *Proceedings of the London Mathematical Society*, s3-2(1):417–439, 1952.
- 7 Guy Even, Moti Medina, and Dana Ron. Distributed maximum matching in bounded degree graphs. In *Proceedings of the International Conference on Distributed Computing and Networking (ICDCN)*, pages 18:1–18:10, 2015.
- 8 Silvio Lattanzi and Sergei Vassilvitskii. Consistent k -clustering. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pages 1975–1984, 2017.
- 9 Andrew McGregor. Finding graph matchings in data streams. In *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, pages 170–181, 2005.
- 10 Michal Parnas and Dana Ron. Approximating the minimum vertex cover in sublinear time and a connection to distributed algorithms. *Theoretical Computer Science*, 381(1-3):183–196, 2007.
- 11 Pan Peng and Yuichi Yoshida. Average sensitivity of spectral clustering. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 1132–1140, 2020.
- 12 Nithin Varma and Yuichi Yoshida. Average sensitivity of graph algorithms. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2021. to appear.