Towards the Locality of Vizing's Theorem

Hsin-Hao Su Boston College

suhx@bc.edu

Hoa T. Vu Boston College vuhd@bc.edu

Abstract

Vizing [Viz64] showed that it suffices to color the edges of a simple graph using $\Delta + 1$ colors, where Δ is the maximum degree of the graph. However, up to this date, no efficient distributed edge-coloring algorithm is known for obtaining such coloring, even for constant degree graphs. The current algorithms that get closest to this number of colors are the randomized $(\Delta + \tilde{\Theta}(\sqrt{\Delta}))$ edge-coloring algorithm that runs in polylog(n) rounds by [CHL⁺18] and the deterministic (Δ + polylog(n))-edge-coloring algorithm that runs in poly(Δ , log n) rounds by [GKMU18].

We present two distributed edge-coloring algorithms that run in $poly(\Delta, \log n)$ rounds. The first algorithm, with randomization, uses only $\Delta + 2$ colors. The second algorithm is a deterministic algorithm that uses $\Delta + O(\log n / \log \log n)$ colors. Our approach is to reduce the distributed edge-coloring problem into an online and restricted version of balls-into-bins problem. If ℓ is the maximum load of the bins, our algorithm uses $\Delta + 2\ell - 1$ colors. We show how to achieve $\ell = 1$ with randomization and $\ell = O(\log n / \log \log n)$ without randomization.

1 Introduction

Problem description: Given a simple graph G = (V, E), a k-edge-coloring is a mapping $\phi : E \to \{1, \ldots, k\}$ that maps each edge to a color where no two adjacent edges are mapped to the same color.

We study the edge-coloring problem in the distributed LOCAL model. In this model, vertices host processors and operate in synchronized rounds. In each round, each vertex sends a message of arbitrary size to each of its neighbors, receives messages from its neighbors, and performs local computations. The time complexity of an algorithm is defined to be the number of rounds used. In the end, each vertex produces its own answer. In the edge-coloring problem, the coloring of an edge uv can either be computed by vertex u or vertex v.

The LOCAL model aims to investigate the *locality* of a problem. An r-round local algorithm in the LOCAL model implies that each vertex only uses information in its r-neighborhood to compute the answer, and vice versa. Therefore, a faster algorithm in the LOCAL model would mean that each vertex uses less local information to compute the answer.

Computing edge-coloring in the distributed setting has applications in scheduling problems of wireless networks [Ram99, GDP08]. It is usually the case that the quality of the solution in these applications depends on the number of colors being used. Therefore, we hope to minimize the number of colors while keeping the coloring to be locally computable.

Let Δ denote the maximum degree of G. It takes at least Δ colors to color the edges of G, since the incident edges to a vertex must be colored with distinct colors. Vizing showed that the edges in every simple graph G can be colored with $\Delta + 1$ colors [Viz64]. The best sequential algorithm for obtaining such a coloring runs in $O(\min(\Delta m \log n, m\sqrt{n \log n}))$ time [Arj82, GNK⁺85]. However, in order to obtain a $(\Delta + 1)$ -edge-coloring in the LOCAL model, the only known method up to this date is a O(diameter(G)) solution – a leader collects the topology of the whole graph and then computes the answer.

"How close can we get to this [Vizing's edge-coloring], while remaining in polylogarithmictime?"

This is an open problem raised in [FGK17]. The followings are the thresholds on the number of colors that have been encountered or tackled by existing algorithms.

 $2\Delta - 1$ Threshold: $(2\Delta - 1)$ is a natural threshold to be investigated because $(2\Delta - 1)$ is the minimum number of colors that can be obtained by the greedy algorithm. Panconesi and Rizzi [PR01] gave a deterministic $(2\Delta - 1)$ -edge-coloring algorithm that runs in $O(\Delta + \log^* n)$ rounds.

The $(2\Delta - 1)$ -edge coloring problem translates to the $(\Delta + 1)$ -vertex-coloring problem on its line graph, L(G), where $\hat{\Delta} = 2\Delta - 2$ is an upper bound on the maximum degree of L(G). There have been extensive studies on the $(\Delta + 1)$ -vertex-coloring problem [Lub86, ABI86, Joh99, GP87, GPS88, Lin92, PS92, KW06, BEK14, Bar15, BEPS16, FHK16, HSS18, CLP18, BEG18]. All the results can be applied to obtain a $(2\Delta - 1)$ -edge-coloring.

Getting a polylog(n)-rounds deterministic algorithm for $(2\Delta - 1)$ -edge-coloring had been a major open problem. After the progress made in [GS17, GHK⁺17] for $(2+o(1))\Delta$ -edge-coloring, the problem was settled recently by [FGK17] and later improved in [GHK18].

 $\Delta + \Theta(\sqrt{\Delta})$ Threshold: Based on a randomized approach, Panconesi and Srinivasan [PS97] first gave an algorithm that uses $1.6\Delta + O(\log^{1+\delta} n)$ colors and runs in $O(\log n)$ rounds w.h.p.¹, where $\delta > 0$ is a constant. Later, Dubhashi, Grable, and Panconesi [DGP98] gave an algorithm that uses $(1 + \epsilon)\Delta$ colors and runs in $O(\log n)$ rounds w.h.p. for constant ϵ and $\Delta > \log^{1+\delta} n$. Later, Elkin, Pettie, and Su [EPS15] removed the requirement that $\Delta \ge \log^{1+\delta} n$ and showed that for constant ϵ , there exists Δ_{ϵ} such that for all $\Delta \ge \Delta_{\epsilon}$, a $(1 + \epsilon)\Delta$ -edge-coloring can be obtained w.h.p. by solving $O(\log^* \Delta)$

¹W.h.p. denotes with high probability, which means with probability at least $1 - 1/n^c$ for some constant c.

Lovasz Local Lemma (LLL) instances in the distributed setting. An LLL instance can be computed in polylog(n) rounds in the distributed setting [CPS14]. This line of research culminated in the work of Chang et al. [CHL⁺18], who showed that for any $\epsilon = \tilde{\Omega}(1/\sqrt{\Delta})$ (not necessarily a constant), it is possible to obtain a $(1 + \epsilon)\Delta$ -edge-coloring w.h.p. by $O(\log(1/\epsilon))$ invocations of distributed LLL with an additive poly(log log n) rounds.

The algorithm of $[CHL^+18]$ allows one to find a $\Delta + \tilde{O}(\sqrt{\Delta})$ -edge-coloring in polylog(n) rounds. However, this number of colors seems to be the limit of such randomized approach. An intuitive reason is because in these coloring algorithms, each edge is colored with a (almost) random color out of the $(1 + \epsilon)\Delta$ colors. Once an edge is assigned a color, it is permanently colored. Consider a stage of the algorithm where an edge uv is still uncolored but all the incident edges are colored. Let M(x) denote the colors that are not assigned to any incident edges of x. The size of M(u) and M(v) are about $\epsilon\Delta$. Note that $M(u) \cap M(v)$ must be non-empty in order to color uv. The $\epsilon\Delta$ colors in M(u) and M(v)are likely to be "randomly" sampled out of the $(1 + \epsilon)\Delta$ colors. Therefore, the expected number of colors in $M(u) \cap M(v)$ is $O(\epsilon^2\Delta)$. This implies ϵ has to be at least $\Omega(1/\sqrt{\Delta})$.

Very recently, in the dynamic setting, Duan et al. [DHZ19] gave a randomized $(1+\epsilon)\Delta$ -edge-coloring algorithm with $O(\operatorname{poly}(1/\epsilon, \log n))$ amortized update time when $\epsilon = \Omega(\sqrt{\Delta}/\log n)$. Incidentally, the number of colors they are able to obtain is also capped at the $\Delta + \Theta(\sqrt{\Delta})$ threshold.

 $\Delta + \text{poly}\log(n)$ Threshold: Recently, Ghaffari et al. [GKMU18] gave a $\text{poly}((1/\epsilon), \log n)$ -round deterministic algorithm that uses $(1 + \epsilon)\Delta$ rounds provided that $\frac{\epsilon}{\log(1/\epsilon)} = \Omega(\log n/\Delta)$. A corollary of the result is that it is possible to obtain $\Delta + O(\log n \cdot \log\left(2 + \frac{\Delta}{\log n}\right))$ colors in $\text{poly}(\Delta, \log n)$ rounds. This breaks the $\Delta + \tilde{\Theta}(\sqrt{\Delta})$ barrier, provided that $\Delta = \omega(\log^c n)$ for a large enough c. However, when Δ is small, say in a constant degree graph, it is still unclear what the minimum possible number of colors is to color the graph in polylog(n) rounds.

1.1 Our Results

We show that by using $poly(\Delta, \log n)$ rounds, the number of colors can be pushed down to $\Delta + 2$, which is merely just one more color than that in Vizing's Theorem.

Theorem 1.1. There exists a randomized distributed $(\Delta + 2)$ -edge-coloring algorithm that runs in $\operatorname{poly}(\Delta, \log n)$ rounds w.h.p. Furthermore, for any $\epsilon \geq (2/\Delta)$, this can be turned into a randomized distributed $(1 + \epsilon)\Delta$ -edge-coloring algorithm that runs in $\operatorname{poly}((1/\epsilon), \log n)$ rounds w.h.p.

Chang et al. [CHL⁺18] showed that any algorithm for $(\Delta + c)$ -edge-coloring based on "extending partial colorings by recoloring subgraphs" needs $\Omega(\frac{\Delta}{c} \log \frac{cn}{\Delta})$ rounds. Our algorithms belongs to such category and so a polynomial dependency on Δ and $\log n$ is necessary.

Since the edge-coloring problem is locally-checkable, by using the derandomization result from [GHK18], we can convert the randomized algorithm to a deterministic algorithm with a $2^{O(\sqrt{\log n})}$ factor blow-up.

Corollary 1.2. For any $\epsilon \geq (2/\Delta)$, there exists a deterministic distributed $(1 + \epsilon)\Delta$ -edge-coloring algorithm that runs in $poly((1/\epsilon)) \cdot 2^{O(\sqrt{\log n})}$ time.

For deterministic algorithms, we also show that it is possible to obtain a $\Delta + O(\log n / \log \log n)$ coloring in $O(\Delta^6 \cdot \log^{O(1)} n)$ rounds.

Theorem 1.3. For any $\lambda > 1$, there exists a deterministic distributed $\Delta + O(\log n/(\log \lambda + \log \log n))$ edge-coloring algorithm that runs in $O(\lambda \cdot \Delta^6 \cdot \log^{O(1)} n)$ rounds. Furthermore, this can be turned into a deterministic distributed $(1+\epsilon)\Delta$ -edge-coloring algorithm that runs in $O(\lambda \cdot (1/\epsilon)^6 \cdot \log^{O(1)} n)$ rounds, provided $\epsilon \geq K \log n/(\Delta(\log \lambda + \log \log n)))$ for some constant K > 0.

Table 1: Minimum number of colors that can be obtained in $poly(\Delta, \log n)$ rounds. The deterministic result on the fourth row is obtained by setting $\lambda = \Delta^{\gamma}$ for some constant $\gamma > 0$ in Theorem 1.3.

	Pre	vious		New	Range of Δ
	$\Delta + \tilde{O}(\sqrt{\Delta})$ [CHI	$[1^{+}18]$			$O(\log^c n)$
Randomized	$\Delta + O(\log n \cdot \log(2 + \log(2 + \log(n + (n + $	$\frac{O(\log n \cdot \log(2 + \frac{\Delta}{\log n}))}{[\text{GKMU18}]}$		Thm. 1.1	$\Omega(\log^c n)$
	1.5Δ [GKM	[GKMU18]			$O(\frac{\log n}{\log \log n})$
Deterministic			$\Delta + O(\frac{\log n}{\log \log n})$	Thm. 1.3	$\left[\Omega(\frac{\log n}{\log\log n}), O(\log n)\right]$
	$\Delta + O(\log n \cdot \log(2 + \frac{1}{16}))$	$\frac{\Delta}{\log n}))$	$\Delta + O(\log_{\Delta} n)$	Thm. 1.3	$[\Omega(\log n), 2^{O(\sqrt{\log n})}]$
	[GKM	IŬ18]	$\Delta + 2$	Cor. 1.2	$2^{\Omega(\sqrt{\log n})}$

Comparing Theorem 1.3 (with $\lambda = O(1)$) to the deterministic algorithms in [GKMU18], we are using fewer colors $(\Delta + O(\log n / \log \log n) \text{ v.s. } \Delta + O(\log n \cdot \log(2 + \frac{\Delta}{\log n})))$ and fewer number of rounds $(\tilde{O}(\Delta^6) \text{ v.s. } \tilde{O}(\Delta^9))$. For bipartite graphs, we show that it is possible to further improve the number of rounds to just $\tilde{O}(\Delta^4)$. We summarize what coloring can be obtained in poly $(\Delta, \log n)$ in Table 1.

The second statements of Theorem 1.1 and Theorem 1.3 can be obtained by applying the degree splitting algorithm of [GHK⁺17] to recursively split the edges into subgraphs whose maximum degrees are upper bounded by $O(1/\epsilon)$ (see Lemma 3.4). Then we can apply the algorithm on each subgraph to obtain a poly($(1/\epsilon)$, log n) algorithm.

1.2 A High-Level Description of Our Method

To get a $poly(\Delta, \log n)$ -round algorithm, we employ the construction of Vizing's coloring in the distributed setting. In Vizing's Theorem, it was shown that it is possible to align the endpoints of an uncolored edge with the same missing color. Thus, the edge can be colored with the missing color. The alignment can be done by shifting a fan and augmenting along an alternating path. The edges can be colored by this method one after another in a sequential manner.

However, in the distributed setting, there are two challenges. First, we have to color a large fraction of uncolored edges simultaneously. The difficulties arise when we color the edges together since there are dependencies among the coloring processes. Second, the number of rounds needed to augment along an alternating path must be at least proportional to its length. It is unclear if short alternating paths always exist (potentially the length of an alternating path can be $\Omega(n)$).

While the first challenge can be resolved by decomposing the coloring process into different phases and applying symmetry breaking techniques carefully, overcoming the second challenge relies on the fact all alternating paths are short. Instead of showing this is true, we chop up long alternating paths to stop the augmentation from propagating. This can be done by placing a blocking edge on it and augment on the portion of the path only from the starting vertex to the blocking edge. In the end, if each vertex is adjacent to at most ℓ blocking edges, then the remaining blocking edges can be colored with additional $2\ell - 1$ colors, because the maximum degree graph induced by the blocking edges is at most ℓ . Therefore, the problem reduces to the following load-balancing problem.

Balls into Bins v.s. Paths into Vertices: Consider the following version of balls-into-bins problem. Suppose there are n bins. The balls arrive in an online fashion. When a ball x arrives, it is given a size-T subset of the bins, B_x . We have to decide which bin in B_x to place x into. Each bin can only be the potential choice of at most t balls. The question is how to minimize the maximum load of the bins.

The problem of minimizing the maximum number of incident blocking edges per vertex can be thought as a balls-into-bins problem. In particular, each alternating path of length at least T is a ball and each vertex is a bin. Choosing the blocking edge on the first T edges is equivalent to choosing the bin (thus each ball occupies two bins, but we ignore this fact for now). We will show that each vertex can be passed by at most t alternating paths throughout the algorithm.

If $t \leq T/\lambda$, then by placing each ball into a random bin, we can show that w.h.p. the maximum load is $O(\log n/(\log \lambda + \log \log n))$. Moreover, such randomized algorithm can be derandomized by letting each ball go to the bin with the smallest cost (we will specify the cost function later), breaking ties arbitrarily. This is how we obtain our deterministic algorithm that uses additional $O(\log n/(\log \lambda + \log \log n))$ colors in Theorem 1.3. Note that the $\Theta(\log n/(\log \lambda + \log \log n))$ bound is tight for the balls-into-bins problem. Additional constraints must be considered in order to further reduce the load.

To achieve maximum load of 1, we modify our algorithm such that we place the ball into one of the non-empty bins randomly. If there is always at least one non-empty bin in B_x for every ball x, it guarantees a maximum load of one. Recall that each ball is a path in our scenario. There are at most $n\Delta^{T-1}$ different paths of length T. If T is large enough $(T = \text{poly}(\Delta, \log n) \cdot t)$, we can show that w.h.p. for every path, at most a constant fraction of the bins on the path are occupied (throughout the algorithm) by taking a union bound over all possible paths of length T. Since there are at least a constant fraction of empty bins, each ball is always guaranteed to land in an empty bin. In the end of the algorithm, the blocking edges induce a matching. After recoloring the matching with a new color, we obtain a $(\Delta + 2)$ -edge-coloring.

2 Notation and Vizing's Theorem

Consider a partial edge coloring $\phi : E \to \{\bot, 1, 2, \dots, \Delta + 1, \star\}$ of a graph G. A partial coloring is proper if the subgraph induced by each color $i \in \{1, 2, \dots, \Delta + 1\}$ is a matching. We say an edge e is uncolored if $\phi(e) = \bot$. The color \star is a special color that denotes the color of blocking edges. Let $M(v) \subseteq \{1, 2, \dots, \Delta + 1\}$ denote the missing colors of v (i.e. the colors that are not assigned to the incident edges of v). Note that M(v) is always non-empty. Fix m(v) to be any missing color of v.

Fix a partial edge-coloring ϕ . Suppose that v is missing α , i.e. $\alpha \in M(v)$, and there is at least one uncolored edge that is incident to v. The fan F_v with respect to ϕ is constructed as follow. Let $e_1 = vx_1$ be an uncolored edge incident to v. Our goal is to color the edge vx_1 . If $m(x_1) \notin M(v)$, then there must be an edge vx_2 with the color $m(x_1)$. In particular, we construct vx_i as the edge with the color $m(x_{i-1})$. We stop this process if for some i, we have that $m(x_i) \in M(v)$ or some previous edge vx_i currently has color $m(x_i)$.

More formally, given $e_i = vx_i$, we construct edge e_{i+1} if both of the following two conditions do not occur: (i) $m(x_i) \in M(v)$. (ii) $m(x_i) \in \{m(x_1), \ldots, m(x_{i-1})\}$. Furthermore, e_{i+1} is constructed as the edge that is incident to v and colored in $m(x_i)$. We repeatedly construct edges $e_1 \ldots e_k$ for some $k \ge 1$ until it is not possible to do it further. The vertex set of F_v consists of the *center* v and the *leaves* x_1, \ldots, x_k . The edge set of F_v is $\{e_1, \ldots, e_k\}$. We say k is the degree of F_v , $\deg(F_v)$. A fan F_v is an $\alpha\beta$ -fan if $\alpha \in M(v)$ and $m(x_{\deg(F_v)}) = \beta$. We again note that this procedure is well-defined since with $\Delta + 1$ colors in the palette, $m(x_i)$ is always non-empty.

We emphasize that when we use F_v to refer to a fan, that is a shorthand to say that it has v as the center. The notation F_v alone however does not uniquely define a fan since we may grow multiple fans centering at v throughout the algorithm.

The operation $\operatorname{shift}(F_v, i)$ recolors the edges of F_v as follow. For each $1 \leq j \leq i-1$, color edge e_j with $m(x_j)$. We leave the color of e_i to be the same. An $\alpha\beta$ -alternating path is a path whose edges are colored in α and β alternatively and it is maximal of such paths. That is, no path whose edges are alternating between α and β contains it properly. Given an $\alpha\beta$ -alternating path P, the operation $\operatorname{augment}(P)$ switches the colors from α to β and from β to α for each edge on P. Note that two $\alpha\beta$ -alternating paths cannot possibly intersect with each other.



Figure 1

2.1 Vizing's Theorem and Fan Repair

Vizing's Theorem shows that the uncolored edge in a fan can be properly colored if one recolors some other edges. This is done by using an **augment** operation and a **shift** operation. We call such a procedure *repairing* the fan. We adopt a version of the proof from [MG92], where it only uses one alternating path instead of potentially two. Given a fan F_v with edges vx_1, \ldots, vx_k . If $m(x_k) \in M(v)$, perform $shift(F_v, k)$ and color vx_k with $m(x_k)$ then every edge in F_v is now colored. Otherwise, it must be the case that there exists j, where $1 \leq j \leq k - 2$, such that $m(x_k) = m(x_j)$. Let P_v be the $\alpha\beta$ -alternating path that starts at v. Note that the first edge of P_v is $e_{j+1} = vx_{j+1}$. There are two cases.

Case 1: P_v does not end at x_j . We perform $\operatorname{shift}(F_v, j+1)$ followed by $\operatorname{augment}(P_v)$. After the $\operatorname{shift}(F_v, j+1)$ operation, there is only one conflict – both vx_j and vx_{j+1} are colored with β . This is resolved by $\operatorname{augment}(P_v)$, which changes the color of vx_{j+1} to α . Moreover, $\operatorname{augment}(P_v)$ does not create new conflicts, since its ending vertex must be missing either α or β (depending on the color of its last edge) before the $\operatorname{shift}(F_v, j+1)$ operation. The operation $\operatorname{shift}(F_v, j+1)$ does not color any edge with α or β except for xv_j . Since P_v does not end at x_j as assumed, the partial coloring remains proper after augmenting along P_v .

Case 2: P_v ends at x_j (Figure 1a), we perform $\operatorname{augment}(P_v)$ first. After the augmentation, now v is missing β and x_j is missing α (Figure 1b). We let $m(x_j) = \alpha$ (there can be multiple missing colors, but we choose $m(x_j)$ to be α). Under the new partial coloring and $m(\cdot)$ -values, F_v is still a fan of v by definition. Moreover, $m(x_k) = \beta \in M(v)$. We then perform $\operatorname{shift}(F_v, k)$ and color vx_k with β (Figure 1c).

Therefore, given a fan, we can repair it and thus decrease the number of uncolored edges. Sequentially, if we fix fans one by one, we will be able to color all the edges. To be able to do this efficiently in the distributed setting, we need to address two challenges.

- We need to be able to color a large portion of uncolored edges at the same time.
- The number of rounds needed to augment along a path is proportional to the length of the path. It is unclear if all alternating paths are short (i.e. $poly(\Delta, \log n)$).

3 Algorithm

To tackle the first challenge, we divide the fans into different color classes so that the fans in the same color class can be fixed without interfering with each other. The rough idea is to take advantage of the fact that any two $\alpha\beta$ -alternating paths do not intersect and the fact that if u and v are in the same color class of a 2-hop coloring then F_u and F_v do not intersect. Therefore, we can repair all $\alpha\beta$ -fans

whose centers are in the same color class of a 2-hop coloring together.

To overcome the second challenge, the idea is to truncate long alternating paths. Given an alternating path $P = (v_0, v_1, \dots, v_k)$, let P(i) denote the truncated version of P defined as follow

$$P(i) = \begin{cases} P, & \text{if } k \leq i. \\ (v_0, v_1, \dots, v_i), & \text{otherwise.} \end{cases}$$

We will truncate alternating paths at length $T = \text{poly}(\Delta, \log n)$, where the actual value of T will be specified later. Instead of doing the **augment** operation on P, we will choose one edge in P(T) to be the blocking edge to block the augmentation propagating down the path. Note that we only need to put a blocking edge on P(T) if |P| > T, since if $|P| \leq T$, we can augment P in O(T) rounds.

3.1 The Framework

We describe our algorithm in Algorithm 1. Let G^2 denote the distance-2 graph of G where $(u, v) \in G^2$ if $\operatorname{dist}_G(u, v) \leq 2$. Since the maximum degree of G^2 is at most $\Delta + \Delta \cdot (\Delta - 1) = \Delta^2$, we can color G^2 using $\Delta^2 + 1$ colors. This can be done in $O(\Delta^2 + \log^* n)$ rounds [BEK14].

Alg	gorithm 1 Distributed Fan Repair
1:	Obtain a 2-hop coloring using $\Delta^2 + 1$ colors.
2:	while there exists uncolored edges do
3:	for $i = 1, 2, \dots, \Delta^2 + 1$ do
4:	for $\alpha = \operatorname{color} 1, 2, \dots, \Delta + 1$ do
5:	Let $V_{i,\alpha}$ denote color- <i>i</i> vertices that misses α and have at least one uncolored edges.
6:	Let $\mathcal{F}_{i,\alpha}$ denote the set of fans grown from each vertex $V_{i,\alpha}$.
7:	for $\beta = \operatorname{color} 1, 2, \dots, \Delta + 1$ do
8:	Let $\mathcal{F}_{i,\alpha,\beta}$ denote the set of $\alpha\beta$ -fans in $\mathcal{F}_{i,\alpha}$.
9:	Build a conflict graph $G_{\mathcal{F}_{i,\alpha,\beta}}$ for $\mathcal{F}_{i,\alpha,\beta}$ and color $G_{\mathcal{F}_{i,\alpha,\beta}}$ using $O(1)$ colors.
10:	Let $\mathcal{F}_{i,\alpha,\beta,j} \subseteq \mathcal{F}_{i,\alpha,\beta}$ denote the fans that are colored in j .
11:	for $j = \operatorname{color} 1, 2, \dots, O(1)$ do
12:	Repair each fan in $\mathcal{F}_{i,\alpha,\beta,j}$ that has not been destroyed.

The outermost while loop repeats until there is no uncolored edge. We will show in Lemma 3.1 that each iteration colors a constant fraction of the uncolored edges. Thus, the while loop uses $O(\log n)$ iterations. Inside the while loop, we iterate through each color of G^2 (Line 3). For each color class, we further iterate through all possible values for α (Line 4). Let $V_{i,\alpha}(\phi)$ denote the set of color-*i* vertices that have at least one incident uncolored edge and miss α w.r.t. the partial coloring ϕ . We often omit the parameter ϕ and write it as $V_{i,\alpha}$ when the reference is w.r.t. the current coloring.

For each $v \in V_{i,\alpha}$, we grow a fan F_v centered at v. Let $\mathcal{F}_{i,\alpha}$ be a collection of all such fans. It is guaranteed that the fans do not intersect, since $V_{i,\alpha}$ is an independent set in G^2 . Next, we iterate β from color 1 to $\Delta + 1$. Let $\mathcal{F}_{i,\alpha,\beta}$ denote all the $\alpha\beta$ -fans in $\mathcal{F}_{i,\alpha}$. Ideally, we want to be able to repair all fans in $\mathcal{F}_{i,\alpha,\beta}$ simultaneously, since their alternating paths do not intersect. However, it is possible that an alternating path of a fan F_v ends at a node of another fan $F_w \in \mathcal{F}_{i,\alpha,\beta}$. In this case, we might not be able to repair them together, because the $\operatorname{augment}(P_v)$ step may destroy the structure of F_w (i.e. change M(x) for $x \in F_w$ or change the color of the edges in F_w). Note that if P_v crosses a node in F_w and does not end at a node in F_w , augmenting along P_v does not affect the structure of F_w , since in this case, P_v cannot contain any edge of F_w . In other words, if P_v crosses but does not end at a node $x \in F_w$, then $m(x) \neq \alpha$ and $m(x) \neq \beta$. Therefore, augmenting P(v) or the prefix of P_v before the blocking edge (if $|P_v| > T$) cannot change m(x).

To resolve this issue, we consider the directed graph $G_{\mathcal{F}_{i,\alpha,\beta}} = (V_{\mathcal{F}_{i,\alpha,\beta}}, E_{\mathcal{F}_{i,\alpha,\beta}})$, where $V_{\mathcal{F}_{i,\alpha,\beta}} = \{F_v \mid F_v \in \mathcal{F}_{i,\alpha,\beta}\}$ and $(F_u, F_w) \in E_{\mathcal{F}_{i,\alpha,\beta}}$ if $|P_v| \leq T$ and P_v ends at any nodes of F_w . Note that we

only consider $|P_v| \leq T$, because if $|P_v| > T$, we will only do modification on $P_v(T)$. Such modification will not affect the fans that intersect with $P_v(T)$.

Since the out-degree of every vertex in $G_{\mathcal{F}_{i,\alpha,\beta}}$ is at most 1, the arborcity of $G_{\mathcal{F}_{i,\alpha,\beta}}$ is O(1). If we consider $G_{\mathcal{F}_{i,\alpha,\beta}}$ as an undirected graph, we can color $G_{\mathcal{F}_{i,\alpha,\beta}}$ using O(1) colors in $O(\log n)$ rounds [BE10], if $G_{\mathcal{F}_{i,\alpha,\beta}}$ were the underlying communication graph. Since each edge in $G_{\mathcal{F}_{i,\alpha,\beta}}$ is stretched by a factor of at most T in G, we can simulate the coloring algorithm in $O(T \cdot \log n)$ rounds in G. Now we iterate through each color j of the coloring and process the fans colored in j together (Line 11). Since the alternating path of each fan colored in j does not end at any other fan colored in j, nor do the alternating paths intersect, we can repair them simultaneously.

Therefore, it is possible to repair all undestroyed fans in $\mathcal{F}_{i,\alpha,\beta,j}$ simultaneously, using a **shift** operation and possibly an **augment** operation for each fan. For a fan F_v with its alternating path P_v whose first edge is the edge between v and the (j + 1)'th leave of F_v , if $|P_v| \leq T$, we repair the fan as outlined in Section 2. Otherwise, we select an edge of $P_v(T) = (v_0, v_1, \ldots, v_T)$ to be the blocking edge, say (v_{i-1}, v_i) , where $1 \leq i \leq T$. We color (v_{i-1}, v_i) with a special color \star and then perform shift(v, j + 1) followed by $\texttt{augment}(P_v(i-1))$ (Case 1 of the repairing step). After the repairing step, the uncolored edge in the fan becomes colored and the partial coloring remains proper. We will show that after $O(\log n)$ iterations of the outermost while loop, all the edges become colored (some possibly in color \star).

Lemma 3.1. In each iteration of the while loop, at least a constant fraction of uncolored edges become colored. Thus, it takes $O(\log n)$ iterations to color all the edges.

Proof. Consider an iteration of the while loop. Let $\phi_0^{-1}(\perp)$ denote the set of uncolored edges at the beginning of the while loop. Let $\phi_{i,\alpha}$ be the partial coloring during the *i*'th iteration of the outer **for loop** at Line 3 and at the beginning of α 'th iteration of the inner **for loop** at Line 4. Let ϕ_i denote the partial coloring at the beginning of *i*'th iteration of the outer **for loop** at Line 3.

For each $uv \in \phi_0^{-1}(\bot)$, we orient the edge from u to v if the color of u is smaller than that of v; otherwise we orient it from v to u. We will show that

$$\sum_{\alpha=1}^{\Delta+1} |\mathcal{F}_{i,\alpha}| \ge \sum_{u \in V_i} \text{outdeg}(u)/2,$$

where V_i denotes the set of color-*i* vertices.

For each $v \in V_i$, let C(v) denote the first outdeg(v) missing colors of v w.r.t. ϕ_i so that |C(v)| = outdeg(v). For each $\alpha \in C(v)$, if $v \in V_{i,\alpha}(\phi_{i,\alpha})$, then a fan centered at v will be created during iteration α (Recall that $V_{i,\alpha}(\phi_{i,\alpha})$ denotes the set of color *i*-vertices that have at least one uncolored incident edge and are missing α with respect to the coloring $\phi_{i,\alpha}$). Otherwise, if $v \notin V_{i,\alpha}(\phi_{i,\alpha})$, it must be caused by the repairing step of some color-*i* vertex *u* that changes the color of an incident edge of v to α during iteration α' for some $\alpha' < \alpha$. In other words, there is a fan F_u that was grown in iteration $\alpha' < \alpha$ and repairing F_u causes $v \notin V_{i,\alpha}(\phi_{i,\alpha})$. Repairing a fan F_u can only change the missing color set (i.e. $M(\cdot)$) of its leaves, its center, and the other endpoint of the alternating path. Since u and v are both in V_i , they cannot be neighbors. Therefore, v must be either the endpoint of the alternating path P_u or v = u.

Case 1: v is the end of the alternating path of F_u . It must be the **augment** operation of F_u that changes the missing color of v during the repairing step of F_u .

Case 2: u = v. Let $x_1 \ldots x_k$ be the leaves of F_u . If an edge incident to u becomes colored in α during iteration α' , it must be the case that $\alpha = m(x_k)$ and $\alpha \in M(u)$. We perform $\text{shift}(F_u, k)$ operation and color ux_k in α . Note that in this case the augment operation is not performed.

Therefore, if $v \notin V_{i,\alpha}(\phi_{i,\alpha})$, we can blame it to some fan F_u . Note that each fan F_u can be blamed at most once, since Case 1 and Case 2 are mutually exclusive. Since for each v and $\alpha \in C(v)$, either it creates a fan or it blames another fan, we have

$$\sum_{u \in V_i} \text{outdeg}(u) = \sum_{u \in V_i} |C(u)| \le \sum_{\alpha=1}^{\Delta+1} |\mathcal{F}_{i,\alpha}| + (\#\text{blames})$$
$$\le 2\sum_{\alpha=1}^{\Delta+1} |\mathcal{F}_{i,\alpha}| .$$

The last inequality follows since $(\#\text{blames}) \leq \sum_{\alpha=1}^{\Delta+1} |\mathcal{F}_{i,\alpha}|$. Note that at most half of the fans are destroyed, since repairing a fan can only destroy at most one other fan. If a fan is not destroyed, then the uncolored edge in the fan must become colored. Therefore, at least $\sum_{i=1}^{\Delta^2+1} \sum_{\alpha=1}^{\Delta+1} |\mathcal{F}_{i,\alpha}|/2$ uncolored edges are colored. The total number of edges that are colored during the while loop is at least

$$\sum_{i=1}^{\Delta^2+1} \sum_{\alpha=1}^{\Delta+1} |\mathcal{F}_{i,\alpha}|/2 \ge \sum_{i=1}^{\Delta^2+1} \sum_{v \in V_i} \text{outdeg}(v)/4 = |\phi_0^{-1}(\bot)|/4 .$$

This implies that at least 1/4 of the uncolored edges become colored during the iteration. Thus, after $O(\log n)$ iterations, all edges are colored.

3.2Load Balancing the Blocking Edges

By Lemma 3.1, Line 12 of Algorithm 1 is executed for at most $t = K\Delta^4 \log n$ times for some constant K > 0. Note that in each execution of Line 12, the alternating path of each fan in $\mathcal{F}_{i,\alpha,\beta,j}$ does not intersect alternating paths of other fans in $\mathcal{F}_{i,\alpha,\beta,j}$.

Let $\ell(v)$ denote the load of v, which is defined to be the number of incident edges to v colored in \star . Also, let $\ell(G) = \max_{v \in V} \ell(v)$. The hope is that at the end of the algorithm, $\ell(G)$ will be small. Since the maximum degree induced by \star -edges is $\ell(G)$, we can recolor \star -edges using $2\ell(G) - 1$ additional colors using the algorithm of Panconesi and Rizzi [PR01] in $O(\ell(G) + \log^* n)$ rounds.

In each execution of Line 12 of Algorithm 1, each vertex can only increase its \star degree by one. This is because the alternating paths of the fans in $\mathcal{F}_{i,\alpha,\beta,j}$ do not intersect, and thus each vertex can be passed by at most one alternating path. Recall that t is the number of times Line 12 of Algorithm 1 is executed. Therefore, each vertex can be passed by at most t alternating paths throughout the algorithm.

On the other hand, for each alternating path P of length at least T, it has T choices to place the blocking edge on P(T). If we place the blocking edge uniformly at random on each path, the expected load for each vertex v is, $\mathbb{E}[\ell(v)] \leq \frac{2t}{T} = O(1/\lambda)$, provided $T \geq t \cdot \lambda$. We can show by Chernoff bound that $\ell(G) = O(\log n/(\log \lambda + \log \log n))$ w.h.p. However, to obtain a $(\Delta + 2)$ -edge coloring, we need to make $\ell(G) = 1$ so that the \star -edges form a matching. In the rest of this section, we show how to achieve $\ell(G) = 1$ using randomization and $\ell(G) = O(\log n/(\log \lambda + \log \log n))$ without randomization. Achieving $\ell(G) = 1$: Let $T = K\Delta^7 \cdot t = K\Delta^{11} \log n$ for some sufficiently large constant K. Instead of randomly placing the blocking edge, we do the following. Given a path P, let $E_0(P)$ be the set of edges whose endpoints have zero loads. In each alternating path P, we choose an edge from $E_0(P)$ uniformly at random to be the blocking edge. We will show that at any time during the algorithm, for every path P of length T (P is not necessarily an alternating path), $|E_0(P)| \ge T/15$ w.h.p. Therefore, given an alternating path of length at least T, it is always possible to choose a blocking edge whose both endpoints have zero loads w.hp. If some alternating path of length at least T with $|E_0(P)| < T/15$ occurs during Line 12, we say the algorithm *fails* and let our graph and coloring *freeze* there. In that case, P will not be processed. If the algorithm did not fail, then we have achieved $\ell(G) = 1$.

Lemma 3.2. After the last execution of Line 12 of Algorithm 1, it holds w.h.p. that for every path P of length T, we have that $|E_0(P)| \ge T/15$.

Proof. Fix a path P of length T. We will show that the probability $\Pr(|E_0(P)| < T/15) \le 1/(\operatorname{poly}(n) \cdot \Delta^T)$ after the *t*'th execution of Line 12 of Algorithm 1. By taking a union bound over all possible paths of length T, we will conclude that w.h.p. for every path P of length T, $|E_0(P)| < T/15$. This also implies the algorithm did not fail throughout execution $1 \ldots t$ w.h.p.

We say a vertex v is occupied if $\ell(v) \ge 1$. Otherwise, it is *empty*. Similarly, we say an edge is occupied if any of its endpoints is occupied. Otherwise, it is empty.

We say an alternating path is adjacent to an edge uv in P if u or v is in the alternating path. We divide all the alternating paths from the beginning up to after the t'th execution of Line 12 of Algorithm 1 that intersect P into two groups:

$$\mathcal{P}_1 = \{ \text{alt. paths that are adjacent to at least } \frac{T}{1000} \text{ edges in } P \}$$

 $\mathcal{P}_2 = \{ \text{alt. paths that are adjacent to fewer than } \frac{T}{1000} \text{ edges in } P \}.$

Because the alternating paths in each execution of Line 12 of Algorithm 1 are vertex-disjoint, for each edge in P, at most 2 alternating paths are adjacent to it in each execution. Thus, at most 2talternating paths are adjacent to that edge in total. This implies the number of pairs $(e, \tilde{P}) \in P \times \mathcal{P}_1$ where \tilde{P} is adjacent to e is at most 2tT. Since each alternating path in \mathcal{P}_1 consumes at least T/1000pairs, we have

$$|\mathcal{P}_1| \le \frac{2tT}{T/1000} = 2000t$$
.

Note that we only place at most one blocking edge on an alternating path (i.e., if the path is truncated). Each blocking edge can cause at most three edges in P to be occupied (this corresponds to the case the blocking edge is in P). Thus, the alternating paths in \mathcal{P}_1 can cause at most 6000t edges in P to be occupied.

Now we bound the number of occupied edges in P caused by the blocking edges of alternating paths in \mathcal{P}_2 . Let $P = u_0 \ldots u_T$. Let $E'(P) = \{u_0 u_1, u_2 u_3, \ldots, u_{2 \cdot (\lceil T/2 \rceil - 1)} u_{2 \cdot (\lceil T/2 \rceil - 1) + 1}\}$ denote the set of every other edge in P so that $|E'(P)| = \lceil T/2 \rceil$. We say an edge is **adjacent** to E'(P) if it shares an endpoint with an edge in E'(P).

We say a vertex u is in E'(P) if it is an endpoint of an edge in E'(P). We orient every edge (including those not in P) e = uv that is adjacent to E'(P) as follows: Without loss of generality, let u be the vertex that is in E'(P). If v is not in E'(P), we orient e from u to v. Otherwise, we orient the edge from the vertex that appears earlier in P to the vertex that appears later in P.

Given an edge $e = uv \in E'(P)$, we say e is **outwardly occupied** if any of the **outgoing edges** from u or v is a blocking edge of some alternating path in \mathcal{P}_2 . Next, we will bound the number of outwardly occupied edges in E'(P).

The reason we are using "outwardly occupied" instead of "occupied" is because it resolves some dependency issues. In particular, each alternating path in \mathcal{P}_2 can only make at most one edge in E'(P) outwardly occupied, while it can make two edges occupied. Moreover, as we will see later, the number of occupied edges in E'(P) is at most 3 times the number of outwardly occupied edges.

Now we hope to upper bound the total number of outwardly occupied edges $X = \sum_{e \in E'(P)} X_e$ where X_e is the indicator variable for the event e is outwardly occupied. Consider an edge e, the following is a rough argument to bound $\Pr(X_e = 1)$. The probability that an alternating path makes e outwardly occupied is O(1/T), since it has $\Omega(T)$ empty edges to choose from. Moreover, there are at most O(t) alternating paths that are adjacent to e throughout the algorithm. Thus, by taking a union bound over these paths, we have $\Pr(X_e = 1) = O(t) \cdot O(1/T) = O(t/T)$. As a result, $\mathbb{E}[X] = O(T \cdot (t/T)) = O(t)$. However, it will be difficult to obtain a tail bound on X directly due to dependency issues. Since we need to apply the strongest form of Chernoff bound in order to obtain a very sharp bound $(1/(\text{poly}(n) \cdot \Delta^T))$ on the probability that X deviates from $\mathbb{E}[X]$, we couple random variables $\{X_e\}_{e \in E'(P)}$ with independent Bernoulli random variables $\{Y_e\}_{e \in E'(P)}$ with success probability at most 150t/T in the following lemma whose proof is deferred to the end of this proof.

Lemma 3.3. For every $e \in E'(P)$, let $X_e = 1$ if e is outwardly occupied by a blocking edge in an alternating path in \mathcal{P}_2 and $X_e = 0$ otherwise. Let $\{Y_e\}_{e \in E'(P)}$ be independent Bernoulli random variables with success probability at most 150t/T. Let $X = \sum_{e \in E'(P)} X_e$ and $Y = \sum_{e \in E'(P)} Y_e$. For any $M \ge 0$, we have $\Pr(X > M) \le \Pr(Y > M)$.

Therefore,

$$\mathbb{E}[Y] \le \lceil T/2 \rceil \cdot \frac{150t}{T} < 100t \; .$$

Let $\delta = (T/700t) - 1$, and so $(1 + \delta) = (T/700t)$. Let M = 100t be an upper bound of $\mathbb{E}[Y]$. We can apply a variant of Chernoff bound (where one could replace $\mathbb{E}[Y]$ by M if $M \ge \mathbb{E}[Y]$, see Corollary 4 of [PS15]) to obtain:

$$\begin{aligned} \Pr(Y > T/7) &= \Pr(Y > (T/700t) \cdot 100t) \\ &= \Pr(Y > (1 + \delta) \cdot M) \\ &\leq \Pr(Y > \mathbb{E}[Y] + \delta M) \\ &\leq \left(\frac{e}{1 + \delta}\right)^{(1+\delta) \cdot M} \end{aligned} \qquad \text{Corollary 4 of [PS15]} \\ &= \left(\frac{e}{1 + \delta}\right)^{T/7} \\ &\leq \left(\frac{1}{e\Delta^7}\right)^{T/7} \qquad (1 + \delta) = \frac{T}{700t} > e^2 \cdot \Delta^7 \\ &\leq \frac{1}{\text{poly}(n)} \cdot \frac{1}{\Delta^T} \qquad T = \Omega(\log n) \;. \end{aligned}$$

By Lemma 3.3, $\Pr(X > T/7) \le \Pr(Y > T/7) \le 1/\operatorname{poly}(n) \cdot 1/\Delta^T$.

Suppose that the number of outwardly occupied edges in E'(P) is at most T/7. Consider an occupied edge that is not outwardly occupied. Its adjacent blocking edge must be adjacent to an edge $e \in E'(P)$ that is outwardly occupied. Moreover, each $e \in E'(P)$ can only be adjacent to at most two blocking edges. Therefore, the number of occupied edges is at most T/7 + 2(T/7) + 6000t = 3T/7 + 6000t. The number of empty edges in E'(P) must be at least $\lceil T/2 \rceil - 3T/7 - 6000t > T/15$. Thus, the probability that there are fewer than T/15 empty edges on P is at most $1/(\text{poly}(n) \cdot \Delta^T)$.

There are at most $n \cdot \Delta^{T-1}$ paths of length T. By taking a union bound all length T paths, the probability that there is a length T path with fewer than T/15 empty edges is at most $(n \cdot \Delta^{T-1})/(\text{poly}(n) \cdot \Delta^T) = 1/\text{poly}(n)$. This also implies that the algorithm did not fail w.h.p.

Proof of Lemma 3.3. First, we create i.i.d. random variables $\theta_{e,i}$ for each edge $e \in E'(P)$ and $1 \leq i \leq 2t$, where $\Pr(\theta_{e,i} = 1) = 75/T$. For each $e \in E'(P)$, we let $Y_e = \bigvee_{1 \leq i \leq 2t} (\theta_{e,i} = 1)$. Clearly, $\Pr(Y_e = 1) \leq 2t \cdot (75/T) = 150t/T$ and $\{Y_e\}_{e \in E'(P)}$ are independent.

The next step is to couple $\{X_e\}_{e \in E'(P)}$ with $\{\theta_{e,i}\}_{e,i}$. Recall that in our algorithm, each alternating path $\tilde{P} \in \mathcal{P}_2$ (that is processed) selects the blocking edge uniformly at random from $E_0(\tilde{P})$ (unoccupied edges at the execution when \tilde{P} occurs). Now we consider a modified algorithm that has the same outcome with the original algorithm. It allows us to do the coupling more easily. Given $e \in E'(P)$, let $out_{\tilde{P}}(e)$ denote the set of edges that are in $E_0(\tilde{P})$ and are outgoing edges from any endpoint of e. Note that if $e, e' \in E'(P)$ and $e \neq e'$, then $out_{\tilde{P}}(e)$ and $out_{\tilde{P}}(e')$ are disjoint.

Let $E'_{\tilde{P}}(P) \subseteq E'(P)$ denote the set of edges with non-empty $out_{\tilde{P}}(e)$. Since $\tilde{P} \in \mathcal{P}_2$, $|E'_{\tilde{P}}(P)| \leq T/1000$. Also, since $|out_{\tilde{P}}(e)| \leq 4$, we have $\sum_{e \in E'_{\tilde{P}}(P)} |out_{\tilde{P}}(e)| \leq 4T/1000$.

The modified algorithm maintains a counter, occurred, which is initially set to 0. Then it goes through each edge $e \in E'_{\tilde{P}}(P)$ and does the following: Throw a coin independently with success probability $|out_{\tilde{P}}(e)|/(|E_0(\tilde{P})| - occurred))$. If it succeeds, it selects a blocking edge randomly from $out_{\tilde{P}}(e)$ and then it is done with processing \tilde{P} . Otherwise, it sets occurred $\leftarrow occurred + |out_{\tilde{P}(e)}|$ and continues to process the next edge in $E'_{\tilde{P}}(P)$. If the algorithm did not successfully select a blocking edge after it processed every edge in $E'_{\tilde{P}}(P) \subseteq E'(P)$, it selects an edge from $E_0(\tilde{P}) \setminus (\bigcup_{e \in E'_{\tilde{P}}(P)} out_{\tilde{P}}(e))$ uniformly at random as the blocking edge.

It is easy to see that in the modified algorithm, the probability that any edge in $E_0(P)$ is selected as the blocking edge is exactly $1/|E_0(\tilde{P})|$. Therefore, the modified algorithm is equivalent to the original algorithm where \tilde{P} selects a random edge in $E_0(\tilde{P})$ as the blocking edge.

Now we couple the randomness of the modified algorithm with $\{\theta_{e,i}\}_{i,e}$. When the modified algorithm processes an edge $e \in E'_{\tilde{P}}(P)$, instead of throwing a new coin with probability $|out_{\tilde{P}}(e)|/(|E_0(\tilde{P})| - occurred)$, we will couple it with the randomness from $\theta_{e,i}$. For each $e \in E'(P)$, we maintain a counter q(e) that points to the smallest index such that $\theta_{e,q(e)}$ is not yet exposed. Initially, q(e) are 1 for every $e \in E'(P)$. Now when the algorithm processes $\tilde{P} \in \mathcal{P}_2$ and an edge $e \in E'_{\tilde{P}}(P)$, we will expose $\theta_{e,q(e)}$. Let \mathcal{E}_1 denote the event that $\theta_{e,q(e)} = 1$. Also, we set $q(e) \leftarrow q(e) + 1$. Let \mathcal{E}_2 denote the event that a coin with success probability $\frac{1}{\Pr(\mathcal{E}_1)} \cdot \frac{|out_{\tilde{P}}(e)|}{(|E_0(\tilde{P})| - occurred)}$ succeeds. Note that $\Pr(\mathcal{E}_1) \cdot \Pr(\mathcal{E}_2) = |out_{\tilde{P}}(e)|/(|E_0(\tilde{P})| - occurred)$. The original coin with probability $|out_{\tilde{P}}(e)|/(|E_0(\tilde{P})| - occurred)$ will be simulated by the trial that both the two events \mathcal{E}_1 and \mathcal{E}_2 happen.

Note that given $e \in E'(P)$, at most 2t alternating paths in \mathcal{P}_2 may be adjacent to e in total. This implies $\theta_{e,q(e)}$ is always well-defined when the algorithm refers to it. It remains to show that $\Pr(\mathcal{E}_2)$ is a valid probability, namely, $|out_{\tilde{P}}(e)|/(|E_0(\tilde{P})| - occurred) \leq \Pr(\mathcal{E}_1)$.

$$\begin{aligned} \Pr(\mathcal{E}_1) &= 75/T \\ &\geq \left(\frac{|out_{\tilde{P}}(e)|}{T/15 - 4T/1000}\right) & |out_{\tilde{P}}(e)| \leq 4 \\ &\geq \left(\frac{|out_{\tilde{P}}(e)|}{T/15 - occurred}\right) & occurred \leq 4T/1000 \\ &\geq \left(\frac{|out_{\tilde{P}}(e)|}{|E_0(\tilde{P})| - occurred}\right) & |E_0(\tilde{P})| \geq T/15 . \end{aligned}$$

The second to the last line holds since occurred $\leq \sum_{e \in E'_{\tilde{P}}(P)} |out_{\tilde{P}}(e)| \leq 4T/1000$. Also, the last line holds since if $|E_0(\tilde{P})| < T/15$, the algorithm freezes and \tilde{P} will not be processed.

Finally, we point out that the coupling argument implies whenever $X_e = 1$, it must be the case that $\theta_{e,i} = 1$ for some $1 \le i \le 2t$. This in turns implies that $Y_e = 1$ and thus $X_e \le Y_e$. Therefore, for any M > 0, $\Pr(X > M) \le \Pr(Y > M)$.

By Lemma 3.1, Line 9 is executed $O(\Delta^4 \log n)$ times. Each execution uses $O(T \cdot \log n)$ time, where $T = O(\Delta^{11} \log n)$. By Lemma 3.1 again, Line 12 is also executed $O(\Delta^4 \log n)$ times. Each execution takes O(T) time. Therefore, the number of rounds used by our algorithm to obtain a $(\Delta + 2)$ -edge-coloring is $O(\Delta^{15} \log^3 n)$. The following lemma (whose proof we postpone to the end of the section) completes the proof of Theorem 1.1, which shows the algorithm can be converted to a $(1 + \epsilon)\Delta$ -edge-coloring algorithm that runs in $O((1/\epsilon)^{15} \log^3 n)$ rounds for any $\epsilon \geq 2/\Delta$.

Lemma 3.4. Suppose that there is a deterministic (randomized) $(\Delta + z)$ -edge-coloring algorithm that runs in $O(\Delta^x \log^y n)$ rounds. Then there exists a deterministic (randomized) $(1 + \epsilon)\Delta$ -edge-coloring algorithm that runs in $O((1/\epsilon)^x \cdot (z+1)^x \cdot \log^y n) + \tilde{O}(\log^2 \Delta \cdot \epsilon^{-1} \cdot \log \epsilon^{-1} \cdot \log n)$ rounds for any $\epsilon \geq z/\Delta$, where $\tilde{O}(\cdot)$ hides a poly(log log n) factor.

Achieving $\ell(G) = O(\log n/(\log \lambda + \log \log n))$ deterministically: We will choose $T = 2t \cdot \lambda$. Let $(1 + \delta) = K \cdot \lambda \cdot \log n/(\log \lambda + \log \log n)$ for some constant K > 0. Given an alternating path P with length at least T, let $\ell_q(v)$ denote the load of v after q'th execution of Line 12 of Algorithm 1.

During the q'th execution of Line 12 of Algorithm 1, we select the blocking edge as follows. Given an edge e = uv, define the cost c(e) to be $(1 + \delta)^{\ell_{q-1}(v)} + (1 + \delta)^{\ell_{q-1}(u)}$. Given an alternating path P, we select the edge $e \in P(T)$ with the minimum cost c(e) to be the blocking edge. We will show that this greedy algorithm achieves a maximum load of $O(\log n/(\log \lambda + \log \log n))$.

Lemma 3.5. If $T = 2t\lambda$, the greedy algorithm achieves a load of $O(\log n/(\log \lambda + \log \log n))$.

Proof. We use the method of conditional expectation to show that the greedy algorithm is a derandomization of the randomized algorithm. Define the pessimistic estimator

$$\Phi(q) = \sum_{v \in V} \frac{(1+\delta)^{\ell_q(v)} \cdot (1+\delta \cdot (2/T))^{t-q}}{(1+\delta)^{(1+\delta)/\lambda}} \,.$$

We show that $\Phi(q+1) \leq \Phi(q)$ by the choice of the greedy algorithm. Consider the (q+1)'th iteration of Line 12 of Algorithm 1.

Let B denote a set of blocking edges chosen during (q + 1)'th execution. Given an edge e, if v is an endpoint of e, we write $v \in e$. Moreover, we write $v \in B$, if there exists $e \in B$ such that $v \in e$. We use $[v \in B]$ denote the indicator variable for the event $v \in B$. Let Q(B) denote the following quantity:

$$Q(B) = \sum_{v \in V} (1+\delta)^{\ell_q(v) + [v \in B]}$$

Suppose that we choose each blocking edge randomly in (q + 1)'th execution of Line 12, we have

$$\begin{split} \mathbb{E}[Q(B)] \\ &= \sum_{v \in V} \left(\Pr(v \in B) \cdot (1+\delta)^{\ell_q(v)+1} + \Pr(v \notin B) \cdot (1+\delta)^{\ell_q(v)} \right) \\ &= \sum_{v \in V} (1+\delta)^{\ell_q(v)} \cdot \left(\Pr(v \in B) \cdot (1+\delta) + (1-\Pr(v \in B)) \right) \\ &= \sum_{v \in V} (1+\delta)^{\ell_q(v)} \cdot (1+\delta \cdot \Pr(v \in B)) \\ &\leq \sum_{v \in V} (1+\delta)^{\ell_q(v)} \cdot \left(1+\delta \cdot \frac{2}{T} \right) \; . \end{split}$$

Observe that given e and e', if c(e) < c(e'), then $\sum_{v \in V} (1+\delta)^{\ell_q(v)+[v \in e]} < \sum_{v \in V} (1+\delta)^{\ell_q(v)+[v \in e']}$. Moreover, the alternating paths generated in the (q+1)'s execution of Line 12 do not intersect. Therefore, our greedy strategy is always choosing a set of blocking edges to minimize Q(B). Thus,

$$\begin{split} \Phi(q+1) &= \frac{(1+\delta \cdot (2/T))^{t-q-1}}{(1+\delta)^{(1+\delta)/\lambda}} \cdot \min_{B} Q(B) \\ &\leq \frac{(1+\delta \cdot (2/T))^{t-q-1}}{(1+\delta)^{(1+\delta)/\lambda}} \cdot \mathbb{E}_{B}[Q(B)] \\ &\leq \frac{(1+\delta \cdot (2/T))^{t-q-1}}{(1+\delta)^{(1+\delta)/\lambda}} \cdot \sum_{v \in V} (1+\delta)^{\ell_{q}(v)} \left(1+\delta \cdot (2/T)\right) \\ &= \Phi(q) \;. \end{split}$$

Let $X = \log \lambda + \log \log n$. We have

$$\begin{split} &\sum_{v \in V} (1+\delta)^{\ell_t(v)-(1+\delta)/\lambda} = \Phi(t) \leq \Phi(t-1) \leq \ldots \leq \Phi(0) \\ &= \sum_{v \in V} \frac{(1+\delta \cdot (2/T))^t}{(1+\delta)^{(1+\delta)/\lambda}} \\ &\leq \sum_{v \in V} \frac{e^{\delta \cdot (2t/T)}}{(1+\delta)^{(1+\delta)/\lambda}} \qquad (\text{since } 1+x \leq e^x) \\ &\leq \sum_{v \in V} \frac{e^{\delta/\lambda}}{(1+\delta)^{(1+\delta)/\lambda}} \\ &\leq \sum_{v \in V} \left[\frac{e}{K\lambda \log n/(\log \lambda + \log \log n)} \right]^{K \log n/(\log \lambda + \log \log n)} \\ &= \sum_{v \in V} \exp(-(X + \log K - \log(X) - 1) \cdot K \log n/X) \\ &\leq \sum_{v \in V} \exp(-(K/2) \log n) \qquad (\text{for sufficiently large } X) \\ &\leq \sum_{v \in V} \frac{1}{\text{poly}(n)} < 1 \,. \end{split}$$

This implies $\ell_t(v) \leq (1+\delta)/\lambda = O(\log n/(\log \lambda + \log \log n))$ for all v. Thus, we have achieved $\ell(G) = O(\log n/(\log \lambda + \log \log n))$.

The running time is dominated by Line 9, which is executed $O(\Delta^4 \log n)$ times. Each execution uses $O(T \cdot \log n)$ time. The total running time for obtaining a $\Delta + O(\log n/(\log \lambda + \log \log n))$ coloring is $O(\Delta^4 T \cdot \log^2 n) = O(\lambda \cdot \Delta^8 \log^3 n)$. We show how to reduce the running time to $O(\lambda \cdot \Delta^6 \log^3 n)$ in the next section. By Lemma 3.4, it can be converted to a deterministic $(1 + \epsilon)\Delta$ -edge-coloring algorithm that runs in $O(\lambda \cdot (1/\epsilon)^6 \log^9 n)$ algorithm for $\epsilon = \Omega(\log n/(\Delta(\log \lambda + \log \log n)))$.

Proof of Lemma 3.4. W.L.O.G. we may assume $\epsilon \leq 1/8$. If $\Delta \leq (16(z+8))/\epsilon$, then by applying the $(\Delta + z)$ -edge-coloring algorithm, we obtain an $(1+\epsilon)\Delta$ -edge-coloring in $O(\Delta^x \log^y n) = O((1/\epsilon)^x \cdot (z+1)^x \log^y n)$ rounds, since $z \leq \epsilon \Delta$ and $\epsilon \leq (16(z+8))/\Delta$.

Otherwise, let $\epsilon' = \epsilon/(32 \log_{4/3} \Delta)$. By using the splitting algorithm of [GHK⁺17], we obtain two subgraphs of maximum degree $(1/2 + \epsilon')\Delta + 4$ in $\tilde{O}(\epsilon'^{-1} \cdot \log \epsilon'^{-1} \cdot \log n)$ rounds. Let $\Delta_0 = \Delta$ and $\Delta_i = (1/2 + \epsilon')\Delta_{i-1} + 4$. If we apply the algorithm recursively for h levels, we obtain 2^h subgraphs whose maximum degree are bounded by Δ_h . We choose h to be the smallest number that $(1/2 + \epsilon')^h \cdot \Delta \in$ $[(8(z+8))/\epsilon, (16(z+8))/\epsilon]$. Note that $h \leq \log_{4/3} \Delta$, since $\epsilon' \leq \epsilon \leq 1/4$. Note that

$$\begin{aligned} \Delta_h &= (1/2 + \epsilon')^h \cdot \Delta + \sum_{i=0}^{h-1} (1/2 + \epsilon')^i \cdot 4 \\ &\leq (1/2 + \epsilon')^h \cdot \Delta + \sum_{i=0}^{h-1} (3/4)^i \cdot 4 \qquad (\text{since } \epsilon' < 1/4) \\ &\leq (1/2 + \epsilon')^h \cdot \Delta + 16 \\ &\leq (1 + \epsilon/4)(1/2 + \epsilon')^h \cdot \Delta \\ &\quad (\text{since } (1/2 + \epsilon')^h \cdot \Delta \ge (8(z+8))/\epsilon \ge 64/\epsilon) . \end{aligned}$$

We run the $(\Delta_h + z)$ -edge-coloring algorithm on the 2^h subgraphs using disjoint palettes to obtain a $(\Delta_h + z)$ -edge-coloring for each. The total number of colors that are used is

$$2^{h} \cdot (\Delta_{h} + z) \leq 2^{h} \cdot (1 + \epsilon/4) \cdot \Delta_{h} \qquad (\text{since } \Delta_{h} \geq (1/2 + \epsilon')^{h} \cdot \Delta \geq 8z/\epsilon)$$

$$= 2^{h} \cdot (1 + \epsilon/4)^{2} \cdot (1/2 + 2\epsilon')^{h} \cdot \Delta \qquad (\text{since } \Delta_{h} \leq (1 + \epsilon/4) \cdot (1/2 + 2\epsilon')^{h} \cdot \Delta)$$

$$= (1 + \epsilon/4)^{2} \cdot (1 + 4\epsilon')^{h} \cdot \Delta \qquad (\text{since } (1 + a)^{b} \leq 1 + 2ab \text{ for } 0 \leq ab \leq 1)$$

$$\leq (1 + \epsilon/4)^{2} \cdot (1 + \epsilon/4) \cdot \Delta \qquad (\text{since } \epsilon' = \epsilon/(32 \log_{4/3} \Delta) \leq \epsilon/(32h))$$

$$\leq (1 + \epsilon)\Delta \qquad (\text{when } \epsilon < 1/8) .$$

Since $\Delta_h = O(z/\epsilon)$, the running time is $\tilde{O}(\log \Delta \cdot \epsilon'^{-1} \cdot \log \epsilon'^{-1} \cdot \log n) + O((1/\epsilon)^x (z+1)^x \log^y n)$. \Box

4 Faster Algorithms

In this section, we provide more efficient algorithms that use fewer number of rounds. Our technique is based on another type of fans that was introduced by Gabow et al. [GNK⁺85]. This allows us to reduce the number of rounds by a factor Δ^2 for general graphs and Δ^4 for bipartite graphs.

A reverse αB -fan with center v and leaves x_1, x_2, \ldots, x_k (where $B = \{\beta_1, \beta_2, \ldots, \beta_{k+1}\}$ and $k \ge 2$) satisfies: All edges vx_i are uncolored. The color α is missing at all leaves, i.e., $\alpha \in M(x_i)$ for all $1 \le i \le k$. The colors in B are missing at v, i.e., $\beta \in M(v)$ for all $\beta \in B$. Finally, there is an edge with color α that is incident to the center v, i.e., $\alpha \notin M(v)$. The above definition is well-defined since there must be at least k + 1 missing colors at v. We use normal fans and reverse fans to distinguish between the two types. We use fans to refer to fans of either types. We use incomplete vertices to refer to vertices that are incident to at least one uncolored edge.

Algorithm outline: We outline our algorithm (Algorithm 2) and analysis as follows. Similar to the previous section, we will show that each iteration of the **while loop** in Line 1 will color at least a constant fraction of the uncolored edges and therefore the algorithm will terminate after $O(\log n)$ iterations. Roughly speaking, we iterate through the incomplete vertices that miss the colors $\alpha = 1, 2, \ldots, \Delta + 1$ (i.e., the **for loop** in Line 5). For each color α , we try to make a maximal set of vertex-disjoint α -fans (i.e., every vertex in $V_{\gamma,\alpha}$ belongs to a fan) that consists of normal fans and reverse fans. This is done by merging intersecting fans in a clever way. Then, the algorithm tries to repair these fans simultaneously. The idea is that we can repair the fans simultaneously in fewer rounds since they are vertex-disjoint. There are two new key challenges in the LOCAL model (beside long alternating paths):

• We need to implement the merge step in the distributed setting efficiently. Gabow et al. [GNK⁺85] only provided a sequential algorithm for this.

Algorithm	2	Faster	Distributed	Fan	Repair
-----------	----------	--------	-------------	-----	--------

1:	while there exists uncolored edges do
2:	Color the graph using $\Delta + 1$ colors.
3:	for $\gamma = \operatorname{color} 1, 2, \dots, \Delta + 1$ do
4:	Let V_{γ} be the vertices with color γ .
5:	for $\alpha = \operatorname{color} 1, 2, \dots, \Delta + 1$ do
6:	Let $E_{\gamma,\alpha} = \{ uv \in E \mid u \in V_{\gamma} \text{ and } uv \text{ is uncolored and } \alpha \in M(u) \cap M(v) \}.$
7:	Let $M_{\gamma,\alpha}$ be a maximal matching in $E_{\gamma,\alpha}$.
8:	Color $M_{\gamma,\alpha}$ with α .
9:	Let $V_{\gamma,\alpha}$ be the set of incomplete vertices in V_{γ} that now miss the color α .
10:	Color $G^4[V_{\gamma,\alpha}]$ using at most $4\Delta^4$ colors.
11:	${f for}\;i=1,2,\ldots,4\Delta^4\;{f do}$
12:	Grow a normal α -fan F_v from each vertex v with color i in $G^4[V_{\gamma,\alpha}]$.
13:	merge intersecting fans.
14:	for $\beta = \operatorname{color} 1, 2, \dots, \Delta + 1$ do
15:	Build a conflict graph $G_{\mathcal{F}_{\gamma,\alpha,\beta}}$ for normal and sub-reverse $\alpha\beta$ -fans in $\mathcal{F}_{\gamma,\alpha,\beta}$.
16:	Color $G_{\mathcal{F}_{\gamma,\alpha,\beta}}$ using $O(1)$ colors.
17:	for $\eta = \operatorname{color} 1, 2, \dots O(1)$ do
18:	Repair undestroyed normal and sub-reverse fans in $\mathcal{F}_{\gamma,\alpha,\beta}$ with color η in $G_{\mathcal{F}_{\gamma,\alpha,\beta}}$.
19:	Repair semi-destroyed sub-reverse fans in $G_{\mathcal{F}_{\gamma,\alpha,\beta}}$.

• With two types of fans, there are more cases to analyze when repairing a fan destroys the structure of another fan. This can become tricky since a destroyed reverse fan may lead to many edges remain uncolored.

4.1 Making a Maximal Set of α -Fans

We consider a fixed iteration of the outermost **while loop**. We first color the vertices using $\Delta + 1$ colors which could be done deterministically in $O(\Delta + \log^* n)$ rounds [BEK14]. We then process vertices in the same color class γ , denoted by V_{γ} , together. The reason for this step will become clear later.

Fix a nested iteration γ and α of the **for loop** in Line 3 and the **for loop** in Line 5. We first consider the set $E_{\gamma,\alpha}$ of all uncolored edges uv that are incident to one vertex in V_{γ} such that both uand v miss the color α . More formally,

 $E_{\gamma,\alpha} = \{uv \in E \mid u \in V_{\gamma} \text{ and } uv \text{ is uncolored and } \alpha \in M(u) \cap M(v)\}$.

We find a maximal matching $M_{\gamma,\alpha}$ in $E_{\gamma,\alpha}$ and color the matching with the color α in $O(\Delta + \log n)$ rounds [BEK14].

After coloring $M_{\gamma,\alpha}$, recall from Algorithm 2 that $V_{\gamma,\alpha}$ denotes the set of incomplete vertices in V_{γ} that miss the color α . Note that $V_{\gamma,\alpha}$ is an independent set since the γ -color vertices must not be neighbors.

We now describe a sequential procedure that finds a maximal set of α -fans. Specifially, each vertex in $V_{\gamma,\alpha}$ is either a center of a fan or a leaf of a reverse fan. Furthermore, all the fans are vertex-disjoint and the number of uncolored edges remain the same.

We want to maintain the following invariant: a) all active fans are vertex-disjoint and b) all vertices $v \in V_{\gamma,\alpha}$ that have been considered must either be a center of a normal fan or a leaf of a reverse fan. Now, suppose we consider the next vertex v in $V_{\gamma,\alpha}$ and grow a normal α -fan from v. We have argued that $V_{\gamma,\alpha}$ is an independent set. Assuming the invariant holds up to this point, we have:

Claim 4.1. Consider a normal α -fan F_v grown from $v \in V_{\gamma,\alpha}$ that intersects with a currently active α -fan F_u . If F_u is a normal fan, they must intersect at leaf node(s). On the other hand, if F_u is a reverse fan, then they only intersect at u.

Suppose F_v intersects with at least one other α -fan. Let the leaves of F_v be x_1, \ldots, x_k and let $w = x_j$ be the first leaf that is in an intersection, i.e., j is the smallest possible. Suppose F_v intersects with F_u at w.

If F_u is a normal fan, then we shift F_v and F_u from w and uncolor the edges vw and uw. We then deactivate F_u and F_v and grow a reverse fan centering at w with the leaves u and v. We note that w cannot miss α otherwise the matching $M_{\gamma,\alpha}$ is not maximal.

If F_u is a reverse fan, then note that u = w based on Claim 4.1. We shift F_v from u and uncolor the edge vu. We then deactivate F_v and add vu to F_u .

The fact that w is the first leaf that is in an intersection is important since F_v may intersect with many current active fans and shifting at another leaf vertices may destroy other fans that are not F_v and F_u . It is easy to see that the invariant that the number of uncolored edges remain the same and the fans are vertex-disjoint holds after we process each vertex v in $V_{\gamma,\alpha}$. Hence, we obtain a maximal set of vertex-disjoint α -fans.

Merging fans the distributed setting: We still need to implement this idea in the LOCAL model. We now describe the merge step that corresponds to Line 10 to Line 13 in our algorithm.

First, we obtain a 4-hop coloring of $V_{\gamma,\alpha}$ using at most $4\Delta^4$ colors which can be done in $O(\Delta^4 + \log^* n)$ rounds [BEK14]. In particular, this is a vertex coloring of the graph $G^4[V_{\gamma,\alpha}]$ in which there is an edge between uv if the distance between u and v is less than five in G. We then go through vertices in each color class $i = 1, \ldots, 4\Delta^4$ and grow a normal α -fan from vertices in that class.

The key observation is as follows. Suppose u and v are from the same color class i. Since $\operatorname{dist}_G(u, v) > 4$, the normal fans F_u and F_v that are grown in Line 12 do not intersect and furthermore, F_u and F_v cannot intersect the same currently active α -fan F_w . Therefore, we can simultaneously merge the newly grown fans with the existing active fans as described in a constant number of rounds (Algorithm 3).

Algorithm 3 merge

3: for each F_v in $A \setminus B$ that intersects with another fan in B do	
4: Let x be the first leaf of F_v that is in an intersection.	
5: Let F_w be the fan that intersects with F_v at x .	
6: if F_w is a normal fan then	
7: Shift F_w and F_v from x .	
8: Uncolor wx and vx .	
9: Grow a reverse α -fan centering at x with leaves w and v.	
10: Deactivate F_v and F_w .	
11: else	
12: Shift F_v from x .	
13: Uncolor vx .	
14: Add vx to F_w (note that $w = x$).	
15: Deactivate F_v .	

Hence, we obtain a maximal set of vertex-disjoint α -fans after we iterate through all color classes $i = 1, 2, \ldots, 4\Delta^4$. We summarize this as the following claim.

Claim 4.2. After the merge step (just before Line 14), we have a maximal set of vertex-disjoint α -fans with respect to $V_{\gamma,\alpha}$.

4.2 Repairing Fans

Recall that we currently fix a nested iteration γ and α of the for loops in Line 3 and Line 5. We furthermore have a maximal set of vertex-disjoint α -fans.

Sub-reverse fans: For the sake of a simpler analysis, we further decompose a reverse αB -fan that F_w with k leaves into $\lfloor k/2 \rfloor$ edge-disjoint sub-reverse fans. The first sub-reverse fan consists of $\{w, x_1, x_2\}$. The second sub-reverse fan consists of $\{w, x_3, x_4\}$, and so on. The last sub-reverse fan may consist of three leaves if k is odd.

Recall that $B = \{\beta_1, \beta_2, \dots, \beta_{k+1}\}$ is a set of k+1 colors missing at w. We associate β_1 with the first sub-reverse fan, β_2 with the second sub-reverse fan, and so on. We use $F_{w,1}, F_{w,2}, \dots, F_{w,\lfloor k/2 \rfloor}$ to denote these sub-reverse fans. In particular, we call $F_{w,i}$ a sub-reverse $\alpha\beta_i$ -fan.

We now focus on normal fans and sub-reverse fans. Let $\alpha\beta$ -fans denote normal $\alpha\beta$ -fans or subreverse $\alpha\beta$ -fans. We consider a simple repair for a sub-reverse $\alpha\beta_i$ -fan with center vertex w and (at least) two leaf vertices x_{2i}, x_{2i+1} . Consider the maximal alternating $\alpha\beta_i$ -path P_w from the center w. We can first $\operatorname{augment}(P_w)$. If P_w does not end at x_{2i} , then color wx_{2i} with α . Otherwise, if P_w ends at x_{2i} , then color wx_{2i+1} with α . This repair will color an uncolored edge. After the algorithm repairs $F_{w,i}$, then either wx_{2i} or wx_{2i+1} will be colored with the color α . Therefore, we still have $\alpha \notin M(w)$. As a result, other sub-reverse fans centering at w are not destroyed.

Claim 4.3. Repairing a sub-reverse fan $F_{w,i}$ does not destroy another sub-reverse fan $F_{w,j}$.

Repairing fans in the distributed setting: As outlined in the previous section, if the length of a maximal alternating path $P_v = (v_0, v_1, ...)$ is longer than T, we truncate P_v and pick a blocking edge (v_{i-1}, v_i) accordingly (i.e., $i \leq T$ is chosen based on whether the goal is a deterministic algorithm with $\ell(G) = O(\log n/(\log \lambda + \log \log n))$ or a randomized algorithm with $\ell(G) = 1$). We then $\operatorname{augment}(P_v(i))$.

We make the following observation: Based on the definition, a vertex in a sub-reverse $\alpha\beta$ -fan misses either β if it is the center or α if it is a leaf. Therefore, a maximal $\alpha\beta$ -path cannot intersect a sub-reverse $\alpha\beta$ -fan except at the last vertex of the path. Thus, a truncated $\alpha\beta$ -path $P_v(T)$ where $|P_v| > T$ cannot intersect with a sub-reverse $\alpha\beta$ -fan.

Next, let us fix an iteration β in the **for loop** of Line 14. We again observe that the alternating paths of $\alpha\beta$ -fans do not intersect. Augmenting an alternating path can only destroy the structure of at most one other fan. This happens when an alternating path of length at most T ends at another fan.

Let the set of the current normal and sub-reverse $\alpha\beta$ -fans be $\mathcal{F}_{\gamma,\alpha,\beta}$. We build a conflict graph on $G_{\mathcal{F}_{\gamma,\alpha,\beta}} = (V_{\mathcal{F}_{\gamma,\alpha,\beta}}, E_{\mathcal{F}_{\gamma,\alpha,\beta}})$, where $V_{\mathcal{F}_{\gamma,\alpha,\beta}} = \{F_v \mid F_v \in \mathcal{F}_{\gamma,\alpha,\beta}\}$ and $(F_u, F_w) \in E_{\mathcal{F}_{\gamma,\alpha,\beta}}$ if $|P_u| \leq T$ and P_u ends at a vertex in F_w . If $|P_u| > T$ and the truncated path $P_u(T)$ ends at any node of F_w , then F_w must be a normal fan and furthermore the modification on $P_u(T)$ will not affect the structure of F_w as pointed out in the previous section.

Since each vertex in $G_{\mathcal{F}_{\gamma,\alpha,\beta}}$ contributes at most one edge, the arboricity of $G_{\mathcal{F}_{\gamma,\alpha,\beta}}$ is constant. Therefore, as argued in the previous section, we can build and color the conflict graph using O(1) colors in $O(T \log n)$ rounds [BE10]. Now we iterate through each color j of the coloring and process the fans colored in j together. Since the alternating paths of the fans colored in j neither intersect nor end at another fan, we can repair them simultaneously.

There are three cases to take into consideration when an alternating path of length at most T ends at another fan F_w :

- Case 1: F_w is a normal $\alpha\beta$ -fan.
- Case 2: F_w is a sub-reverse $\alpha\beta$ -fan and the alternating path ends at a leaf of F_w .
- Case 3: F_w is a sub-reverse $\alpha\beta$ -fan and the alternating path ends at w.

We say that F_w is destroyed if Case (1) or Case (2) happens. This is because augmenting that path will destroy the structure of F_w .

We however cannot ignore Case (3) since it may destroy $\Omega(\Delta)$ sub-reverse fans centering at w. If case (3) occurs, we say that F_w is *semi-destroyed*. Fortunately, it is relatively easy to repair all semi-destroyed fans simultaneously. As argued above, if Case (3) happens, the final edge of the alternating path must be a color- α edge. After the alternating path is augmented, the final edge will become a color- β edge and therefore α becomes missing at w. Hence, we can color an edge in F_w with α . We call this step repairing semidestroyed sub-reverse fans. Note that this step can be done simultaneously among $\alpha\beta$ -fans. Based on Claim 4.3, no other sub-reverse fan centering at w is destroyed during this step. We also note that each shift operation to repair a normal fan will not destroy another α -fan since all the fans are vertex-disjoint.

Analysis: Now, we will show that each iteration in the while loop colors at least a constant fraction of the uncolored edges. Let ϕ_0 be the partial coloring at the beginning of a fixed interation of the while loop. Let ϕ_{γ} be the partial coloring just before the iteration γ of the for loop in Line 3. Similarly, let $\phi_{\gamma,\alpha}$ be the partial coloring just before the iteration γ and α of the nested for loops in Line 3 and Line 5. We again use $|\phi_0^{-1}(\perp)|$ to denote the number of uncolored edges at the beginning of current iteration of the while loop.

We use $\mathcal{F}_{\gamma,\alpha,\beta}$ to denote the normal and sub-reverse $\alpha\beta$ -fans considered in the nested iteration γ, α, β . In addition, let $\mathcal{F}_{\gamma,\alpha} = \bigcup_{\beta=1}^{\Delta+1} \mathcal{F}_{\gamma,\alpha,\beta}$. We again let $V_{\gamma,\alpha}(\phi)$ denote the set of color- γ incomplete vertices that miss color α with respect to the partial coloring ϕ .

Lemma 4.4. Each iteration of the while loop of Algorithm 2 colors a constant fraction of uncolored edges. Thus, it takes $O(\log n)$ iterations to color all the edges.

Proof. For each $uv \in \phi_0^{-1}(\bot)$, we orient the edge from u to v if the color of u is smaller than that of v; otherwise we orient it from v to u. Specifically, we orient the edge from u to v if $u \in V_{\gamma}$ and $v \in V_{\gamma'}$ such that $\gamma < \gamma'$. We will show that

$$\sum_{\alpha=1}^{\Delta+1} |M_{\gamma,\alpha}| + \sum_{\alpha=1}^{\Delta+1} |\mathcal{F}_{\gamma,\alpha}| \ge \sum_{u \in V_{\gamma}} \operatorname{outdeg}(u)/8 \ .$$

For each $v \in V_{\gamma}$, let C(v) denote the first outdeg(v) missing colors of v w.r.t. ϕ_{γ} so that |C(v)| = outdeg(v). For each $\alpha \in C(v)$, if $v \in V_{\gamma,\alpha}(\phi_{\gamma,\alpha})$, then there can be two cases. The first case is that an uncolored edge incident to v will be in the matching $M_{\gamma,\alpha}$. The second case is that v will belong to a fan in $\mathcal{F}_{\gamma,\alpha}$. Recall that in the second case, the algorithm will first grow a normal α -fan from v and if the fan intersects with an existing fan, then v will belong to a reverse α -fan after the merge step.

Otherwise, if $v \notin V_{\gamma,\alpha}(\phi_{\gamma,\alpha})$, it must be caused by either the repairing step of some normal or sub-reverse fan that was grown in iteration γ . Repairing a normal fan F can only change the missing colors (i.e. $M(\cdot)$) of its leaves, its center, and the other endpoint of the alternating path. So if M(v)is changed by repairing F, we conclude that v is the other end point of the alternating path of F or v is the center of F. Note that v cannot be among F's leaves since V_{γ} is an independent set. On the other hand, repairing a sub-reverse fan F can only changes the missing colors of at most five vertices (this corresponds to the case of repairing a sub-reverse fan with three leaves; in fact, a more careful argument shows that a sub-reverse fan can be blamed at most twice).

Therefore, if $v \notin V_{\gamma,\alpha}(\phi_{\gamma,\alpha})$, we can blame it to the repair of a normal or sub-reverse fan F grown in iteration γ . We have just argued that each fan F can be blamed at most five times. Furthermore, a normal or sub-reverse fan in $\mathcal{F}_{\gamma,\alpha}$ has at most one or three vertices in $V_{\gamma,\alpha}(\phi_{\gamma,\alpha})$ respectively. An edge in $M_{\gamma,\alpha}$ has one vertex in $V_{\gamma,\alpha}(\phi_{\gamma,\alpha})$. Hence, we have the following relationship

$$\sum_{u \in V_{\gamma}} \operatorname{outdeg}(u) = \sum_{u \in V_{\gamma}} |C(u)| \le \sum_{\alpha=1}^{\Delta+1} |M_{\gamma,\alpha}| + 3\sum_{\alpha=1}^{\Delta+1} |\mathcal{F}_{\gamma,\alpha}| + (\# \text{blames})$$
$$\le \sum_{\alpha=1}^{\Delta+1} |M_{\gamma,\alpha}| + 8\sum_{\alpha=1}^{\Delta+1} |\mathcal{F}_{\gamma,\alpha}| .$$

The second inequality is because (#blames) $\leq 5 \sum_{\alpha=1}^{\Delta+1} |\mathcal{F}_{\gamma,\alpha}|$. Therefore,

$$\sum_{\alpha=1}^{\Delta+1} |M_{\gamma,\alpha}| + \sum_{\alpha=1}^{\Delta+1} |\mathcal{F}_{\gamma,\alpha}| \ge \sum_{u \in V_{\gamma}} \operatorname{outdeg}(u)/8 .$$

Note that at most half of the (normal and sub-reverse) fans are destroyed, since repairing a fan can only destroy at most one other fan. If a fan is not destroyed, then one uncolored edge in the fan must become colored. Therefore, at least $\sum_{\gamma=1}^{\Delta+1} \sum_{\alpha=1}^{\Delta+1} |M_{\gamma,\alpha}| + \sum_{\gamma=1}^{\Delta+1} \sum_{\alpha=1}^{\Delta+1} |\mathcal{F}_{\gamma,\alpha}|/2$ uncolored edges are colored. The total number of edges that are colored during an iteration of the while loop is at least

$$\sum_{\gamma=1}^{\Delta+1} \sum_{\alpha=1}^{\Delta+1} |M_{\gamma,\alpha}| + \sum_{\gamma=1}^{\Delta+1} \sum_{\alpha=1}^{\Delta+1} |\mathcal{F}_{\gamma,\alpha}|/2 \ge \sum_{\gamma=1}^{\Delta+1} \sum_{v \in V_{\gamma}} \operatorname{outdeg}(v)/16 = |\phi_0^{-1}(\bot)|/16$$

This implies at least a fraction 1/16 of the uncolored edges become colored during the iteration. Thus, after $O(\log n)$ iterations, all edges are colored.

It is easy to check that the number of rounds is $O\left(\log n \cdot \Delta^2 \left(\Delta^4 + \Delta \cdot T \cdot \log n\right)\right)$ and the number of parallel augmentation in Line 18 and Line 19 is $t = O(\Delta^3 \log n)$.

If we opt for a randomized algorithm with $\ell(G) = 1$ or a deterministic algorithm with $\ell(G) = O(\log n/(\log \lambda + \log \log n))$, then we set $T = O(\Delta^7 t)$ or $T = 2t\lambda$ respectively as in the previous section. Since we managed to reduce t by a factor Δ^2 , we obtain the following improvements.

Theorem 4.5. There exists a deterministic distributed $\Delta + O(\log n/(\log \lambda + \log \log n))$ -edge-coloring algorithm that runs in $O(\Delta^6 \cdot \lambda \log^3 n)$ rounds. Furthermore, there exists a deterministic distributed algorithm that colors all the edges using $\Delta + 2$ colors and runs in $O(\Delta^{13} \cdot \log^3 n)$ rounds.

In general, we save a factor Δ^2 for both the deterministic and randomized algorithms. One may ask if we can make further improvement. We next show that it is possible to save another factor Δ^2 if the graph is bipartite.

4.3 Algorithm for Bipartite Graphs

Algorithm outline: We now give an even faster and simpler algorithm for edge coloring bipartite graphs. This algorithm is also based on a similar idea. The pseudo-code is outlined in Algorithm 4.

We first fix an iteration of the **while loop**. Then, fix an iteration α of the **for loop** in Line 2. We first find a maximal matching among the uncolored edges uv where $\alpha \in M(u) \cap M(v)$ in $O(\Delta + \log^* n)$ rounds [BEK14]. We color the matching with α . Then, let V_{α} the independent set of the remaining incomplete vertices that are still missing α .

For bipartite graphs, we however do not need to grow a fan from each vertex in V_{α} . We consider uncolored edges vu for each v in V_{α} . Some of these edges may share an end point that is not in V_{α} . We can merge these intersecting edges to form reverse α -fans. This can easily be implemented in a constant number of rounds. Now, we are left with a set of disjoint uncolored α -edges Y_{α} that have one end point in V_{α} and a set of reverse α -fans \mathcal{F}_{α} . All the reverse α -fans and α -edges are vertex-disjoint. They are also maximal, i.e., all vertices in V_{α} belong to a reverse α -fan or an α -edge.

In the normal setting, we can color an edge vu in Y_{α} as follows. Suppose $\alpha \in M(v)$ and $\beta \in M(u)$. We call vu an $\alpha\beta$ -edge. First, find the maximal alternating $\alpha\beta$ -path P_u from u. We observe that because the graph is bipartite, this path cannot end at u or v. Hence, we can perform $\operatorname{augment}(P_u)$ and color vu with α . Implementation in the distributed setting: Recall that we currently fix an iteration α in Line 2. We again split reverse fans into sub-reverse fans. We use $\alpha\beta$ -items to refer to $\alpha\beta$ -edges and sub-reverse $\alpha\beta$ -fans. We aim to repair $\alpha\beta$ -items, i.e., color an $\alpha\beta$ -edge or repair a sub-reverse $\alpha\beta$ -fan in iteration β of the for loop in Line 8. We want to color $\alpha\beta$ -items simultaneously and efficiently in the LOCAL model. We again truncate an alternating path if its length exceeds T and pick a blocking edge accordingly (based on whether the goal is a deterministic algorithm with $\ell(G) = \log n/(\log \lambda + \log \log n)$ or a randomized algorithm with $\ell(G) = 1$).

A truncated alternating path $P_v(T)$ where $|P_v| > T$ cannot intersect another $\alpha\beta$ -item. To this end, we build a conflict graph $G_{\alpha,\beta}$ in which the vertices are $\alpha\beta$ -items. There is an edge between two vertices if a truncated alternating path P where $|P| \leq T$ of an $\alpha\beta$ -item ends at another $\alpha\beta$ -item. With a similar argument as in the previous sections, this graph has bounded arboricity and therefore can be colored using O(1) colors in $O(T \log n)$ rounds [BE10]. We then iterate through each color class and repair $\alpha\beta$ -edges and sub-reverse $\alpha\beta$ -fans with this color. Since in each color class, the alternating paths neither intersect nor run into another item, we can repair them simultaneously. Similar to the previous algorithm, we also need to repair semi-destroyed sub-reverse $\alpha\beta$ -fans.

Alg	gorithm 4 Distributed Item-Repair for Bipartite Graphs
1:	while there exists uncolored edges do
2:	for $\alpha = \operatorname{color} 1, 2, \dots, \Delta$ do
3:	Let M_{α} be a maximal matching of $E_{\alpha} = \{vu \in E \mid vu \text{ is uncolored and } \alpha \in M(v) \cap M(u)\}.$
4:	Color M_{α} with α .
5:	Let V_{α} be the set of incomplete vertices that now miss the color α .
6:	For each $v \in V_{\alpha}$, pick an incident uncolored edge vu . Activate vu .
7:	Merge intersecting active edges into reverse α -fans.
8:	for $\beta = \operatorname{color} 1, 2, \dots, \Delta + 1$ do
9:	Build a conflict graph $G_{\alpha\beta}$ for $\alpha\beta$ -edges and sub-reverse $\alpha\beta$ -fans.
10:	Color $G_{\alpha\beta}$ using $O(1)$ colors.
11:	for $\eta = \operatorname{color} 1, 2, \dots O(1)$ do
12:	Repair undestroyed $\alpha\beta$ -edges and sub-reverse fans in $G_{\alpha\beta}$ with color η .
13:	Repair semi-destroyed sub-reverse fans in $G_{\alpha\beta}$.

Let $V_{\alpha}(\phi)$ be the set of incomplete vertices missing α with respect to the partial coloring ϕ . Let ϕ_0 be the partial coloring at the beginning of a fixed while loop and ϕ_{α} be the partial coloring right before the iteration α of the loop in Line 2.

Lemma 4.6. Each iteration of the while loop of Algorithm 4 colors a constant fraction of uncolored edges. Thus, it takes $O(\log n)$ iterations to color all the edges.

Proof. Let d(v) be the number of uncolored edges incident to v with respect to ϕ_0 . Let C(v) be the first d(v) missing color at v with respect to ϕ_0 . Let \mathcal{I}_{α} be the set that includes all $\alpha\beta$ -items and the edges in the matching M_{α} .

Consider $\alpha \in C(v)$. If $v \in V_{\alpha}(\phi_{\alpha})$, then v belongs to an element in \mathcal{I}_{α} . Specifically, v may belong to an edge in the matching M_{α} , an $\alpha\beta$ -edge or a sub-reverse $\alpha\beta$ -fan for some β .

If $v \notin V_{\alpha}(\phi_{\alpha})$, then we can blame this to an element in $\mathcal{I}_{\alpha'}$ for $\alpha' < \alpha$. In particular, this could be caused by either a) the repair of an α' -item or b) the coloring of some edge in $M_{\alpha'}$.

The coloring step of each edge in a matching $M_{\alpha'}$ can be blamed at most twice since it can only change the missing colors of its endpoints. The repair of an $\alpha\beta$ -edge can be blamed at most three times, since this can only change the missing colors of two of its endpoinds and the other endpoint of its alternating path. Finally, the repair of a sub-reverse $\alpha\beta$ -fan can be blamed at most five times. Thus, each element in $\bigcup_{\alpha=1}^{\Delta+1} \mathcal{I}_{\alpha}$ can be blamed at most five times overall.

We also note that an edge in M_{α} has two vertices in $V_{\alpha}(\phi_{\alpha})$. An α -edge has exactly one vertex in $V_{\alpha}(\phi_{\alpha})$ and a sub-reverse α -fan has at most three vertices in $V_{\alpha}(\phi_{\alpha})$. Therefore, we have the following relationship

$$\sum_{u \in V} d(u) = \sum_{u \in V} |C(u)| \le 3 \sum_{\alpha=1}^{\Delta} |\mathcal{I}_{\alpha}| + (\#\text{blames})$$
$$\le 8 \sum_{\alpha=1}^{\Delta} |\mathcal{I}_{\gamma,\alpha}| \qquad (\#\text{blames}) \le 5 \sum_{\alpha=1}^{\Delta+1} |\mathcal{I}_{\gamma,\alpha}| .$$

Using the same argument as in previous sections, at least half of α -items are repaired. Therefore, the number of edges being colored in this while iteration is at least

$$\sum_{\alpha=1}^{\Delta} |\mathcal{I}_{\alpha}|/2 \ge \sum_{v \in V} d(v)/16 = |\phi_0^{-1}(\bot)|/8 .$$

The number of parallel augmentations in Line 12 and Line 13 is $t = O(\log n \cdot \Delta^2)$. Similarly, by setting $T = O(\lambda t)$ or $T = O(t\Delta^7)$, we achieve the following respective improvements.

Theorem 4.7. Consider bipartite graphs. There exists a deterministic distributed $\Delta + O(\log n/(\log \lambda + \log \log n))$ -edge-coloring algorithm that runs in $O(\Delta^4 \cdot \lambda \log^3 n)$ rounds. Furthermore, there exists a randomized, $(\Delta + 1)$ -edge-coloring distributed algorithm that runs in $O(\Delta^{11} \log^3 n)$ rounds.

Note that for the randomized case, we use one fewer color in the randomized algorithm since the coloring algorithm requires m(v) to be non-empty only when v is incident to an uncolored edge. Hence, a palette of Δ colors, excluding the special color \star is sufficient.

References

Noga Alon, László Babai, and Alon Itai. A fast and simple randomized parallel algorithm [ABI86] for the maximal independent set problem. Journal of algorithms, 7(4):567–583, 1986. [Arj82] Eshrat Arjomandi. An efficient algorithm for colouring the edges of a graph with $(\Delta + 1)$ colours. INFOR: Information Systems and Operational Research, 20(2):82–101, 1982. [Bar15] Leonid Barenboim. Deterministic $(\Delta + 1)$ -coloring in sublinear (in Δ) time in static, dynamic and faulty networks. In Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, pages 345–354, 2015. [BE10] Leonid Barenboim and Michael Elkin. Sublogarithmic distributed mis algorithm for sparse graphs using nash-williams decomposition. Distributed Computing, 22(5-6):363–379, 2010. [BEG18] Leonid Barenboim, Michael Elkin, and Uri Goldenberg. Locally-iterative distributed ($\Delta +$ 1)-coloring below szegedy-vishwanathan barrier, and applications to self-stabilization and to restricted-bandwidth models. In the Proc. of ACM Symp. on Princ. of Dist. Comp. (*PODC*), pages 437–446, 2018. [BEK14] L. Barenboim, M. Elkin, and F. Kuhn. Distributed $(\Delta + 1)$ -coloring in linear (in Δ) time. SIAM Journal on Computing, 43(1):72-95, 2014. L. Barenboim, M. Elkin, S. Pettie, and J. Schneider. The locality of distributed symmetry [BEPS16] breaking. Journal of the ACM (JACM), 63(3), 2016. Article 20.

- [CHL⁺18] Y.-J. Chang, Q. He, W. Li, S. Pettie, and J. Uitto. The complexity of distributed edge coloring with small palettes. In *Proceedings 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2633–2652, 2018.
- [CLP18] Y.-J. Chang, W. Li, and S. Pettie. An optimal distributed $(\Delta + 1)$ -coloring algorithm? In Proceedings 50th ACM Symposium on Theory of Computing (STOC), pages 445–456, 2018.
- [CPS14] Kai-Min Chung, Seth Pettie, and Hsin-Hao Su. Distributed algorithms for the lovász local lemma and graph coloring. In the Proc. of ACM Symp. on Princ. of Dist. Comp. (PODC), pages 134–143. ACM, 2014.
- [DGP98] Devdatt Dubhashi, David A Grable, and Alessandro Panconesi. Near-optimal, distributed edge colouring via the nibble method. *Theoretical Computer Science*, 203(2):225–251, 1998.
- [DHZ19] Ran Duan, Haoqing He, and Tianyi Zhang. Dynamic edge coloring with improved approximation. In *SODA*, pages 1937–1945. SIAM, 2019.
- [EPS15] Michael Elkin, Seth Pettie, and Hsin-Hao Su. $(2\Delta 1)$ -edge-coloring is much easier than maximal matching in the distributed setting. In Symp. on Discrete Algorithms (SODA), pages 355–370. SIAM, 2015.
- [FGK17] Manuela Fischer, Mohsen Ghaffari, and Fabian Kuhn. Deterministic distributed edgecoloring via hypergraph maximal matching. In Proc. of the Symp. on Found. of Comp. Sci. (FOCS), pages 180–191, 2017.
- [FHK16] Pierre Fraigniaud, Marc Heinrich, and Adrian Kosowski. Local conflict coloring. In Proc. of the Symp. on Found. of Comp. Sci. (FOCS), pages to appear, arXiv:1511.01287, 2016.
- [GDP08] Shashidhar Gandham, Milind Dawande, and Ravi Prakash. Link scheduling in wireless sensor networks: Distributed edge-coloring revisited. *Journal of Parallel and Distributed Computing*, 68(8):1122 – 1134, 2008.
- [GHK⁺17] Mohsen Ghaffari, Juho Hirvonen, Fabian Kuhn, Yannic Maus, Jukka Suomela, and Jara Uitto. Improved distributed degree splitting and edge coloring. In Proc. of the Int'l Symp. on Dist. Comp. (DISC), pages 1–15, 2017.
- [GHK18] Mohsen Ghaffari, David Harris, and Fabian Kuhn. On derandomizing local distributed algorithms. In *Proc. of the Symp. on Found. of Comp. Sci. (FOCS)*, pages 662–673, 2018.
- [GKMU18] Mohsen Ghaffari, Fabian Kuhn, Yannic Maus, and Jara Uitto. Deterministic distributed edge-coloring with fewer colors. In Proc. 50th ACM Symposium on Theory of Computing, STOC 2018, pages 418–430, 2018.
- [GNK⁺85] H. N. Gabow, T. Nishizeki, O. Kariv, D. Leven, and O. Terada. Algorithms for edgecoloring graphs. *Technical Report TRECIS-8501, Tohoku University*, 1985.
- [GP87] A. V. Goldberg and S. A. Plotkin. Parallel $(\Delta + 1)$ -coloring of constant-degree graphs. Information Processing Letters, 25(4):241 – 245, 1987.
- [GPS88] A. V. Goldberg, S. A. Plotkin, and G. E. Shannon. Parallel symmetry-breaking in sparse graphs. SIAM J. Discrete Math., 1(4):434–446, 1988.

- [GS17] Mohsen Ghaffari and Hsin-Hao Su. Distributed degree splitting, edge coloring, and orientations. In *Proc. 28th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2505–2523, 2017.
- [HSS18] David Harris, Johannes Schneider, and Hsin-Hao Su. Distributed ($\Delta + 1$)-coloring in sublogarithmic rounds. *Journal of the ACM*, 65(4), 2018. Article 19.
- [Joh99] Öjvind Johansson. Simple distributed $\Delta + 1$ -coloring of graphs. Inf. Process. Lett., 70, 1999.
- [KW06] Fabian Kuhn and Roger Wattenhofer. On the complexity of distributed graph coloring. In Symp. on Principles of Distributed Computing (PODC), 2006.
- [Lin92] Nati Linial. Locality in Distributed Graph Algorithms. SIAM Journal on Computing, 21(1):193–201, 1992.
- [Lub86] Michael Luby. A simple parallel algorithm for the maximal independent set problem. SIAM journal on computing, 15(4):1036–1053, 1986.
- [MG92] J. Misra and David Gries. A constructive proof of vizing's theorem. *Inf. Process. Lett.*, 41(3):131–133, March 1992.
- [PR01] Alessandro Panconesi and Romeo Rizzi. Some simple distributed algorithms for sparse networks. *Distributed computing*, 14(2):97–100, 2001.
- [PS92] Alessandro Panconesi and Aravind Srinivasan. Improved distributed algorithms for coloring and network decomposition problems. In Proc. of the Symp. on Theory of Comp. (STOC), pages 581–592. ACM, 1992.
- [PS97] Alessandro Panconesi and Aravind Srinivasan. Randomized distributed edge coloring via an extension of the chernoff-hoeffding bounds. *SIAM Journal on Computing*, 26(2):350– 368, 1997.
- [PS15] S. Pettie and H.-H. Su. Distributed algorithms for coloring triangle-free graphs. Information and Computation, 243:263–280, 2015.
- [Ram99] S. Ramanathan. A unified framework and algorithm for channel assignment in wireless networks. *Wireless Networks*, 5(2):81–94, 1999.
- [Viz64] V. G. Vizing. On an estimate of the chromatic class of a p-graph. *Diskret. Analiz*, (3):25–30, 1964.