

# Improved Deterministic Distributed Matching via Rounding\*

Manuela Fischer

ETH Zurich, Switzerland  
manuela.fischer@inf.ethz.ch

---

## Abstract

We present improved deterministic distributed algorithms for a number of well-studied matching problems, which are simpler, faster, more accurate, and/or more general than their known counterparts. The common denominator of these results is a *deterministic distributed rounding* method for certain *linear programs*, which is the first such rounding method, to our knowledge. A sampling of our end results is as follows.

- An  $O(\log^2 \Delta \cdot \log n)$ -round deterministic distributed algorithm for computing a maximal matching, in  $n$ -node graphs with maximum degree  $\Delta$ . This is the first improvement in about 20 years over the celebrated  $O(\log^4 n)$ -round algorithm of Hańćkowiak, Karoński, and Panconesi [SODA'98, PODC'99].
- A deterministic distributed algorithm for computing a  $(2 + \varepsilon)$ -approximation of maximum matching in  $O(\log^2 \Delta \cdot \log \frac{1}{\varepsilon} + \log^* n)$  rounds. This is exponentially faster than the classic  $O(\Delta + \log^* n)$ -round 2-approximation of Panconesi and Rizzi [DIST'01]. With some modifications, the algorithm can also find an  $\varepsilon$ -maximal matching which leaves only an  $\varepsilon$ -fraction of the edges on unmatched nodes.
- An  $O(\log^2 \Delta \cdot \log \frac{1}{\varepsilon} + \log^* n)$ -round deterministic distributed algorithm for computing a  $(2 + \varepsilon)$ -approximation of a maximum weighted matching, and also for the more general problem of maximum weighted  $b$ -matching. These improve over the  $O(\log^4 n \cdot \log_{1+\varepsilon} W)$ -round  $(6 + \varepsilon)$ -approximation algorithm of Panconesi and Sozio [DIST'10], where  $W$  denotes the maximum normalized weight.
- A deterministic local computation algorithm for a  $(2 + \varepsilon)$ -approximation of maximum matching with  $2^{O(\log^2 \Delta)} \cdot \log^* n$  queries. This improves almost exponentially over the previous deterministic constant approximations which have query-complexity of  $2^{\Omega(\Delta \cdot \log \Delta)} \cdot \log^* n$ .

A full version of this paper with all proofs is available on arXiv.org [9].

**1998 ACM Subject Classification** G.2.2 Graph Algorithms

**Keywords and phrases** distributed graph algorithms, deterministic distributed algorithms, rounding linear programs, maximal matching, maximum matching approximation

**Digital Object Identifier** 10.4230/LIPIcs.DISC.2017.17

## 1 Introduction and Related Work

We work with the standard LOCAL model of distributed computing [24]: the network is abstracted as a graph  $G = (V, E)$ , with  $n = |V|$ ,  $m = |E|$ , and maximum degree  $\Delta$ . Each node has a unique identifier. In each round, each node can send a message to each of its neighbors. We do not limit the message sizes, but for all the algorithms that we present,

---

\* A full version of the paper is available at <http://arxiv.org/abs/1703.00900>.



$O(\log n)$ -bit messages suffice. We assume that all nodes have knowledge of  $\log \Delta$  up to a constant factor. If this is not the case, it is enough to try exponentially increasing estimates for  $\log \Delta$ .

### 1.1 Broader Context, and Deterministic Distributed Rounding

Efficient deterministic distributed graph algorithms remain somewhat of a rarity, despite the intensive study of the area since the 1980's. In fact, among the four classic problems of the area – maximal independent set,  $(\Delta + 1)$ -vertex-coloring, maximal matching, and  $(2\Delta - 1)$ -edge-coloring – only for maximal matching a poly  $\log n$ -round deterministic algorithm is known, due to a breakthrough of Hańćkowiak, Karoński, and Panconesi [14, 16]. Finding poly  $\log n$ -round deterministic algorithms for the other three problems remains a long-standing open question, since [24]. In a stark contrast, in the world of randomized algorithms, all these problems have  $O(\log n)$ -round [28, 1] or even more efficient algorithms [3, 10, 17].

Despite this rather bleak state of the art for deterministic algorithms, there is immense motivation for them. Here are three sample reasons: (1) One traditional motivation is rooted in the classic complexity-theoretic quest which seeks to understand the difference between the power of randomized and distributed algorithms. (2) Another traditional motivation comes from practical settings where even small error probabilities cannot be tolerated. (3) Nowadays, there is also a more modern motive: we now understand that in order to have faster randomized algorithms, we *must* come up with faster deterministic algorithms.<sup>1</sup> This connection goes in two directions: (A) Almost all the recent developments in randomized algorithms use the *shattering technique* [3, 10, 17, 12] which randomly breaks down the graph into small components, typically of size poly  $\log n$ , and then solves them via a deterministic algorithm. Speeding up (the  $n$ -dependency in) these randomized algorithms needs faster deterministic algorithms. (B) The more surprising direction is the reverse. Chang et al. [4] recently showed that for a large class of problems the randomized complexity on  $n$ -node graphs is at least the deterministic complexity on  $\Theta(\sqrt{\log n})$ -node graphs. Hence, if one improves over (the  $n$ -dependency in) the current randomized algorithms, one has inevitably improved the corresponding deterministic algorithm.

Ghaffari, Kuhn, and Maus [11] recently proved a *completeness*-type result which shows that “*the only obstacle*” for efficient deterministic distributed graph algorithms is deterministically *rounding* fractional values to integral values while approximately preserving some linear constraints.<sup>2</sup> To put it more positively, if we find an efficient deterministic method for rounding, we would get efficient algorithms for essentially all the classic local graph problems, including the four mentioned above.

Our results become more instructive when viewed in this context. The common denominator of our results is a deterministic distributed method which allows us to round fractional matchings to integral matchings. This can be more generally seen as rounding the fractional solutions of a special class of *linear programs* (LPs) to integral solutions. To the best of our knowledge, this is the first known *deterministic distributed rounding* method. We can now say that

*matching admits an efficient deterministic algorithm because  
matching admits an efficient deterministic distributed rounding.*

<sup>1</sup> For instance, our improvement in the deterministic complexity of maximal matching directly improves the randomized complexity of maximal matching, as we formally state in Corollary 3.

<sup>2</sup> Stating this result formally and in full generality requires some definitions. See [11] for the precise statement.

## 1.2 Our Results

We provide improved distributed algorithms for a number of matching problems, as we overview next.

### 1.2.1 Approximate Maximum Matching

► **Theorem 1.** *There is an  $O(\log^2 \Delta \cdot \log \frac{1}{\varepsilon} + \log^* n)$ -round deterministic distributed algorithm for a  $(2 + \varepsilon)$ -approximate maximum matching, for any  $\varepsilon > 0$ .*

There are three remarks in order, regarding this result:

- For constant  $\varepsilon > 0$ , this  $O(\log^2 \Delta + \log^* n)$ -round algorithm is significantly faster than the previously best known deterministic constant approximations, especially in low-degree graphs: the  $O(\Delta + \log^* n)$ -round 2-approximation of Panconesi and Rizzi [30], the  $O(\log^4 n)$ -round 2-approximation of Hańćkowiak et al. [16], the  $O(\log^4 n)$ -round  $(3/2)$ -approximation of Czygrinow et al. [6, 7], and its extension [5] which finds a  $(1 + \varepsilon)$ -approximation in  $\log^{O(\frac{1}{\varepsilon})} n$  rounds.
- This  $O(\log^2 \Delta + \log^* n)$ -round complexity gets close to the lower bound – due to the celebrated results of Kuhn et al. [22, 23] and Linial [24] – of  $\Omega(\log \Delta / \log \log \Delta + \log^* n)$  that holds for any constant approximation of matching, even for randomized algorithms.
- This distributed LOCAL algorithm can be transformed to a deterministic *Local Computation Algorithm* (LCA) [2, 33] for a  $(2 + \varepsilon)$ -approximation of maximum matching, with a query complexity of  $2^{O(\log^3 \Delta)} \cdot \log^* n$ . This is essentially by using the standard method of Parnas and Ron [32], with an additional idea of [8]. Using slightly more care, the query complexity can be improved to  $2^{O(\log^2 \Delta)} \cdot \log^* n$ . Since formally stating this result requires explaining the computational model of LCAs, we defer that to the journal version. We remark that this query complexity improves almost exponentially over the previous deterministic constant approximations with  $2^{\Omega(\Delta \cdot \log \Delta)} \cdot \log^* n$  [8].

### 1.2.2 (Almost) Maximal Matching, and Edge Dominating Set

**Maximal Matching.** Employing our approximation algorithm for maximum matching, we get an  $O(\log^2 \Delta \cdot \log n)$ -round deterministic distributed algorithm for maximal matching.

► **Theorem 2.** *There is an  $O(\log^2 \Delta \cdot \log n)$ -round deterministic maximal matching algorithm.*

This is the first improvement in about 20 years over the breakthroughs of Hańćkowiak et al., which presented first an  $O(\log^7 n)$ - [14] and then an  $O(\log^4 n)$ -round [16] algorithm for maximal matching.

As alluded to before, this improvement in the deterministic complexity directly leads to an improvement in the  $n$ -dependency of the randomized algorithms. In particular, plugging in our improved deterministic algorithm in the maximal matching algorithm of Barenboim et al. [3] improves their round complexity from  $O(\log^4 \log n + \log \Delta)$  to  $O(\log^3 \log n + \log \Delta)$ .

► **Corollary 3.** *There is an  $O(\log^3 \log n + \log \Delta)$ -round randomized distributed algorithm that with high probability<sup>3</sup> computes a maximal matching.*

<sup>3</sup> As standard, *with high probability* means with probability at least  $1 - 1/n^c$ , for a desirably large constant  $c \geq 2$ .

**Almost Maximal Matching.** Recently, there has been quite some interest in characterizing the  $\Delta$ -dependency in the complexity of maximal matching, either with no dependency on  $n$  at all or with at most an  $O(\log^* n)$  additive term [18, 13]. Göös et al. [13] conjectured that

*there should be no  $o(\Delta) + O(\log^* n)$  algorithm for computing a maximal matching.*

Theorem 2 does not provide any news in this regard, because of its multiplicative  $\log n$ -factor. Indeed, our findings also seem to be consistent with this conjecture and do not suggest any way for breaking it. However, using some extra work, we can get a faster algorithm for  $\varepsilon$ -maximal matching, a matching that leaves only  $\varepsilon$ -fraction of edges among unmatched nodes, for a desirably small  $\varepsilon > 0$ .

► **Theorem 4.** *There is an  $O(\log^2 \Delta \cdot \log \frac{1}{\varepsilon} + \log^* n)$ -round deterministic distributed algorithm for an  $\varepsilon$ -maximal matching, for any  $\varepsilon > 0$ .*

This theorem statement is interesting because of two aspects: (1) This faster almost maximal matching algorithm sheds some light on the difficulties of proving the aforementioned conjecture. In a sense, any conceivable proof of this conjectured lower bound must distinguish between maximal and almost maximal matchings and rely on the fact that precisely a maximal matching is desired, and not just something close to it. Notice that since the complexity of Theorem 4 grows slowly as a function of  $\varepsilon$ , we can choose  $\varepsilon$  quite small. By setting  $\varepsilon = \Delta^{-\text{poly log } \Delta}$ , we get an algorithm that, in  $O(\text{poly log } \Delta + \log^* n)$  rounds, produces a matching that seems to be maximal for almost all nodes, even if they look up to their  $\text{poly log } \Delta$ -hop neighborhood. (2) Perhaps, in some practical settings, this almost maximal matching, which practically looks maximal for essentially all nodes, may be as useful as maximal matching, especially since it can be computed much faster.

**Approximate Minimum Edge Dominating Set.** As a corollary of the almost maximal matching algorithm of Theorem 4, we get a fast algorithm for approximating *minimum edge dominating set*, which is the smallest set of edges such that any edge shares at least one endpoint with them. The proof can be found in the full version [9].

► **Corollary 5.** *2+eps-approx-EDS There is an  $O(\log^2 \Delta \cdot \log \frac{\Delta}{\varepsilon} + \log^* n)$ -round deterministic distributed algorithm for a  $(2 + \varepsilon)$ -approximate minimum edge dominating set, for any  $\varepsilon > 0$ .*

Previously, the fastest algorithms ran in  $O(\Delta + \log^* n)$  rounds [30] or  $O(\log^4 n)$  rounds [16], providing 2-approximations. Moreover, Suomela [34] provided roughly 4-approximations in  $O(\Delta^2)$  rounds, in a restricted variant of the LOCAL model with only port numberings.

### 1.2.3 Approximate Maximum Weighted Matching and B-Matching

An interesting aspect of the method we use is its flexibility and generality. In particular, the algorithm of Theorem 1 can be easily extended to computing a  $(2 + \varepsilon)$ -approximation of maximum weighted matching, and more interestingly, to a  $(2 + \varepsilon)$ -approximation of maximum weighted *b-matching*. These extensions can be found in the full version [9].

► **Theorem 6.** *There is an  $O(\log^2 \Delta \cdot \log \frac{1}{\varepsilon} + \log^* n)$ -round deterministic distributed algorithm for a  $(2 + \varepsilon)$ -approximate maximum weighted matching, or *b-matching*, for any  $\varepsilon > 0$ .*

To the best of our knowledge, this is the first distributed deterministic algorithm for approximating maximum (weighted) *b-matching*. Moreover, even in the case of standard matching, it improves over the previously best-known algorithm: A deterministic algorithm

for  $(6 + \varepsilon)$ -approximation of maximum weighted matching was provided by Panconesi and Sozio [31], with a round complexity of  $O(\log^4 n \cdot \log_{1+\varepsilon} W)$ , where  $W$  denotes the maximum normalized weight. However, that deterministic algorithm does not extend to  $b$ -matching.

### 1.3 Related Work, Randomized Distributed Matching Approximation

Aside from the deterministic algorithms discussed above, there is a long line of research on randomized distributed approximation algorithms for matching: for the unweighted case, [19] provide a 2-approximation in  $O(\log n)$  rounds, and [26] a  $(1 + \varepsilon)$ -approximation in  $O(\log n)$  for any constant  $\varepsilon > 0$ . For the weighted case, [35, 27, 26] provide successively improved algorithms, culminating in the  $O(\log \frac{1}{\varepsilon} \cdot \log n)$ -round  $(2 + \varepsilon)$ -approximation of [26]. Moreover, [21] present an  $O(\log n)$ -round randomized algorithm for 2-approximate weighted  $b$ -matching.

## 2 Our Deterministic Rounding Method, in a Nutshell

The main ingredient in our results is a simple deterministic method for *rounding* fractional solutions to integral solutions. We believe that this *deterministic distributed rounding* will be of interest well beyond this paper. To present the flavor of our deterministic rounding method, here we overview it in a simple special case: we describe an  $O(\log^2 \Delta)$ -round algorithm for a constant approximation of the maximum unweighted matching in 2-colored bipartite graphs. The precise algorithm and proof appear in Section 4.1.1.

**Fractional Solution.** First, notice that finding a fractional approximate maximum matching is straightforward. In  $O(\log \Delta)$  rounds, we can compute a fractional matching  $\mathbf{x} \in [0, 1]^m$  whose total value  $\sum_e x_e$  is a constant approximation of maximum matching. One standard method is as follows: start with all edge values at  $x_e = 2^{-\lceil \log \Delta \rceil}$ . Then, for  $O(\log \Delta)$  rounds, in each round raise all edge values  $x_e$  by a 2-factor, except for those edges that are incident to a node  $v$  such that  $\sum_{e \in E(v)} x_e \geq 1/2$ . Throughout,  $E(v) := \{e \in E : v \in e\}$  denotes the set of edges incident to node  $v$ . One can easily see that this fractional matching has total value  $\sum_e x_e$  within a 4-factor of the maximum matching.

**Gradual Rounding.** We gradually round this fractional matching  $\mathbf{x} \in [0, 1]^m$  to an integral matching  $\mathbf{x}' \in \{0, 1\}^m$  while ensuring that we do not lose much of the value, i.e.,  $\sum_e x'_e \geq (\sum_e x_e)/C$ , for some constant  $C$ . We have  $O(\log \Delta)$  rounding phases, each of which takes  $O(\log \Delta)$  rounds. In each phase, we get rid of the smallest (non-zero) values and thereby move closer to integrality. The initial fractional matching has<sup>4</sup> only values  $x_e = 2^{-i}$  for  $i \in \{0, \dots, \lceil \log \Delta \rceil\}$  or  $x_e = 0$ . In the  $k^{\text{th}}$  phase, we partially round the edge values  $x_e = 2^{-i}$  for  $i = \lceil \log \Delta \rceil - k + 1$ . Some of these edges will be raised to  $x_e = 2 \cdot 2^{-i}$ , while others are dropped to  $x_e = 0$ . The choices are made in a way that keeps  $\sum_e x_e$  essentially unchanged, as we explain next.

Consider the graph  $H$  edge-induced by edges  $e$  with value  $x_e = 2^{-i}$ . For the sake of simplicity, suppose all nodes of  $H$  have even degree. Dealing with odd degrees requires some delicate care, but it will not incur a loss worse than an  $O(2^{-i})$ -fraction of the total value. In this even-degree graph  $H$ , we effectively want that for each node  $v$  of  $H$ , half of its edges

<sup>4</sup> Any fractional maximum matching can be transformed to this format, with at most a 2-factor loss in the total value: simply round down each value to the next power of 2, and then drop edges with values below  $2^{-(\lceil \log \Delta \rceil + 1)}$ .

raise  $x_e = 2^{-i}$  to  $x_e = 2 \cdot 2^{-i}$  while the others drop it to  $x_e = 0$ . For that, we generate a degree-2 graph  $H'$  by replacing each node  $v$  of  $H$  with  $d_H(v)/2$  nodes, each of which gets two of  $v$ 's edges<sup>5</sup>. Notice that the edge sets of  $H'$  and  $H$  are the same. Graph  $H'$  is simply a set of cycles of even length, as  $H$  was bipartite.

In each cycle of  $H'$ , we would want that the raise and drop of edge weights is alternating. That is, odd-numbered, say, edges are raised to  $x_e = 2 \cdot 2^{-i}$  while even-numbered edges are dropped to  $x_e = 0$ . This would keep  $\mathbf{x}$  a valid fractional matching – meaning that each node  $v$  still has  $\sum_{e \in E(v)} x_e \leq 1$  – because the summation  $\sum_{e \in E(v)} x_e$  does not increase, for each node  $v$ . Furthermore, it would keep the total weight  $\sum_e x_e$  unchanged. If the cycle is shorter than length  $O(\log \Delta)$ , this raise/drop sequence can be identified in  $O(\log \Delta)$  rounds. For longer cycles, we cannot compute such a perfect alternation in  $O(\log \Delta)$  rounds. However, one can do something that does not lose much<sup>6</sup>: imagine that we chop the longer cycles into edge-disjoint paths of length  $\Theta(\log \Delta)$ . In each path, we drop the endpoints to  $x_e = 0$  while using a perfect alternation inside the path. These border settings mean we lose  $\Theta(1/\log \Delta)$ -fraction of the weight. Thus, even over all the  $O(\log \Delta)$  iterations, the total loss is only a small constant fraction of the total weight.

### 3 Preliminaries

**Matching and Fractional Matching.** An integral matching  $M$  is a subset of  $E$  such that  $e \cap e' = \emptyset$  for all  $e \neq e' \in M$ . It can be seen as an assignment of values  $x_e \in \{0, 1\}$  to edges, where  $x_e = 1$  iff  $e \in M$ , such that  $c_v := \sum_{e \in E(v)} x_e \leq 1$  for all  $v \in V$ . When the condition  $x_e \in \{0, 1\}$  is relaxed to  $0 \leq x_e \leq 1$ , such an assignment is called a fractional matching.

**B-Matching.** A  $b$ -matching for  $b$ -values  $\{1 \leq b_v \leq d_G(v) : v \in V\}$  is an assignment of values  $x_e \in \{0, 1\}$  to edges  $e \in E$  such that  $\sum_{e \in E(v)} x_e \leq b_v$  for all  $v \in V$ . Throughout,  $d_G(v)$  denotes the degree of  $v$  in  $G$ . Again, one can relax this to fractional  $b$ -matchings by replacing  $x_e \in \{0, 1\}$  with  $0 \leq x_e \leq 1$ .

**Maximal and  $\varepsilon$ -Maximal Matching.** An integral matching is called maximal if we cannot add any edge to it without violating the constraints. For  $\varepsilon > 0$ , we say that  $M \subseteq E$  is an  $\varepsilon$ -maximal matching if  $|\Gamma^+(M)| \geq (1 - \varepsilon)|E|$  for  $\Gamma^+(M) := \{e \in E \mid \exists e' \in M : e \cap e' \neq \emptyset\}$ , that is, if after removing the edges in and incident to  $M$  from  $G$ , at most  $\varepsilon|E|$  edges remain.


**Maximum and Approximate Maximum Matching.** A matching  $M^*$  is called maximum if it is the/a largest matching in terms of cardinality. For any  $c > 1$ , we say that a matching is  $c$ -approximate if  $c \sum_{e \in E} x_e \geq |M^*|$  for a maximum matching  $M^*$ . In a weighted graph where each edge  $e$  is assigned a weight  $w_e \geq 0$ , we say that  $M^*$  is a maximum weighted matching if it is the/a matching with maximum weight  $w(M^*) := \sum_{e \in M^*} w_e$ . An integral matching  $M$  is a  $c$ -approximate weighted matching if  $c \sum_{e \in M} w_e \geq w(M^*)$ .

We now state some simple and well-known facts about matchings.

► **Lemma 7.** For a maximal matching  $M$  and a maximum matching  $M^*$  in  $G = (V, E)$ , we have the following two properties: (i)  $|M| \geq \frac{|E|}{2\Delta - 1}$ , and (ii)  $\frac{|M^*|}{2} \leq |M| \leq |M^*|$ .

<sup>5</sup> This simple idea has been used frequently before. For instance, it gives an almost trivial proof of Petersen's 2-factorization theorem from 1891 [29]. It has also been used by [20, 14, 16].

<sup>6</sup> Our algorithm actually does something slightly different, but describing this ideal procedure is easier.



■ **Figure 1** A graph and its 2-decomposition.

► **Lemma 8** (Panconesi and Rizzi [30]). *There is an  $O(\Delta + \log^* n)$ -round deterministic distributed algorithm for maximal matching. Furthermore, if a  $q$ -coloring of the graph is provided, then the algorithm runs in  $O(\Delta + \log^* q)$  rounds.*

Many problems are easier in small-degree graphs. To exploit this fact, we sometimes use the following simple transformation which decomposes a graph into graphs with maximum degree 2 – that is, node-disjoint paths and cycles – with the same edge set, in zero rounds. As mentioned before, this has been used frequently in prior work [29, 20, 14, 16].

**2-decomposition.** We 2-decompose graph  $G$  as follows. For every node  $v \in V$ , introduce  $\lceil \frac{d_G(v)}{2} \rceil$  copies and arbitrarily split its incident edges among these copies in such a way that every copy has degree 2, with the possible exception of one copy which has degree 1 (when  $v$  has odd degree). The graph on these copy nodes is what we call a 2-decomposition of  $G$ . See Figure 1 for an example.

## 4 Approximate Maximum Matching

We present a  $(2 + \varepsilon)$ -approximation algorithm for maximum matching, proving Theorem 1. The first step towards this goal is finding a constant approximation, explained in Section 4.1. We show in Section 4.2 how to further improve this approximation ratio to  $2 + \varepsilon$ .

### 4.1 Constant Approximate Maximum Matching

In this subsection, we show how to compute a constant approximation.

► **Lemma 9.** *There is an  $O(\log^2 \Delta + \log^* n)$ -round deterministic distributed algorithm for a  $c$ -approximate maximum matching, for some constant  $c$ .*

The key ingredient for our  $c$ -approximation algorithm of Lemma 9 is a distributed algorithm that computes a constant approximate maximum matching in the special case of a 2-colored bipartite graph. We first present the algorithm for this special case in Section 4.1.1, and then explain in Section 4.1.2 how to reduce the general graph case to the bipartite case, hence proving Lemma 9.

#### 4.1.1 Constant Approximate Maximum Matching in Bipartite Graphs

Next, we show how to find a  $c$ -approximate matching in a 2-colored bipartite graph.

► **Lemma 10.** *There is an  $O(\log^2 \Delta)$ -round deterministic distributed algorithm for a  $c$ -approximate maximum matching in a 2-colored bipartite graph, for some constant  $c$ .*

**Roadmap.** The proof of Lemma 10 is split into three parts. In the first step, explained in Lemma 13, we compute a  $2^{-\lceil \log \Delta \rceil}$ -fractional 4-approximate maximum matching in  $O(\log \Delta)$  rounds. The second step, which is also the main step of our method and is formalized in Lemma 14, is a method to round these fractional values to almost integrality in  $O(\log^2 \Delta)$  rounds. In the third step, presented in Lemma 15, we resort to a simple constant-round algorithm to transform the almost integral matching that we have found up to this step into an integral matching. As a side remark, we note that we explicitly state some of the constants in this part of the paper, for the sake of readability. We remark that these constants are not the focus of this work, and we have not tried to optimize them.

We start with some helpful definitions.

► **Definition 11** (Loose and tight nodes and edges). Given a fractional matching, we call a node  $v$  *loose* if  $c_v = \sum_{e \in E(v)} x_e \leq \frac{1}{2}$ , and *tight* otherwise, where  $E(v) := \{e \in E : v \in e\}$ . We call an edge *loose* if both of its endpoints are loose; otherwise, the edge is called *tight*.

► **Definition 12** (The fractionality of a fractional matching). We call a fractional matching  $2^{-i}$ -fractional for an  $i \in \mathbb{N}$  if  $x_e \in \{0\} \cup \{2^{-j} : 0 \leq j \leq i\}$ . Notice that a  $2^{-0}$ -fractional matching is simply an integral matching.

**Step 1, Fractional Matching.** We show that a simple greedy algorithm already leads to a fractional 4-approximate maximum matching.

► **Lemma 13.** *There is an  $O(\log \Delta)$ -round deterministic distributed algorithm for a  $2^{-\lceil \log \Delta \rceil}$ -fractional 4-approximate maximum matching.*

**Proof.** Initially, set  $x_e = 2^{-\lceil \log \Delta \rceil}$  for all  $e \in E$ . This trivially satisfies the constraints  $c_v = \sum_{e \in E(v)} x_e \leq 1$ . Then, we iteratively raise the value of all loose edges in parallel by a 2-factor. This can be done in  $O(\log \Delta)$  rounds, since at the latest when the value of an edge is  $1/2$ , both endpoints would be tight. Once all edges are tight, for a maximum matching  $M^*$  we have  $\sum_{e \in E} x_e = \frac{1}{2} \sum_{v \in V} c_v \geq \frac{1}{2} \sum_{e=\{u,v\} \in M^*} (c_u + c_v) > \frac{|M^*|}{4}$ . ◀


**Step 2, Main Rounding.** The heart of our approach, the Rounding Lemma, is a method that successively turns a  $2^{-i}$ -fractional matching into a  $2^{-i+1}$ -fractional one, for decreasing values of  $i$ , while only sacrificing the approximation ratio by a little.

► **Lemma 14** (Rounding Lemma). *There is an  $O(\log^2 \Delta)$ -round deterministic distributed algorithm that transforms a  $2^{-\lceil \log \Delta \rceil}$ -fractional 4-approximate maximum matching in a 2-colored bipartite graph into a  $2^{-4}$ -fractional 14-approximate maximum matching.*


**Proof.** Iteratively, for  $k = 1, \dots, \lceil \log \Delta \rceil - 4$ , in phase  $k$ , we get rid of edges  $e$  with value  $x_e = 2^{-i}$  for  $i = \lceil \log \Delta \rceil - k + 1$  by either increasing their values by a 2-factor to  $x_e = 2^{-i+1}$  or setting them to  $x_e = 0$ . In the following, we describe the process for one phase  $k$ , thus a fixed  $i$ .

Let  $H$  be the graph induced by the set  $E_i := \{e \in E : x_e = 2^{-i}\}$  of edges with value  $2^{-i}$  and use  $H'$  to denote its 2-decomposition. Notice that  $H'$  is a node-disjoint union of paths and even-length cycles. Set  $\ell = 12 \log \Delta$ . We call a path/cycle *short* if it has length at most  $\ell$ , and *long* otherwise. We now process short and long cycles and paths, by distinguishing three cases, as we discuss next. Each of these cases will be done in  $O(\log \Delta)$  rounds, which implies that the complexity of one phase is  $O(\log \Delta)$ . Thus, over all the  $O(\log \Delta)$  phases, this rounding algorithm takes  $O(\log^2 \Delta)$  rounds.





■ **Figure 2** The edge values of a short and a long cycle induced by edges in  $E_i$  after rounding. In the long cycle, nodes of color 1 are depicted as blue squares and nodes of color 2 as red circles.



■ **Figure 3** The edge values of two short paths induced by edges in  $E_i$  after rounding. Tight endpoints are depicted as (unfilled) boxes and loose endpoints as (filled) squares.

**Case A, Short Cycles.** Alternately set the values of the edges to 0 and to  $2^{-i+1}$ . Since the cycle has even length, the values  $c_v = \sum_{e \in E(v)} x_e$  for all nodes  $v$  in the cycle remain unaffected by this update. Moreover, the total value of the edges in the cycle stays the same.

**Case B, Long Cycles and Long Paths.** We first orient the edges in a manner that ensures that each maximal directed path has length at least  $\ell$ . This is done in  $O(\ell)$  rounds. For that purpose, we start with an arbitrary orientation of the edges. Then, for each  $j = 1, \dots, \lceil \log \ell \rceil$ , we iteratively merge two (maximal) directed paths of length  $< 2^j$  that are directed towards each other by reversing the shorter one, breaking ties arbitrarily. For more details of this orientation step, we refer to [15, Fact 5.2].

Given this orientation, we determine the new values of  $x_e$  as follows. Recall that we are given a 2-coloring of nodes. Set the value of all border edges (that is, edges that have an incident edge such that they are either oriented towards each other or away from each other) to 0, increase the value of a non-border edge to  $2^{-i+1}$  if it is oriented towards a node of color 1, say, and set it to 0 otherwise.

Now, we show that this process generates a valid fractional matching while incurring only a small loss in the value. Observe that no constraint is violated, as for each node the value of at most one incident edge can be raised to  $2^{-i+1}$  while the other is dropped to 0. Moreover, in each maximal directed path, we can lose at most  $3 \cdot 2^{-i}$  in the total sum of edge values. This happens in the case of an odd-length path starting with a node of color 2. Hence, we lose at most a  $\frac{3}{\ell}$ -fraction of the total sum of the edge values in long cycles and long paths.

**Case C, Short Paths.** Give the path an arbitrary direction, that is, identify the first and the last node. Set the value of the first edge to  $2^{-i+1}$  if the first node is loose, and to 0 otherwise. Alternately, starting with value 0 for the second edge, set the value of every even edge to 0 and of every odd edge to  $2^{-i+1}$ . If the last edge should be set to  $2^{-i+1}$  (that is, the path has odd length) but the last node is tight, set the value of that last edge to 0 instead.

We now discuss the validity of the new fractional matching. If a node  $v$  is in the interior of the path, that is, not one of the endpoints, then  $v$  can have at most one of its incident edges

## 17:10 Improved Deterministic Distributed Matching via Rounding

increased to  $2^{-i+1}$  while the other one decreases to 0. Hence the summation  $c_v = \sum_{e \in E(v)} x_e$  does not increase. If  $v$  is the first or last node in the path, the value of the edge incident to  $v$  is increased only if  $v$  was loose, i.e., if  $c_v = \sum_{e \in E(v)} x_e \leq \frac{1}{2}$ . In this case, we still have  $c_v = \sum_{e \in E(v)} x_e \leq 1$  after the increase, as the value of the edge raises by at most a 2-factor.

We now argue that the value of the matching has not decreased by too much during this update. For that, we group the edges into blocks of two consecutive edges, starting from the first edge. If the path has odd length, the last block consists of a single edge. It is easy to see that the block value, that is, the sum of the values of its two edges, of every interior (neither first nor last) block is unaffected.

If an endpoint  $v$  of a path is loose, the value of the block containing  $v$  remains unchanged or increases (in the case of an odd-length path ending in  $v$ ). If  $v$  is tight, then the value of its block stays the same or decreases by  $2^{-i+1}$ , which is at most a  $2^{-i+2}$ -fraction of the value  $c_v$ .

This allows us to bound the loss in terms of these tight endpoints. The crucial observation is that every node can be endpoint of a short path at most once. This is because, in the 2-decomposition, a node can be the endpoint of a path only if it has a degree-1 copy. This happens only if it has odd degree, and in that case, it has exactly one degree-1 copy, hence, also exactly one endpoint of a short path. Therefore, we lose at most a  $2^{-i+2}$ -fraction in  $\sum_{v \in V} c_v$  when updating the values in short paths.

**Analyzing the Overall Effect of Rounding.** First, we show that over all the rounding phases, the overall loss is only a constant fraction of the total value  $\sum_{e \in E} x_e$ .

Let  $x_e^{(i)}$  and  $c_v^{(i)}$  denote the value of edge  $e$  and node  $v$ , respectively, before eliminating all the edges with value  $2^{-i}$ . Putting together the loss analyses discussed above, we get

$$\sum_{e \in E} x_e^{(i-1)} \geq \sum_{e \in E} x_e^{(i)} - \frac{3}{\ell} \sum_{e \in E} x_e^{(i)} - 2^{-i+2} \sum_{v \in V} c_v^{(i)} \geq \left(1 - \frac{3}{\ell} - 2^{-i+3}\right) \sum_{e \in E} x_e^{(i)}.$$

It follows that

$$\begin{aligned} \sum_{e \in E} x_e^{(4)} &\geq \left( \prod_{i=5}^{\lceil \log \Delta \rceil} \left(1 - \frac{3}{\ell} - 2^{-i+3}\right) \right) \sum_{e \in E} x_e^{(\lceil \log \Delta \rceil)} \geq \left( \prod_{i=5}^{\lceil \log \Delta \rceil} e^{-2(\frac{3}{\ell} + 2^{-i+3})} \right) \sum_{e \in E} x_e^{(\lceil \log \Delta \rceil)} \\ &\geq e^{-\frac{1}{4} - 16 \sum_{i=5}^{\lceil \log \Delta \rceil} 2^{-i}} \sum_{e \in E} x_e^{(\lceil \log \Delta \rceil)} \geq \frac{1}{e^{\frac{5}{4}}} \sum_{e \in E} x_e^{(\lceil \log \Delta \rceil)} \geq \frac{1}{4e^{\frac{5}{4}}} |M^*| \geq \frac{1}{14} |M^*| \end{aligned}$$

for a maximum matching  $M^*$ , recalling that we started with a 4-approximate maximum matching. Here, the second inequality holds because  $\frac{3}{\ell} + 2^{-i+3} \leq \frac{1}{2}$ , as  $i \geq 5$ .

Finally, observe that in all the rounding phases the constraints  $c_v = \sum_{e \in E(v)} x_e \leq 1$  are preserved, since the value  $c_v$  can increase by at most a 2-factor and only when  $v$  is loose.  $\blacktriangleleft$

**Step 3, Final Rounding.** So far, we have an almost integral matching. Next, we round all edges to either 0 or 1, by finding a maximal matching in the subgraph induced by edges with positive value.

**► Lemma 15.** *There is an  $O(1)$ -round deterministic distributed algorithm that, given a  $2^{-4}$ -fractional 14-approximate maximum matching in a 2-colored bipartite graph, computes an integral 434-approximate maximum matching.*

**Proof.** In the given  $2^{-4}$ -fractional matching,  $x_e \neq 0$  means  $x_e \geq \frac{1}{16}$ . Thus, a node cannot have more than 16 incident edges with non-zero value in this fractional matching. In this constant-degree subgraph, a maximal matching  $M$  can be found in  $O(1)$  rounds using the algorithm in Lemma 8, recalling that we are given a 2-coloring. We have  $|M| \geq \frac{|\{e \in E : x_e > 0\}|}{31} \geq \frac{1}{31} \sum_{e \in E} x_e$  by Lemma 7 (i), and, since we started with a 14-approximation,  $M$  is 434-approximate.  $\blacktriangleleft$

### 4.1.2 Constant Approximate Maximum Matching in General Graphs

We now explain how the approximation algorithm for maximum matchings in 2-colored bipartite graphs can be employed to find approximate maximum matchings in general graphs. The main idea is to transform the given general graph into a bipartite graph with the same edge set in such a way that a matching in this bipartite graph can be easily turned into a matching in the general graph.

**Proof of Lemma 9.** Let  $\vec{E}$  be an arbitrary orientation of the edges  $E$ . Split every node  $v \in V$  into two siblings  $v_{\text{in}}$  and  $v_{\text{out}}$ , and add an edge  $\{v_{\text{out}}, v_{\text{in}}\}$  to  $E_B$  for every oriented edge  $(u, v) \in \vec{E}$ . Let  $V_{\text{in}} := \{v_{\text{in}} : v \in V\}$  and  $V_{\text{out}} := \{v_{\text{out}} : v \in V\}$  be the nodes with color 1 and 2, respectively. By Lemma 10, a  $c$ -approximate maximum matching  $M_B$  in the bipartite graph  $B = (V_{\text{in}} \cup V_{\text{out}}, E_B)$  can be computed in  $O(\log^2 \Delta)$  rounds. We now go back to  $V$ , that is, merge  $v_{\text{in}}$  and  $v_{\text{out}}$  back into  $v$ . This makes the edges of  $M_B$  incident to  $v_{\text{in}}$  or  $v_{\text{out}}$  now be incident to  $v$ , leaving us with a graph  $G' = (V, M_B) \subseteq G$  with maximum degree 2.

We compute a maximal matching  $M'$  in  $G'$ . Using the algorithm of Lemma 8, this can be done in  $O(\log^* n)$  rounds. If a poly  $\Delta$ -coloring of  $G$  is provided, which implies a coloring of  $G'$  with poly  $\Delta$  colors, the round complexity of this step is merely  $O(\log^* \Delta)$ .

It follows from Lemma 7 (i) that  $|M'| \geq \frac{|M_B|}{3} \geq \frac{|M_B^*|}{3c} \geq \frac{|M^*|}{3c}$  for maximum matchings  $M_B^*$  in  $B$  and  $M^*$  in  $G$ , respectively. Thus,  $M'$  is a  $3c$ -approximate maximum matching in  $G$ . The last inequality is true since by introducing additional nodes but leaving the edge set unchanged (when going from  $G$  to  $B$ ), the maximum matching size cannot decrease. ◀

## 4.2 Wrap-Up: $(2 + \varepsilon)$ -Approximate Matching and Maximal Matching

In this section, we iteratively invoke the constant approximation algorithm of 9 to obtain algorithms for a  $(2 + \varepsilon)$ -approximate maximum matching (Theorem 1) and a maximal matching (Theorem 2).

The approximation ratio of a matching algorithm can be improved from  $c$  to  $2 + \varepsilon$  easily, by  $O(\log \frac{1}{\varepsilon})$  repetitions: each time, we apply the algorithm of Lemma 9 to the remaining graph, and remove the found matching together with its neighboring edges from the graph.

Before explaining the details, we present the following frequently used trick.

► **Remark.** If a poly  $\Delta$ -coloring of a graph is provided, we can go around the  $\Omega(\log^* n)$  lower bound of Linial [24], omitting the additive  $O(\log^* n)$  term from the round complexity of the algorithms presented in this paper. More generally, if such an algorithm is invoked iteratively, one can first precompute an  $O(\Delta^2)$ -coloring in  $O(\log^* n)$  rounds using Linial's algorithm [25], which allows us to replace the  $O(\log^* n)$  term by  $O(\log^* \Delta)$  by Lemma 8 in each iteration.

**Proof of Theorem 1.** Starting with  $G_0 = G$ , for  $i = 0, \dots, k - 1$ , where  $k = O(\log \frac{1}{\varepsilon})$ , iteratively compute a  $c$ -approximate maximum matching  $M_i$  in  $G_i$ , using the algorithm of Lemma 9. We delete  $M_i$  together with its incident edges from the graph, that is, set  $G_{i+1} = (V, E(G_i) \setminus \Gamma^+(M_i))$ .

Now, we argue that the obtained matching  $\bigcup_{i=0}^{k-1} M_i$  is  $(2 + \varepsilon)$ -approximate. To this end, we bound the size of a maximum matching in the remainder graph  $G_k$ .

Let  $M_i^*$  be a maximum matching in  $G_i$ . An inductive argument shows that  $|M_i^*| \leq (1 - \frac{1}{c})^i |M^*|$ . Indeed, observe  $|M_{i+1}^*| \leq |M_i^*| - |M_i| \leq (1 - \frac{1}{c})|M_i^*| \leq (1 - \frac{1}{c})^{i+1} |M^*|$ , where the first inequality holds since otherwise  $M_{i+1}^* \cup M_i$  would be a better matching than  $M_i^*$  in  $G_i$ , contradicting the latter's optimality. For  $k = \log_{1 - \frac{1}{c}} \frac{\varepsilon}{2(2 + \varepsilon)}$ , we thus have  $|M_k^*| \leq \frac{\varepsilon}{2(2 + \varepsilon)} |M^*|$ . As  $\bigcup_{i=0}^{k-1} M_i$  is a maximal matching in  $G \setminus G_k$  by construction,  $(\bigcup_{i=0}^{k-1} M_i) \cup M_k^*$

## 17:12 Improved Deterministic Distributed Matching via Rounding

is a maximal matching in  $G$ . By Lemma 7 (ii), this means that  $|\bigcup_{i=0}^{k-1} M_i| + |M_k^*| \geq \frac{|M^*|}{2}$ , hence  $|\bigcup_{i=0}^{k-1} M_i| \geq \left(\frac{1}{2} - \frac{\varepsilon}{2(2+\varepsilon)}\right) |M^*| \geq \frac{|M^*|}{2+\varepsilon}$ .

We have  $O(\log \frac{1}{\varepsilon})$  iterations, each taking  $O(\log^2 \Delta + \log^* n)$  rounds. As mentioned in Remark 4.2, by precomputing an  $O(\Delta^2)$ -coloring in  $O(\log^* n)$  rounds, the round complexity of each iteration can be decreased to  $O(\log^2 \Delta + \log^* \Delta) = O(\log^2 \Delta)$ , leading to an overall running time of  $O(\log^2 \Delta \cdot \log \frac{1}{\varepsilon} + \log^* n)$  rounds. ◀

► **Remark.** The analysis above shows that the matching  $M$  computed by the algorithm of Theorem 1 is not only  $(2 + \varepsilon)$ -approximate, but also has the property that any matching in the remainder graph (induced by  $E \setminus \Gamma^+(M)$ ) can have size at most  $\varepsilon |M^*|$  for a maximum matching  $M^*$  in  $G$ .

If one increases the number of repetitions to  $O(\log n)$ , the found matching is maximal.

**Proof of Theorem 2.** Apply the  $c$ -approximation algorithm of Lemma 9 for  $k = \log_{1-\frac{1}{c}} \frac{1}{n}$  iterations on the respective remainder graph, as described in the proof of Theorem 1. The same analysis (also adopting the notation from there) shows that a maximum matching  $M_k^*$  in the remainder graph  $G_k$  must have size  $|M_k^*| \leq \frac{|M^*|}{n} < 1$ , which means that  $G_k$  is an empty graph. But then  $\bigcup_{i=1}^{k-1} M_i$  must be maximal. ◀

## 5 Almost Maximal Matching

In the previous section, we have seen how one can obtain a matching that reduces the size of the matching in the remainder graph, that is, the graph after removing the matching and all incident edges, by a constant factor. Intuitively, one would expect that this also reduces the number of remaining edges by a constant factor, which would directly lead to an (almost) maximal matching just by repetitions. However, this is not the case, since not every matched edge removes the same number of edges from the graph, particularly in non-regular graphs. This calls for an approach that weights edges incident to nodes of different degrees differently, which naturally brings into play weighted matchings.

In Lemma 16, we present a fast algorithm that finds a constant approximation of maximum weighted matching based on the algorithm of Theorem 1. Then, we use this algorithm, by assigning certain weights to the edges, to find a matching that removes a constant fraction of the edges in Lemma 17. Via  $O(\log \frac{1}{\varepsilon})$  repetitions of this, each time removing the found matching and its incident edges, we get an  $\varepsilon$ -maximal matching. More details are provided in the proof of Theorem 4 in the end of this section. Observe that when setting  $\varepsilon = \frac{1}{n^2}$ , thus increasing the number of repetitions to  $O(\log n)$ , we obtain a maximal matching.

► **Lemma 16.** *There is an  $O(\log^2 \Delta + \log^* n)$ -round deterministic distributed algorithm for a 256-approximate maximum weighted matching.*

**Proof.** We assume without loss of generality that the edge weights are normalized, that is, from a set  $\{1, \dots, W\}$  for some maximum weight  $W$ . Round the weights  $w_e$  for  $e \in E$  down to the next power of 8, resulting in weights  $w'_e$ . This rounding procedure lets us lose at most a 8-factor in the total weight and provides us with a decomposition of  $G$  into graphs  $C_i = (V, E_i)$  with  $E_i := \{e \in E : w'_e = 8^i\}$  for  $i \in \{0, \dots, \lfloor \log_8 W \rfloor\}$ .

In parallel, run the algorithm of Theorem 1 with  $\varepsilon = 1$  on every  $C_i$  to find a 3-approximate maximum matching  $M_i$  in  $C_i$  in  $O(\log^2 \Delta + \log^* n)$  rounds. Observe that while the edges in  $\bigcup_i M_i$  do not form a matching, since edges from  $M_i$  and  $M_j$  for  $i \neq j$  can be neighboring, a

matching  $M \subseteq \bigcup_i M_i$  can be obtained by deleting all but the highest-index edge in every such conflict, that is, by removing all edges  $e \in M_i$  with an incident edge  $e' \in M_j$  for a  $j > i$ .

In the following, we argue that the weight of  $M$  cannot be too small compared to the weight of  $\bigcup_i M_i$  by an argument based on counting in two ways.

Every edge  $e \in (\bigcup_i M_i) \setminus M$  puts blame  $w'_e$  on an edge in  $M$  as follows. Since  $e \notin M$ , there is an edge  $e'$  incident to  $e$  such that  $e \in M_i$  and  $e' \in M_j$  for some  $j > i$ . If  $e' \in M$ , then  $e$  blames weight  $w_e$  on  $e'$ . If  $e' \notin M$ , then  $e$  puts blame  $w_e$  on the same edge as  $e'$  does.

For an edge  $e \in M \cap E_i$  and  $j \in [i]$ , let  $n_j$  be the maximum number of edges from  $M_{i-j}$  that blame  $e$ . An inductive argument shows that  $n_j \leq 2^j$ . Indeed, there can be at most two edges from  $M_{i-1}$  blaming  $e$ , at most one per endpoint of  $e$ , and, for  $j > 1$ , we have  $n_j \leq 2 + \sum_{j'=1}^{j-1} n_{j'} \leq 2 + \sum_{j'=1}^{j-1} 2^{j'} = 2^j$ , since at most two edges in  $M_{i-j}$  can be incident to  $e$  and at most one further edge can be incident to each edge in  $M_{i-j'}$  for  $j' < j$ .

Therefore, overall, at most  $\sum_{j=1}^i 2^j 8^{i-j} \leq \frac{1}{3} 8^i \leq \frac{1}{3} w'_e$  weight is blamed on  $e$ . This means that  $\sum_{e \in (\bigcup_i M_i) \setminus M} w'_e \leq \frac{1}{3} \sum_{e \in M} w'_e$ , hence  $\sum_{e \in \bigcup_i M_i} w'_e \leq \frac{4}{3} \sum_{e \in M} w'_e$ , and lets us conclude that  $\sum_{e \in M^*} w_e \leq 8 \sum_{e \in M^*} w'_e \leq 24 \sum_{e \in \bigcup_i M_i} w'_e \leq 32 \sum_{e \in M} w'_e \leq 256 \sum_{e \in M} w_e$  for a maximum weighted matching  $M^*$ . ◀

Next, we explain how to use this algorithm to remove a constant fraction of edges, by introducing appropriately chosen weights. We define the weight of each edge to be the number of its incident edges. This way, an (approximate) maximum weighted matching corresponds to a matching that removes a large number of edges.

► **Lemma 17.** *There is an  $O(\log^2 \Delta + \log^* n)$ -round deterministic distributed algorithm for a  $\frac{511}{512}$ -maximal matching.*

**Proof.** For each edge  $e = \{u, v\} \in E$ , introduce a weight  $w_e = d_G(u) + d_G(v) - 1$ , and apply the algorithm of Lemma 16 to find a 256-approximate maximum weighted matching  $M$  in  $G$ .

For the weight  $w(M^*)$  of a maximum weighted matching  $M^*$ , it holds that  $w(M^*) \geq |E|$ , as the following simple argument based on counting in two ways shows. Let every edge in  $E$  put a blame on an edge in  $M^*$  that is responsible for its removal from the graph as follows. An edge  $e \in M^*$  blames itself. An edge  $e \notin M^*$  blames an arbitrary incident edge  $e' \in M^*$ . Notice that at least one such edge must exist, as otherwise  $M^*$  would not even be maximal. In this way,  $|E|$  many blames have been put onto edges in  $M^*$  such that no edge  $e = \{u, v\} \in M^*$  is blamed more than  $w_e$  times, as  $e$  can be blamed by itself and any incident edge. Therefore, indeed  $w(M^*) = \sum_{e \in M^*} w_e \geq |E|$ , and, as  $M$  is a 256-approximate, it follows that  $\sum_{e \in M} w_e \geq \frac{|E|}{256}$ .

Now, observe that  $w_e$  is the number of edges that are deleted when removing  $e$  together with its incident edges from  $G$ . Since every edge can be incident to at most two matched edges (and thus can be deleted by at most two edges in the matching), in total  $|\Gamma^+(M)| \geq \frac{1}{2} \sum_{e \in M} w_e \geq \frac{|E|}{512}$  many edges are removed from  $G$  when deleting the edges in and incident to  $M$ , which proves that  $M$  is a  $\frac{511}{512}$ -maximal matching. ◀

We iteratively invoke this algorithm to successively reduce the number of remaining edges.

**Proof of Theorem 4.** For  $i = 0, \dots, k = O(\log \frac{1}{\varepsilon})$  and  $G_0 = G$ , iteratively apply the algorithm of Lemma 17 to  $G_i$  to get a  $c$ -maximal matching  $M_i$  in  $G_i$ . Set  $G_{i+1} = (V, E(G_i) \setminus \Gamma^+(M_i))$ , that is, remove the matching and its neighboring edges from the graph. Then  $M := \bigcup_{i=0}^{k-1} M_i$  for  $k = \log_c \varepsilon$  is  $\varepsilon$ -approximate, since  $|E \setminus \Gamma^+(M)| = |E(G_k)| \leq c^k |E| \leq \varepsilon |E|$ , using  $|E(G_{i+1})| \leq c |E(G_i)|$ .

Overall, recalling Remark 4.2, this takes  $O(\log^2 \Delta \cdot \log \frac{1}{\varepsilon} + \log^* n)$ . ◀

**Acknowledgment.** I want to thank Mohsen Ghaffari for suggesting this topic, for his guidance and his support, as well as for the many valuable and enlightening discussions. I am also thankful to Seth Pettie for several helpful comments.

---

### References

---

- 1 Noga Alon, László Babai, and Alon Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *Journal of algorithms*, 7(4):567–583, 1986.
- 2 Noga Alon, Ronitt Rubinfeld, Shai Vardi, and Ning Xie. Space-efficient local computation algorithms. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1132–1139, 2012.
- 3 Leonid Barenboim, Michael Elkin, Seth Pettie, and Johannes Schneider. The locality of distributed symmetry breaking. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 321–330, 2012.
- 4 Yi-Jun Chang, Tsvi Kopelowitz, and Seth Pettie. An exponential separation between randomized and deterministic complexity in the LOCAL model. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 615–624, 2016.
- 5 Andrzej Czygrinow and Michał Hańćkowiak. Distributed algorithm for better approximation of the maximum matching. In *International Computing and Combinatorics Conference*, pages 242–251, 2003.
- 6 Andrzej Czygrinow, Michał Hańćkowiak, and Edyta Szymańska. Distributed algorithm for approximating the maximum matching. *Discrete Applied Mathematics*, 143(1):62–71, 2004.
- 7 Andrzej Czygrinow, Michał Hańćkowiak, and Edyta Szymańska. A fast distributed algorithm for approximating the maximum matching. In *Proceedings of the Annual European Symposium on Algorithms (ESA)*, volume 3221, pages 252–263, 2004.
- 8 Guy Even, Moti Medina, and Dana Ron. Deterministic stateless centralized local algorithms for bounded degree graphs. In *Proceedings of the Annual European Symposium on Algorithms (ESA)*, pages 394–405, 2014.
- 9 Manuela Fischer. Improved deterministic distributed matching via rounding. *CoRR*, abs/1703.00900, 2017. URL: <http://arxiv.org/abs/1703.00900>.
- 10 Mohsen Ghaffari. An improved distributed algorithm for maximal independent set. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 270–277, 2016.
- 11 Mohsen Ghaffari, Fabian Kuhn, and Yannic Maus. On the complexity of local distributed graph problems. In *Proceedings of the Symposium on Theory of Computing (STOC)*, pages 784–797, 2017.
- 12 Mohsen Ghaffari and Hsin-Hao Su. Distributed degree splitting, edge coloring, and orientations. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2505–2523, 2017.
- 13 Mika Göös, Juho Hirvonen, and Jukka Suomela. Linear-in-delta lower bounds in the local model. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pages 86–95, 2014.
- 14 Michał Hańćkowiak, Michał Karonski, and Alessandro Panconesi. On the distributed complexity of computing maximal matchings. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 219–225, 1998.
- 15 Michał Hańćkowiak, Michał Karonski, and Alessandro Panconesi. On the distributed complexity of computing maximal matchings. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 219–225, 1998.
- 16 Michał Hańćkowiak, Michał Karoński, and Alessandro Panconesi. A faster distributed algorithm for computing maximal matchings deterministically. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pages 219–228, 1999.

- 17 David G. Harris, Johannes Schneider, and Hsin-Hao Su. Distributed  $(\Delta+1)$ -coloring in sublogarithmic rounds. In *Proceedings of the Symposium on Theory of Computing (STOC)*, pages 465–478, 2016.
- 18 Juho Hirvonen and Jukka Suomela. Distributed maximal matching: Greedy is optimal. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pages 165–174, 2012.
- 19 Amos Israeli and Alon Itai. A fast and simple randomized parallel algorithm for maximal matching. *Information Processing Letters*, 22(2):77–80, 1986.
- 20 Amos Israeli and Yossi Shiloach. An improved parallel algorithm for maximal matching. *Information Processing Letters*, 22(2):57–60, 1986.
- 21 Christos Koufogiannakis and Neal E. Young. Distributed fractional packing and maximum weighted b-matching via tail-recursive duality. In *Proceedings of the International Symposium on Distributed Computing (DISC)*, pages 221–238, 2009.
- 22 Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. The price of being near-sighted. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 980–989, 2006.
- 23 Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. Local computation: Lower and upper bounds. *J. ACM*, 63(2):17:1–17:44, 2016.
- 24 Nathan Linial. Distributive graph algorithms - global solutions from local data. In *Proceedings of the Symposium on Foundations of Computer Science (FOCS)*, pages 331–335, 1987.
- 25 Nathan Linial. Locality in distributed graph algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992.
- 26 Zvi Lotker, Boaz Patt-Shamir, and Seth Pettie. Improved distributed approximate matching. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pages 129–136, 2008.
- 27 Zvi Lotker, Boaz Patt-Shamir, and Adi Rosen. Distributed approximate matching. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pages 167–174, 2007.
- 28 Michael Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM journal on computing*, 15(4):1036–1053, 1986.
- 29 Henry Martyn Mulder. Julius petersen’s theory of regular graphs. *Discrete mathematics*, 100(1-3):157–175, 1992.
- 30 Alessandro Panconesi and Romeo Rizzi. Some simple distributed algorithms for sparse networks. *Distributed computing*, 14(2):97–100, 2001.
- 31 Alessandro Panconesi and Mauro Sozio. Fast primal-dual distributed algorithms for scheduling and matching problems. *Distributed Computing*, 22(4):269–283, 2010.
- 32 Michal Parnas and Dana Ron. Approximating the minimum vertex cover in sublinear time and a connection to distributed algorithms. *Theoretical Computer Science*, 381(1):183–196, 2007.
- 33 Ronitt Rubinfeld, Gil Tamir, Shai Vardi, and Ning Xie. Fast local computation algorithms. In *Proceedings of the Symposium on Innovations in Computer Science (ICS)*, pages 223–238, 2011.
- 34 Jukka Suomela. Distributed algorithms for edge dominating sets. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pages 365–374, 2010.
- 35 Mirjam Wattenhofer and Roger Wattenhofer. Distributed weighted matching. In *Proceedings of the International Symposium on Distributed Computing (DISC)*, pages 335–348, 2004.