



Contents lists available at ScienceDirect

Journal of Combinatorial Theory,
Series B

www.elsevier.com/locate/jctb



A fast distributed algorithm for $(\Delta + 1)$ -edge-coloring[☆]



Anton Bernshteyn^{a,b}

^a Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA, USA

^b School of Mathematics, Georgia Institute of Technology, Atlanta, GA, USA

ARTICLE INFO

Article history:

Received 13 July 2020

Available online 28 October 2021

Keywords:

Edge-coloring

Vizing's theorem

Distributed algorithms

LOCAL model

ABSTRACT

We present a deterministic distributed algorithm in the LOCAL model that finds a proper $(\Delta + 1)$ -edge-coloring of an n -vertex graph of maximum degree Δ in $\text{poly}(\Delta, \log n)$ rounds. This is the first nontrivial distributed edge-coloring algorithm that uses only $\Delta + 1$ colors (matching the bound given by Vizing's theorem). Our approach is inspired by the recent proof of the measurable version of Vizing's theorem due to Grebík and Pikhurko.

© 2021 Elsevier Inc. All rights reserved.

1. Introduction

1.1. Distributed algorithms and edge-coloring

The LOCAL model of distributed computation was introduced by Linial in the seminal paper [14]. In this model an n -vertex (simple undirected) graph G abstracts a communication network where each vertex plays the role of a processor and edges represent communication links. Every vertex is given a unique $\Theta(\log n)$ -bit identifier (used for

[☆] This research was partially supported by the NSF grant DMS-2045412.

E-mail address: bahtoh@gatech.edu.

symmetry-breaking). Initially, each vertex knows its own identifier, as well as n (the number of vertices) and perhaps some other global parameters, such as the maximum degree Δ of G . The computation proceeds in *rounds*. During each round, the vertices first perform arbitrary local computations and then synchronously broadcast messages to all their neighbors. At the end, each vertex should output its part of the global solution (for instance, its own color or, in the context of edge-coloring, the colors of the edges incident to it). The only measure of efficiency for such an algorithm is the worst-case number of communication rounds. The reader is referred to the book [2] by Barenboim and Elkin for an introduction to this subject.

It is clear that in a **LOCAL** algorithm that terminates in T rounds, each vertex only has access to information in its radius- T neighborhood. Furthermore, every T -round **LOCAL** algorithm can be transformed into one in which every vertex first collects all the information about its radius- T neighborhood and then makes a decision, based on this information alone, about its part of the output (see [2, §4.1.2]). In particular, any function of G can trivially be computed by a **LOCAL** algorithm in $O(\text{diameter}(G))$ rounds. While in many cases much more efficient distributed algorithms are known, there are still some classical problems for which the trivial bound $O(\text{diameter}(G))$ is the state of the art. In this paper we study one such problem: edge-coloring using $\Delta + 1$ colors.

Recall that a **proper k -edge-coloring** of a graph G is a map $\varphi: E(G) \rightarrow [k]$ such that $\varphi(e) \neq \varphi(h)$ for every pair of distinct edges $e, h \in E(G)$ that share an endpoint. (Here and in what follows we use the standard notation $[k] := \{1, \dots, k\}$ for $k \in \mathbb{N}^+$.) A celebrated theorem of Vizing [20] (see [18, Appendix A.1] for an English translation of Vizing’s paper) asserts that every graph G of maximum degree Δ has a proper $(\Delta + 1)$ -edge-coloring. However, all heretofore known proofs of Vizing’s theorem have been inherently “non-local” (we discuss this in more detail in §1.2). In particular, prior to this work, no nontrivial distributed $(\Delta + 1)$ -edge-coloring algorithm has been known. Here we provide such an algorithm; specifically, we show that a proper $(\Delta + 1)$ -edge-coloring of an n -vertex graph G can be found by a **LOCAL** algorithm whose number of rounds is polynomial in $\log n$ and Δ .

Theorem 1.1. *There exists a deterministic distributed algorithm in the **LOCAL** model that computes a proper $(\Delta + 1)$ -edge-coloring of an n -vertex graph of maximum degree Δ in $\text{poly}(\Delta, \log n)$ rounds.*

While Theorem 1.1 provides the first efficient distributed algorithm for $(\Delta + 1)$ -edge-coloring, significant progress has been made previously concerning edge-coloring with more colors. In particular, considerable work has been done on the problem of $(2\Delta - 1)$ -edge-coloring, since $2\Delta - 1$ is the number of colors required for a greedy (sequential) edge-coloring algorithm. This is a special case of the more general $(\Delta + 1)$ -vertex-coloring problem, which has also attracted a lot of attention. Among the earliest results in distributed computing is a randomized $O(\log n)$ -round $(\Delta + 1)$ -vertex-coloring algorithm due to Alon, Babai, and Itai [1] and independently Luby [15]. On the other hand, Goldbreg,

Plotkin, and Shannon [10] designed a deterministic $(\Delta + 1)$ -vertex-coloring algorithm that runs in $O(\Delta^2) + \log^* n$ rounds (here $\log^* n$ is the iterated logarithm of n , i.e., the number of times the logarithm function must be applied to n before the result becomes at most 1). The Goldbreg–Plotkin–Shannon algorithm is extremely fast when Δ is small compared to n (e.g., when Δ is constant), but for large Δ it becomes slower than the Alon–Babai–Itai/Luby algorithm. Developing a *deterministic* $(\Delta + 1)$ -vertex-coloring or $(2\Delta - 1)$ -edge-coloring algorithm that runs in $\text{poly}(\log n)$ rounds regardless of the value of Δ has been a major challenge [2, Problems 11.2 and 11.4]. After a long line of contributions by numerous researchers, such algorithms have been discovered by Fischer, Ghaffari, and Kuhn [7] for edge-coloring and by Rozhoň and Ghaffari [17] for vertex-coloring.

There have also been many results concerning edge-coloring with fewer than $2\Delta - 1$ colors. For thorough surveys see, e.g., [5, 9]. Here we only mention a few highlights. Chang, He, Li, Pettie, and Uitto [5] developed a randomized algorithm for $(\Delta + \tilde{O}(\sqrt{\Delta}))$ -edge-coloring that runs in $\text{poly}(\Delta, \log \log n)$ rounds. (Their algorithm relies on efficiently solving instances of the Lovász Local Lemma, which can be done in $\text{poly}(\Delta, \log \log n)$ due to a recent result of Rozhoň and Ghaffari [17].) In the deterministic setting, Ghaffari, Kuhn, Maus, and Uitto [9] gave a $\text{poly}(\Delta, \log n)$ -round algorithm for $\lfloor 3\Delta/2 \rfloor$ -edge-coloring. The smallest number of colors for which an efficient distributed algorithm was known prior to our work is $\Delta + 2$ (i.e., just one more than the bound in Vizing’s theorem). Namely, a randomized $\text{poly}(\Delta, \log n)$ -round algorithm for $(\Delta + 2)$ -edge-coloring was designed by Su and Vu [19]. In a recent breakthrough, Rozhoň and Ghaffari [17] developed a general derandomization technique that, in particular, allows one to modify the Su–Vu algorithm to make it deterministic.

It is natural to wonder how close to optimal the number of rounds required by our algorithm is. Chang, He, Li, Pettie, and Uitto [5] showed that any $(\Delta + 1)$ -edge-coloring algorithm based on “extending partial colorings by recoloring subgraphs” (a class to which our algorithm belongs) must take at least $\Omega(\Delta \log(n/\Delta))$ rounds. Furthermore, they showed that every deterministic LOCAL algorithm for $(2\Delta - 2)$ -edge-coloring requires $\Omega(\log n / \log \Delta)$ rounds, even in the case when the underlying graph G is a tree. The number of rounds in our algorithm is certainly not linear in $\Delta \log n$, although we did not make an attempt to optimize it precisely. The following remains an interesting open problem:

Question 1.2. In the regime when Δ is constant, does there exist a randomized $(\Delta + 1)$ -edge-coloring algorithm that takes only $o(\log n)$ rounds?

1.2. Locality of $(\Delta + 1)$ -edge-coloring

The main ingredient of our algorithm is a certain purely combinatorial result concerning extensions of partial $(\Delta + 1)$ -edge-colorings, namely Theorem 1.3 below. Before stating it, we need to introduce some terminology and explain our motivation. Let G be an n -vertex graph of maximum degree Δ . Fix a partition $E(G) = X \sqcup U$ of the edge

set of G and let $\varphi: X \rightarrow [\Delta + 1]$ be a proper partial edge-coloring with domain X . We call the edges in X (resp. U) **colored** (resp. **uncolored**) by φ . A subgraph $H \subseteq G$ is **augmenting** (for φ) if at least one edge of H is uncolored and there is a proper coloring $\psi: X \cup E(H) \rightarrow [\Delta + 1]$ that agrees with φ on $X \setminus E(H)$; that is, by only modifying φ on the edges of H , it is possible to extend it to a proper partial coloring with domain $X \cup E(H)$. Notice that if $H_1, \dots, H_k$ are vertex-disjoint augmenting subgraphs, then their (disjoint) union $H_1 \sqcup \dots \sqcup H_k$ is also augmenting.

A standard way to prove Vizing’s theorem is to construct, given an uncolored edge $e \in U$, an augmenting graph H with $e \in E(H)$ of a certain special form, called a *Vizing chain* (see Fig. 1a). A Vizing chain H consists of a *fan*—i.e., a set of edges that share the same common vertex—and a *path* that is *alternating* in the sense that the sequence of colors assigned by φ to its edges has the form $\alpha, \beta, \alpha, \beta, \dots$ for a pair of colors α, β . In the distributed setting, this approach raises two difficulties. First, uncolored edges cannot be treated one at a time; it is necessary to be able to extend the coloring to include a non-negligible fraction of the uncolored edges simultaneously. The second difficulty is even more fundamental: the Vizing chain H corresponding to an uncolored edge e may have large diameter (in principle, linear in n), so it cannot be discovered in a small number of rounds in the LOCAL model. We overcome both these difficulties by showing that it is possible to find a “large” collection of vertex-disjoint “small” augmenting subgraphs (although these subgraphs need not be Vizing chains anymore):

Theorem 1.3. *There is $n_0 \in \mathbb{N}$ such that for all $n > n_0$, the following holds. Let G be an n -vertex graph of maximum degree Δ . Fix a partition $E(G) = X \sqcup U$ and let $\varphi: X \rightarrow [\Delta + 1]$ be a proper partial edge-coloring. Then there exists a subset $W \subseteq U$ of size $|W| \geq |U|/((\Delta + 1)^{10}(\log n)^2)$ such that it is possible to assign to each edge $e \in W$ a connected augmenting subgraph $H^e \subseteq G$ with the following properties:*

- for each $e \in W$, $E(H^e) \cap U = \{e\}$, that is, e is the unique uncolored edge in H^e ;
- the vertex sets of the graphs H^e , $e \in W$, are pairwise disjoint;
- for each $e \in W$, $|E(H^e)| \leq (\Delta + 1)^6(\log n)^2$.

Thanks to the recent work of Ghaffari, Harris, and Kuhn [8] and Harris [12] on distributed approximation algorithms for hypergraph maximum matching, it is straightforward to derive Theorem 1.1 from Theorem 1.3. We present the details of this derivation in §2. Note that if $U \neq \emptyset$, then the lower bound on the size of W in Theorem 1.3 is positive. Thus, Theorem 1.3 in particular implies that if G contains only a single uncolored edge e , then there is a connected augmenting subgraph $H \subseteq G$ such that $e \in E(H)$ and $|E(H)| \leq \text{poly}(\Delta, \log n)$.

Let us now say a few words about the proof of Theorem 1.3. Our inspiration comes from recent developments in the area of *descriptive combinatorics*, i.e., the study of combinatorial problems on infinite graphs in the presence of additional topological or measure-theoretic constraints. For a state-of-the-art introduction to this subject, see the




Fig. 1. Types of augmenting subgraphs.

surveys [13] by Kechris and Marks and [16] by Pikhurko. It turns out that many of the challenges one encounters in descriptive combinatorics mirror those in distributed computing; furthermore, there are some explicit implications between the two areas [3,4]. Until very recently, one of the most celebrated open problems in descriptive combinatorics had been to obtain a “measurable” version of Vizing’s theorem for graphs G on standard probability spaces (for instance, with $V(G) = [0, 1]$, the unit interval). This was recently accomplished (under some minor additional assumptions) by Grebík and Pikhurko [11]. Their key idea was to consider a more general type of augmenting graphs, namely *two-step Vizing chains* (called “iterated Vizing chains” in [11]; see Fig. 1b). Roughly speaking, to construct a two-step Vizing chain, one first starts growing a Vizing chain as usual, but then at some point, instead of completing the alternating path (which may be too long), one “changes the direction” and starts growing a second Vizing chain. The flexibility in the choice of the point where the second Vizing chain starts allowed Grebík and Pikhurko to control the total number of edges in the resulting structure. In particular, returning to finite graphs, a relatively straightforward adaptation of the Grebík–Pikhurko argument shows that for each uncolored edge $e \in U$, one can find an augmenting two-step Vizing chain H with $e \in E(H)$ such that $|E(H)| = O(\text{poly}(\Delta)\sqrt{n})$.

Unfortunately, the bound $O(\text{poly}(\Delta)\sqrt{n})$ is not sufficient for our purposes, as we can only work with augmenting graphs of diameter at most $\text{poly}(\Delta, \log n)$. We solve this issue by extending the ideas of Grebík and Pikhurko and considering *multi-step Vizing chains* (see Fig. 1c). Roughly speaking, we prove Theorem 1.3 by showing that it is possible to construct augmenting subgraphs of this form by joining together $O(\log n)$ partial Vizing chains, each consisting of $O(\text{poly}(\Delta) \log n)$ edges. For the reader’s convenience, we give an informal overview of our argument in §3.

Acknowledgments

I am very grateful to Clinton Conley and Oleg Pikhurko for insightful conversations and to the anonymous referees for carefully reading the manuscript and providing helpful suggestions.

2. From Theorem 1.3 to Theorem 1.1

In this section we deduce Theorem 1.1 from Theorem 1.3. To this end, we shall use the distributed approximation algorithm for hypergraph maximum matching due to Harris [12]. (Alternatively, we could have used the earlier and somewhat less efficient algorithm due to Ghaffari, Harris, and Kuhn [8, Theorem 1.2].) To begin with, we need a few definitions and some notation. Recall that $\log^* n$ denotes the iterated logarithm of n . The asymptotic notation $\tilde{O}(x)$ hides polylogarithmic factors, i.e., it stands for $O(\text{poly}(\log x) \cdot x)$. Let $\mathcal{H} = (V(\mathcal{H}), E(\mathcal{H}))$ be a hypergraph. To avoid potential confusion with the graph case, we will call the elements of $E(\mathcal{H})$ **hyperedges**. The **rank** of $\mathcal{H}$, in symbols $r(\mathcal{H})$, is the largest size of a hyperedge of $\mathcal{H}$ (if $E(\mathcal{H}) = \emptyset$, then we set $r(\mathcal{H}) := 0$). The **maximum degree** $d(\mathcal{H})$ of $\mathcal{H}$ is the maximum, over all the vertices $x \in V(\mathcal{H})$, of the number of hyperedges $S \in E(\mathcal{H})$ with $x \in S$. A **matching** in $\mathcal{H}$ is a set $M \subseteq E(\mathcal{H})$ of pairwise disjoint hyperedges. We use $\mu(\mathcal{H})$ to denote the maximum number of edges in a matching in $\mathcal{H}$. While the usual LOCAL model is defined for graphs, there is an analogous model operating on a hypergraph $\mathcal{H}$. Namely, in a single communication round of the LOCAL model on $\mathcal{H}$, each vertex $x \in V(\mathcal{H})$ is allowed to send messages to every vertex $y \in V(\mathcal{H})$ such that x and y are contained in a common hyperedge.

Theorem 2.1 (Harris [12, Theorem 1.1]). *There exists a deterministic distributed algorithm in the LOCAL model on an n -vertex hypergraph $\mathcal{H}$ that outputs a matching $M \subseteq E(\mathcal{H})$ with $|M| = \Omega(\mu/r)$ in*

$$\tilde{O}(r \log d + \log^2 d + \log^* n)$$

rounds, where $r := r(\mathcal{H})$, $d := d(\mathcal{H})$, and $\mu := \mu(\mathcal{H})$.

With Theorem 2.1 in hand, we are ready to derive Theorem 1.1 from Theorem 1.3. Let G be an n -vertex graph of maximum degree Δ . Since we only care about asymptotic bounds, we may assume that n is sufficiently large for Theorem 1.3 to apply. For brevity, set $V := V(G)$ and $E := E(G)$. We shall devise a $\text{poly}(\Delta, \log n)$ -round LOCAL algorithm that, given a proper partial edge-coloring $\varphi: X \rightarrow [\Delta + 1]$, $X \subseteq E$, computes a new proper partial edge-coloring $\psi: Y \rightarrow [\Delta + 1]$, $Y \subseteq E$, such that $Y \supseteq X$ and

$$|E \setminus Y| \leq (1 - 1/\text{poly}(\Delta, \log n))|E \setminus X|.$$

Iterating this algorithm $\text{poly}(\Delta, \log n)$ times starting with the empty coloring will then yield a proper $(\Delta + 1)$ -edge-coloring of the entire graph G , as desired.

So, let us fix a proper partial edge-coloring $\varphi: X \rightarrow [\Delta + 1]$. Define an auxiliary hypergraph $\mathcal{H}$ on the vertex set V by making a subset $S \subseteq V$ a hyperedge of $\mathcal{H}$ if and only if $S = V(H)$ for some connected augmenting subgraph $H \subseteq G$ with $|E(H)| \leq (\Delta + 1)^6 (\log n)^2$. Theorem 1.3 then implies that

$$\mu(\mathcal{H}) \geq |E \setminus X| / ((\Delta + 1)^{10} (\log n)^2) = |E \setminus X| / \text{poly}(\Delta, \log n).$$

By definition, each $S \in E(\mathcal{H})$ satisfies $|S| \leq (\Delta + 1)^6 (\log n)^2 + 1 = \text{poly}(\Delta, \log n)$, so

$$r(\mathcal{H}) \leq \text{poly}(\Delta, \log n).$$

To bound the maximum degree of $\mathcal{H}$, consider any vertex $x \in V$. If $x \in S \in E(\mathcal{H})$, then the vertices of S can be ordered as $x = x_0, x_1, \dots, x_{|S|-1}$ so that each x_i , $i \geq 1$, is adjacent to at least one of $x_0, \dots, x_{i-1}$. This means that once $x_0, \dots, x_{i-1}$ are fixed, there are at most $i\Delta \leq r(\mathcal{H})\Delta$ choices for x_i , and thus

$$d(\mathcal{H}) \leq (r(\mathcal{H})\Delta)^{r(\mathcal{H})} \leq \exp(\text{poly}(\Delta, \log n)).$$

Therefore, Theorem 2.1 provides a LOCAL algorithm on $\mathcal{H}$ that outputs a matching $M \subseteq E(\mathcal{H})$ with

$$|M| \geq |E \setminus X| / \text{poly}(\Delta, \log n)$$

in $\text{poly}(\Delta, \log n)$ rounds. It is clear that a single round of communication in the LOCAL model on $\mathcal{H}$ can be simulated in $\text{poly}(\Delta, \log n)$ rounds of the LOCAL model on G , so M can also be computed in the LOCAL model on G in $\text{poly}(\Delta, \log n)$ rounds. Once such a matching M is fixed, in $\text{poly}(\Delta, \log n)$ rounds it is possible to choose a “leader” in each hyperedge $S \in M$ (by picking, say, the vertex with the lexicographically least identifier). In another $\text{poly}(\Delta, \log n)$ rounds, each “leader” x surveys its corresponding hyperedge S , picks an arbitrary augmenting subgraph H with $V(H) = S$, and modifies the coloring so that every edge of H becomes colored. Since M is a matching in $\mathcal{H}$, all these modifications can happen in parallel without conflicting with each other. In the resulting coloring, every edge in X is still colored, plus there is at least one new colored edge per each $S \in M$, and hence we are done.

3. Informal overview of the proof of Theorem 1.3

In this section we sketch the main ideas behind our proof of Theorem 1.3. Due to the informal character of this section, we will avoid any technicalities and precise definitions, all of which are postponed to the later sections. It should be understood that the terminology we use here might have a slightly different meaning when we formally define it, but only in minor and technical ways.

Let G be an n -vertex graph of maximum degree Δ and let φ be a proper partial $(\Delta + 1)$ -edge-coloring of G . For simplicity, in this section we shall treat Δ as a constant, so that implicit constants in asymptotic notation may depend on Δ . Fix an uncolored edge $e = xy$. We will describe our strategy for proving the following consequence of Theorem 1.3:




Fig. 2. A Vizing chain $H = F + P$ before and after shifting. Greek letters represent colors.

There is a connected augmenting subgraph $H \subseteq G$ such that $e \in E(H)$ and $|E(H)| \leq O((\log n)^2)$.

As mentioned in §1.2, standard proofs of Vizing's theorem (for instance, the one given in [6, §5.3]) proceed by building an augmenting subgraph H consisting of two parts:

- a *fan* F with pivot x , i.e., a sequence of distinct edges incident to x , starting with e ;
- an $\alpha\beta$ -*alternating path* P , i.e., a path whose edges are colored $\alpha, \beta, \alpha, \beta$, etc.

There is some variation between different presentations of the proof of Vizing's theorem in the literature as to the precise way in which the fan F and the path P fit together. For our purposes, it will be most convenient to assume that P starts at the vertex z , where xz is the last edge of F . This situation is illustrated in Fig. 2a. We write $H = F + P$ to indicate that the graph H is built from F and P in this manner and refer to H as a *Vizing chain*. Given a Vizing chain $H = F + P$, we can *shift* the colors in H , as shown in Fig. 2b. One can show (see, e.g., [6, §5.3] or Corollary 4.10 below) that there is such a Vizing chain H that is *shiftable*, meaning that the coloring resulting after the shifting is proper, and thus H is an augmenting subgraph.

The problem, of course, is that there is no *a priori* upper bound on the length of the alternating path P (except $O(n)$). In [19], Su and Vu tackled this problem as follows. Suppose that P is too long, say, much longer than $\log n$. Pick a random edge $x'y'$ on P within distance $O(\log n)$ from z and truncate the path there. Then let P_0 be the truncated path and, instead of shifting the whole chain $F + P$, only shift $F + P_0$. This process is shown in Fig. 3. In the resulting coloring, the edge $x'y'$ becomes uncolored. If we could use one extra color, say $\Delta + 2$, then we would assign it to $x'y'$. By executing this procedure for all the uncolored edges, one eventually obtains a $(\Delta + 2)$ -edge-coloring




Fig. 3. Shifting a truncated Vizing chain $F + P_0$.




Fig. 4. A two-step Vizing chain $F + P_0 + F' + P'$.

of the entire graph that is proper except that there might be adjacent edges of color $\Delta + 2$. However, Su and Vu showed in [19] that with high probability, no two edges of color $\Delta + 2$ will be adjacent, and thus one actually gets a proper $(\Delta + 2)$ -edge-coloring.

In [11], Grebík and Pikhurko encountered a similar problem, but in the context of $(\Delta + 1)$ -edge-coloring (as we already mentioned in §1.2, Grebík and Pikhurko's goal was to prove a version of Vizing's theorem for infinite graphs on probability spaces). Their solution also involved picking a random edge $x'y'$ on P , truncating the path there, and then shifting the chain $F + P_0$. But since they could not use an extra color for the edge $x'y'$, they instead built a *new* Vizing chain $H' = F' + P'$ starting at the edge $x'y'$, and shifted the whole combination of $F + P_0 + F' + P'$ (see Fig. 4). Let us call the sequence $F + P_0 + F' + P'$ a *two-step Vizing chain*. This two-step approach can be used to build an

augmenting subgraph with only $O(\sqrt{n})$ edges, and while our ultimate goal is to reduce $O(\sqrt{n})$ further to $O((\log n)^2)$, it will perhaps be instructive if we explain first how to get the bound $O(\sqrt{n})$.

If the length of P is $O(\sqrt{n})$, then we are done, so assume that P is much longer than $\sqrt{n}$. The vertex x' is chosen at random from among the first $\Theta(\sqrt{n})$ vertices of P , and we want to show that the expected length of P' , relative to this random choice of x' , is $O(\sqrt{n})$. To this end, we fix an arbitrary vertex $u \in V(G)$ and ask, what is the probability that P' passes through u ? Suppose that the path P' starts at a vertex z' that is adjacent to x' and is $\gamma\delta$ -alternating (as is shown in Fig. 4). If we start at u and follow the edges colored γ and δ , then we would stop precisely when we hit z' . Thus, given u , we can uniquely determine x' by specifying the colors γ and δ ($(\Delta + 1)^2$ options) and the location of x' among the neighbors of z' (Δ options). In other words, given u , there are at most $O(\Delta^3) = O(1)$ possibilities for x' . Since x' is chosen randomly from a set of $\Theta(\sqrt{n})$ candidates, we conclude that

$$\mathbb{P}[P' \text{ passes through } u] \leq O(1/\sqrt{n}).$$

Summing over all $u \in V(G)$ yields $\mathbb{E}[|P'|] \leq O(n \cdot 1/\sqrt{n}) = O(\sqrt{n})$, as desired.

The argument in the previous paragraph has a flaw that, even in this informal overview, we cannot overlook. Namely, the path P' may intersect the fan F or the path P_0 . Because of that, some of the edges of P' may be colored differently before and after the chain $F + P_0$ is shifted. This means that, in principle, even if we know u , δ , and γ , we cannot locate z' by tracing the $\gamma\delta$ -alternating path starting at u unless we also know what the coloring looks like after shifting $F + P_0$. To circumvent this issue, Grebík and Pikhurko showed that, roughly speaking, a careful choice of P' ensures that P' is disjoint from $F + P_0$ with high probability. It is unlikely that P' intersects F —essentially because F is small, namely $|F| \leq \Delta = O(1)$. The real challenge is to make P' (edge-)disjoint from P_0 . Since every edge in P_0 is colored α or β , it would be ideal if we could simply pick P' so that $\{\gamma, \delta\} \cap \{\alpha, \beta\} = \emptyset$. Unfortunately, this is not always possible. Thankfully, Grebík and Pikhurko showed that P' can always be chosen so that either $\{\gamma, \delta\} \cap \{\alpha, \beta\} = \emptyset$ or else, $\{\gamma, \delta\} = \{\alpha, \beta\}$ (see Lemma 4.9 below). In the former case, P' and P_0 are disjoint, so we are done. In the latter case, P' and P_0 cannot intersect transversally: if $P' \cap P_0 \neq \emptyset$, then $P' \supseteq P_0$, and one can show that this is unlikely.

In this paper we take a natural next step and consider *multi-step Vizing chains* (see Fig. 5). Fix parameters $\ell = a \log n$ and $T = b \log n$, where a and b are positive constants with $a \gg b$. Let $x_0 := x$ and $y_0 := y$ and build a Vizing chain $F_0 + Q_0$ for $x_0 y_0$. Here Q_0 is an $\alpha_0 \beta_0$ -alternating path that starts at a vertex z_0 adjacent to x_0 . If $|Q_0| \leq \ell$, then stop. Otherwise, pick a random edge $x_1 y_1$ on Q_0 at distance at most ℓ from z_0 , truncate Q_0 at that edge, and let P_0 be the truncated path. After shifting $F_0 + P_0$, the edge $x_1 y_1$ becomes uncolored, and we repeat the process with $x_1 y_1$ in place of $x_0 y_0$, obtaining a fan F_1 and an $\alpha_1 \beta_1$ -alternating path Q_1 starting at a vertex z_1 adjacent to x_1 . Again,




Fig. 5. A multi-step Vizing chain.

if $|Q_1| \leq \ell$, then we stop, and otherwise we pick a random edge x_2y_2 on Q_1 , truncate Q_1 at that edge, and let P_1 be the truncated path. Then we shift $F_1 + P_1$, and repeat the process with x_2y_2 in place of x_1y_1 , etc. If this process terminates within T steps, then it produces an augmenting subgraph with at most $O(T\ell) = O((\log n)^2)$ edges, as desired. Hence, our goal becomes to show that the process terminates within T steps with positive probability.

The basic calculation here is similar to the one in the two-step case. Namely, we fix a vertex $u \in V(G)$ and bound the probability that the process runs for at least T steps and the path P_T passes through u . If we start at u and follow the edges colored α_T and β_T , we will stop at z_T . Since x_T is adjacent to z_T , we can determine x_T by specifying the colors α_T and β_T and the location of x_T among the neighbors of z_T . Then we can trace the $\alpha_{T-1}\beta_{T-1}$ -alternating path starting at x_T to find z_{T-1} , and then locate x_{T-1} among the neighbors of z_{T-1} . Continuing in this manner, we see that the total number of sequences $(x_1, \dots, x_T)$ for which P_T passes through u is at most $\Delta^{O(T)}$. Since each x_i is chosen among ℓ candidates, we conclude that, for some positive constant C ,

$$\mathbb{P}[P_T \text{ passes through } u] \leq \frac{\Delta^{O(T)}}{\ell^T} \leq \exp(-C \log n \log \log n).$$

If the process runs for at least T steps, then P_T has to pass through *some* vertex, so

$$\mathbb{P}[\text{the process runs for at least } T \text{ steps}] \leq n \cdot \exp(-C \log n \log \log n) = o(1).$$

Again, the above reasoning only really works if each P_i is disjoint from $F_0 + P_0 + \dots + F_{i-1} + P_{i-1}$. As in the two-step case, making P_i disjoint from the fans $F_0, \dots, F_{i-1}$ is a bit easier, so let us focus on the problem of making P_i disjoint from $P_0, \dots, P_{i-1}$. For $i > j$, let $\mathcal{E}_{ij}$ be the event that the first intersection that occurs during the process is

between P_i and P_j . We need to bound the probability of $\mathcal{E}_{ij}$. By using the same methods as Grebík and Pikhurko, i.e., by making sure that either $\{\alpha_i, \beta_i\} \cap \{\alpha_{i-1}, \beta_{i-1}\} = \emptyset$ or $\{\alpha_i, \beta_i\} = \{\alpha_{i-1}, \beta_{i-1}\}$, we can eliminate the possibility that $j = i - 1$, so let us assume that $j \leq i - 2$. Fix the outcome of the procedure up to step j , so that we have already settled on the vertices $x_0, x_1, \dots, x_j$ but not on x_{j+1} . At this point, we already know the path Q_j , but we have not yet truncated it. Pick any vertex u on Q_j within distance ℓ from z_j and let $\mathcal{E}_{ij,u}$ be the event that $\mathcal{E}_{ij}$ happens and the first intersection point between P_i and P_j is u . The key observation is that if $\mathcal{E}_{ij,u}$ happens, then, since u is the *first* intersection point, the calculations in the preceding paragraph *can* be used for u . That is, starting at u , we can trace the $\alpha_i\beta_i$ -alternating path back to z_i , choose x_i from the neighbors of z_i , trace the $\alpha_{i-1}\beta_{i-1}$ -alternating path to z_{i-1} , choose x_{i-1} from the neighbors of z_{i-1} , etc. In this way, we see that there is a positive constant c such that the number of sequences $(x_{j+1}, \dots, x_i)$ for which $\mathcal{E}_{ij,u}$ holds is at most $\Delta^{c(i-j)}$. Therefore,

$$\mathbb{P}[\mathcal{E}_{ij,u}] \leq \frac{\Delta^{c(i-j)}}{\ell^{i-j}}.$$

There are ℓ options for u , so, by the union bound,

$$\mathbb{P}[\mathcal{E}_{ij}] \leq \ell \cdot \frac{\Delta^{c(i-j)}}{\ell^{i-j}} = \frac{\Delta^{c(i-j)}}{\ell^{i-j-1}}.$$

Applying the union bound a second time and using the assumption $j \leq i - 2$, we get

$$\begin{aligned} \mathbb{P}[\text{the paths } P_0, \dots, P_T \text{ are not disjoint}] &\leq \sum_{j=0}^{T-2} \sum_{i=j+2}^T \frac{\Delta^{c(i-j)}}{\ell^{i-j-1}} \leq \Delta^c \sum_{j=0}^{T-2} \sum_{k=1}^{\infty} \frac{\Delta^{ck}}{\ell^k} \\ &= \frac{\Delta^{2c}(T-1)}{\ell - \Delta^c}. \end{aligned}$$

Remembering that $\ell = a \log n$ and $T = b \log n$, we can make the last expression as small as desired by increasing the ratio a/b . Notice that $O((\log n)^2)$ is a natural barrier for this argument, since both ℓ and T must be of the order $\Theta(\log n)$. (With more careful computations, it is possible to shave off a $\text{poly}(\log \log n)$ -factor from this bound.) It remains an interesting open problem to see if one can actually reduce the bound to $O(\log n)$ using a different method.

The rest of the paper is organized as follows. In §4, we introduce the terminology needed for the proof, such as the notion of shifting a chain. We also prove Lemmas 4.8 and 4.9, the first of which essentially asserts the existence of a Vizing chain, while the second one comes from the work of Grebík and Pikhurko and allows one to find $\gamma\delta$ -alternating paths with $\{\gamma, \delta\} \cap \{\alpha, \beta\} = \emptyset$ or $\{\gamma, \delta\} = \{\alpha, \beta\}$. Then, in §5, we describe our Multi-Step Vizing Algorithm in detail and verify some of its properties. After that, in §6, we undertake the probabilistic analysis of the Multi-Step Vizing Algorithm and finish the proof of Theorem 1.3.

4. Chains, paths, and fans

4.1. Basics

For the remainder of the paper we fix an n -vertex graph G of maximum degree Δ . For brevity, set $V := V(G)$ and $E := E(G)$. As usual, we say that two edges of G are **adjacent**, or **neighbors** of each other, if they are distinct and have a common endpoint. Throughout, we shall use the phrase “(partial) coloring” to mean a (partial) $(\Delta + 1)$ -edge-coloring of G . For convenience, if φ is a partial coloring and $e \in E \setminus \text{dom}(\varphi)$ is an uncolored edge, then we write $\varphi(e) = \mathbf{blank}$, where **blank** is a special symbol distinct from every color.

4.2. Missing colors and happy edges

Given a proper partial coloring φ and a vertex $x \in V$, we let $M(\varphi, x) \subseteq [\Delta + 1]$ be the set of all colors **missing** at x in the coloring φ , that is, $M(\varphi, x) := [\Delta + 1] \setminus \{\varphi(xy) : xy \in E\}$. Note that $|M(\varphi, x)|$ is equal to $\Delta + 1 - \deg_G(x)$ plus the number of uncolored edges incident to x ; in particular, the set $M(\varphi, x)$ is always nonempty. We say that an uncolored edge $e = xy$ is φ -**happy** if $M(\varphi, x) \cap M(\varphi, y) \neq \emptyset$. If $e = xy$ is φ -happy and α is any color in $M(\varphi, x) \cap M(\varphi, y)$, then we can extend φ by assigning to e the color α . In other words, in this case the subgraph of G with vertex set $\{x, y\}$ and edge set $\{e\}$ is augmenting for φ .

4.3. Shifts and chains

Given a proper partial coloring, we wish to “locally” modify it in order to create a happy edge (which then would allow us to enlarge the domain of the coloring). In order to achieve this, we will be “shifting” colors from colored edges to their uncolored neighbors. Let $e_0, e_1 \in E$ be two adjacent edges and let φ be a proper partial coloring. Define a coloring $\mathbf{Shift}(\varphi, e_0, e_1)$ by setting, for all $e \in E$,

$$\mathbf{Shift}(\varphi, e_0, e_1)(e) := \begin{cases} \varphi(e_1) & \text{if } e = e_0; \\ \mathbf{blank} & \text{if } e = e_1; \\ \varphi(e) & \text{otherwise.} \end{cases}$$

In other words, $\mathbf{Shift}(\varphi, e_0, e_1)$ “shifts” the color from e_1 to e_0 , leaves e_1 uncolored, and keeps the coloring of the rest of the edges unchanged. The pair (e_0, e_1) is said to be φ -**shiftable** if

- $e_0 \notin \text{dom}(\varphi)$, $e_1 \in \text{dom}(\varphi)$, and
- the coloring $\mathbf{Shift}(\varphi, e_0, e_1)$ is proper.




Fig. 6. Shifting a coloring along a chain.

The latter condition is equivalent to saying that $\varphi(e_1) \in M(\varphi, x)$, where x is the unique vertex that belongs to e_0 but not to e_1 . Note that if a pair (e_0, e_1) is φ -shiftable, then the pair (e_1, e_0) is $\text{Shift}(\varphi, e_0, e_1)$ -shiftable and

$$\text{Shift}(\text{Shift}(\varphi, e_0, e_1), e_1, e_0) = \varphi. \quad (4.1)$$

A **chain** of length $\ell \geq 1$ is a sequence $C = (e_0, \dots, e_{\ell-1})$ of edges such that e_i and e_{i+1} are adjacent, for all $0 \leq i \leq \ell - 2$. Note that we do not require the edges in a chain to be pairwise distinct; however, any two *consecutive* edges must be distinct since they are adjacent. We define the **edge set** $E(C)$ and the **vertex set** $V(C)$ of a chain $C = (e_0, \dots, e_{\ell-1})$ by setting $E(C) := \{e_0, \dots, e_{\ell-1}\}$ and making $V(C)$ be the set of all the endpoints of the edges in C . We also let $H(C)$ be the subgraph of G with vertex set $V(C)$ and edge set $E(C)$. It is clear that the graph $H(C)$ is connected. Let $\text{Start}(C) := e_0$ and $\text{End}(C) := e_{\ell-1}$ denote the first and the last edges of C , respectively, and let $\text{length}(C) := \ell$ be the length of C . Given a partial coloring φ , we define $\text{Shift}(\varphi, C)$ by iteratively shifting the colors from e_1 to e_0 , then from e_2 to e_1 , etc. (see Fig. 6). More precisely, we consider the following recursively defined sequence of colorings:

$$\begin{aligned} \text{Shift}_0(\varphi, C) &:= \varphi; \\ \text{Shift}_{i+1}(\varphi, C) &:= \text{Shift}(\text{Shift}_i(\varphi, C), e_i, e_{i+1}) \quad \text{for all } 0 \leq i \leq \ell - 2, \end{aligned}$$

and set $\text{Shift}(\varphi, C) := \text{Shift}_{\ell-1}(\varphi, C)$. We say that C is φ -**shiftable** if for all $0 \leq i \leq \ell - 2$, the pair (e_i, e_{i+1}) is $\text{Shift}_i(\varphi, C)$ -shiftable. Given a chain $C = (e_0, \dots, e_{\ell-1})$, we let $C^* := (e_{\ell-1}, \dots, e_0)$ denote the **reverse** of C . Repeated applications of (4.1) show that if C is φ -shiftable, then C^* is $\text{Shift}(\varphi, C)$ -shiftable and $\text{Shift}(\text{Shift}(\varphi, C), C^*) = \varphi$.

Fact 4.2. *Let φ be a proper partial coloring and let C be a chain. If C is φ -shiftable, then the following statements are valid:*

- (i) *the unique edge in $E(C)$ that is uncolored by φ is $\text{Start}(C)$;*
- (ii) *$\text{Shift}(\varphi, C)$ is a proper partial coloring that agrees with φ on $E \setminus E(C)$;*
- (iii) *the unique edge in $E(C)$ that is uncolored by $\text{Shift}(\varphi, C)$ is $\text{End}(C)$.*

Proof. Let $C = (e_0, \dots, e_{\ell-1})$, where $\ell = \text{length}(C)$. A straightforward induction shows that for each $0 \leq i \leq \ell - 1$, the coloring $\text{Shift}_i(\varphi, C)$ is proper, which, in particular, yields (ii). To see (i), note first that e_0 must be uncolored by φ since the pair (e_0, e_1) is φ -shiftable. Suppose now that $e \in E(C)$ is a φ -uncolored edge and let i be the least index such that $e = e_i$. If $i > 0$, then e is still uncolored by $\text{Shift}_{i-1}(\varphi, C)$, which means that the pair (e_{i-1}, e_i) is not $\text{Shift}_{i-1}(\varphi, C)$ -shiftable; a contradiction. Part (iii) follows by applying (i) with $\text{Shift}(\varphi, C)$ and C^* in place of φ and C . ■

We say that a chain C is φ -**happy** if C is φ -shiftable and the edge $\text{End}(C)$ is $\text{Shift}(\varphi, C)$ -happy.

Fact 4.3. *Let φ be a proper partial coloring and let C be a chain. If C is φ -happy, then the graph $H(C)$ is augmenting for φ .*

Proof. The coloring φ can be modified to include all the edges in $E(C)$ as follows: First, replace φ by $\psi := \text{Shift}(\varphi, C)$. Now the only uncolored edge in $E(C)$ is $\text{End}(C)$, which, by assumption, is ψ -happy, so ψ can be extended to include it. ■

Thus, our goal becomes to develop a technique for constructing “short” happy chains.

4.4. Initial segments and combinations

For a chain $C = (e_0, \dots, e_{\ell-1})$ and $1 \leq k \leq \ell$, the **initial segment** of C of length k is the chain

$$C|k := (e_0, \dots, e_{k-1}).$$

Note that if C is φ -shiftable for some proper partial coloring φ , then every initial segment of C is also φ -shiftable. Observe also that $\text{Shift}(\varphi, C|k) = \text{Shift}_{k-1}(\varphi, C)$ for all $1 \leq k \leq \ell$. We additionally consider an operation that is in some sense the opposite of taking initial segments. Namely, let C_0 and C_1 be chains such that $\text{End}(C_0) = \text{Start}(C_1)$. Then we can **combine** C_0 and C_1 into a single chain $C_0 + C_1$ by identifying the last edge of C_0 with the first edge of C_1 . More precisely, if $C_0 = (e_0, \dots, e_{\ell-1})$, $C_1 = (h_0, \dots, h_{k-1})$, and $e_{\ell-1} = h_0$, then we set

$$C_0 + C_1 := (e_0, \dots, e_{\ell-1}, h_1, \dots, h_{k-1}).$$

Notice that C_0 is an initial segment of $C_0 + C_1$. It is clear that for any proper partial coloring φ , we have

$$\text{Shift}(\varphi, C_0 + C_1) = \text{Shift}(\text{Shift}(\varphi, C_0), C_1),$$

and if C_0 is φ -shiftable and C_1 is $\text{Shift}(\varphi, C_0)$ -shiftable, then $C_0 + C_1$ is also φ -shiftable.

4.5. Path-chains

A particular type of a chain is a *path-chain*, i.e., a nonempty sequence of edges that forms a path in G in the usual graph-theoretic sense. Explicitly, a chain $P = (e_0, \dots, e_{\ell-1})$ is called a **path-chain**, or simply a **path**, if there is a sequence $(x_0, x_1, \dots, x_\ell)$ of pairwise distinct vertices such that $e_i = x_i x_{i+1}$ for all $1 \leq i \leq \ell - 1$. Such a sequence $(x_0, x_1, \dots, x_\ell)$ is necessarily unique unless $\ell = 1$, while if $P = (xy)$ is a path-chain of length 1, then there are two such sequences, namely (x, y) and (y, x) . When dealing with such length-1 path-chains, we shall still tacitly assume that one of these two sequences is chosen, so that every path-chain P has a well-defined first and a well-defined last vertex, denoted by $\text{vStart}(P)$ and $\text{vEnd}(P)$, respectively.

The aim of this subsection is to describe a particular way of constructing shifttable path-chains by considering edges of two specific colors. For the remainder of §4.5, we fix the following data:

- a proper partial coloring φ ;
- a pair of distinct colors $\alpha, \beta \in [\Delta + 1]$.

Let $G(\varphi, \alpha\beta)$ denote the spanning subgraph of G with

$$E(G(\varphi, \alpha\beta)) := \{e \in E : \varphi(e) \in \{\alpha, \beta\}\}.$$

Since φ is proper, the maximum degree of $G(\varphi, \alpha\beta)$ is at most 2, which means that every connected component of $G(\varphi, \alpha\beta)$ is either a cycle or a path (where we view isolated vertices as paths of length 0). For $x \in V$, let $G(x, \varphi, \alpha\beta)$ denote the connected component of x in $G(\varphi, \alpha\beta)$, and let $\deg(x, \varphi, \alpha\beta)$ be the degree of x in $G(\varphi, \alpha\beta)$. Note that $\deg(x, \varphi, \alpha\beta) < 2$ if and only if at least one of α, β is missing at x . We say that two vertices $x, y \in V$ are $(\varphi, \alpha\beta)$ -**related** if $G(x, \varphi, \alpha\beta) = G(y, \varphi, \alpha\beta)$.

Suppose that $xy \in E$ is an uncolored edge that is not φ -happy. We say that xy is:

- $(\varphi, \alpha\beta)$ -**hopeful** if $\deg(x, \varphi, \alpha\beta) < 2$ and $\deg(y, \varphi, \alpha\beta) < 2$;
- $(\varphi, \alpha\beta)$ -**successful** if it is $(\varphi, \alpha\beta)$ -hopeful and x and y are not $(\varphi, \alpha\beta)$ -related.

For a $(\varphi, \alpha\beta)$ -hopeful edge xy , we define a chain $P(x, y, \varphi, \alpha\beta)$ as follows. Since $\deg(y, \varphi, \alpha\beta) < 2$, the graph $G(y, \varphi, \alpha\beta)$ is a path, one of whose endpoints is y . Let $(e_1, \dots, e_k)$ be the sequence in which the edges of this path appear as it is traversed starting from y . Note that $k \geq 1$, since otherwise $\{\alpha, \beta\} \subseteq M(\varphi, y)$ and $\{\alpha, \beta\} \cap M(\varphi, x) \neq \emptyset$, which would mean that the edge xy is φ -happy. The sequence $(e_1, \dots, e_k)$ is $\alpha\beta$ -*alternating*, in the sense that the sequence of colors $\varphi(e_1), \varphi(e_2), \varphi(e_3), \varphi(e_4), \dots$ has the form $\alpha, \beta, \alpha, \beta, \dots$ or $\beta, \alpha, \beta, \alpha, \dots$. Then we define

$$P(x, y, \varphi, \alpha\beta) := (xy, e_1, \dots, e_k).$$




Fig. 7. The chain $P(x, y, \varphi, \alpha\beta)$.

This construction is illustrated in Fig. 7.

Fact 4.4. *If xy is a $(\varphi, \alpha\beta)$ -hopeful edge, then the chain $P(x, y, \varphi, \alpha\beta)$ is φ -shiftable.*

Proof. Since the edge xy is not φ -happy, $M(\varphi, x) \cap M(\varphi, y) = \emptyset$. As xy is $(\varphi, \alpha\beta)$ -hopeful, this implies that x misses one of the colors α, β and y misses the other one. Without loss of generality, assume that $\alpha \in M(\varphi, x)$ and $\beta \in M(\varphi, y)$. Set $e_0 := xy$ and $P := P(x, y, \varphi, \alpha\beta) = (e_0, e_1, \dots, e_k)$. Then e_0 is uncolored and $\varphi(e_1) = \alpha$, $\varphi(e_2) = \beta$, $\varphi(e_3) = \alpha$, $\varphi(e_4) = \beta$, etc. Since $\alpha \in M(\varphi, x)$, the pair of edges (e_0, e_1) is φ -shiftable. Straightforward induction then shows that the whole chain P is φ -shiftable. ■

If an edge xy is not just $(\varphi, \alpha\beta)$ -hopeful, but in fact $(\varphi, \alpha\beta)$ -successful, then x cannot appear on the path $G(y, \varphi, \alpha\beta)$, which implies that $P(x, y, \varphi, \alpha\beta)$ is a path-chain. Note that the first vertex of $P(x, y, \varphi, \alpha\beta)$ is x , while the last one is the endpoint of the path $G(y, \varphi, \alpha\beta)$ distinct from y .

Fact 4.5. *If xy is a $(\varphi, \alpha\beta)$ -successful edge, then the path-chain $P(x, y, \varphi, \alpha\beta)$ is φ -happy.*

Proof. Set $e_0 := xy$ and $P := P(x, y, \varphi, \alpha\beta) = (e_0, e_1, \dots, e_k)$. We already know, from Fact 4.4, that P is φ -shiftable. In the coloring $\text{Shift}(\varphi, P)$, both endpoints of the edge e_k miss the same one of the colors α, β , so e_k is $\text{Shift}(\varphi, P)$ -happy, as desired. ■

4.6. Fans

Another useful type of a chain is a *fan*. A **fan** is a chain of the form $F = (xy_0, xy_1, \dots, xy_{\ell-1})$, where $x \in V$ and $y_0, \dots, y_{\ell-1}$ are pairwise distinct neighbors of x . Given a fan $F = (xy_0, xy_1, \dots, xy_{\ell-1})$, we call the vertex x the **pivot** of F and write $\text{Pivot}(F) := x$. We also refer to $y_{\ell-1}$ as the last vertex of F and write $\text{vEnd}(F) := y_{\ell-1}$. As in the case of path-chains, this creates some ambiguity when $F = (xy)$ is a fan of length 1. When dealing with such a length-1 fan, we shall still tacitly assume that it has been decided which one of x, y is the pivot. Note that for every fan F , its last edge $\text{End}(F)$ is $\{\text{Pivot}(F), \text{vEnd}(F)\}$.




Fig. 8. The process of shifting a fan for i steps.

Fact 4.6. Let φ be a proper partial coloring and consider a fan $F = (xy_0, xy_1, \dots, xy_{\ell-1})$. Then F is φ -shiftable if and only if the following two conditions are satisfied:

- (i) $xy_0 \notin \text{dom}(\varphi)$, while $xy_i \in \text{dom}(\varphi)$ for all $1 \leq i \leq \ell - 1$;
- (ii) for each $0 \leq i \leq \ell - 2$, $\varphi(xy_{i+1}) \in M(\varphi, y_i)$.

Proof. Condition (i) is necessary for F to be φ -shiftable due to Fact 4.2, so we may assume that (i) holds. For brevity, set $\psi_i := \text{Shift}_i(\varphi, F)$ for each $0 \leq i \leq \ell - 1$. By definition, F is φ -shiftable if for all $0 \leq i \leq \ell - 2$, the pair (xy_i, xy_{i+1}) is ψ_i -shiftable. It is straightforward to check that $\psi_i(xy_i) = \text{blank}$, while $\psi_i(xy_{i+1}) = \varphi(xy_{i+1})$ (see Fig. 8). Thus, (xy_i, xy_{i+1}) is ψ_i -shiftable if and only if $\varphi(xy_{i+1}) \in M(\psi_i, y_i)$. For $i = 0$, this is just saying that $\varphi(xy_1) \in M(\varphi, y_0)$, as desired. For $i > 0$, observe that $M(\psi_i, y_i) = M(\varphi, y_i) \cup \{\varphi(xy_i)\}$. Since the coloring φ is proper, $\varphi(xy_{i+1}) \neq \varphi(xy_i)$, so $\varphi(xy_{i+1}) \in M(\psi_i, y_i)$ if and only if $\varphi(xy_{i+1}) \in M(\varphi, y_i)$, and we are done. ■

Let φ be a partial proper coloring and let $\alpha, \beta \in [\Delta + 1]$ be two distinct colors. Let F be a φ -shiftable fan that is not φ -happy. Set $x := \text{Pivot}(F)$ and $y := \text{vEnd}(F)$ (so $\text{End}(F) = xy$). By analogy with the definitions in §4.5, we say that F is:

- $(\varphi, \alpha\beta)$ -hopeful if $\deg(x, \varphi, \alpha\beta) < 2$ and $\deg(y, \varphi, \alpha\beta) < 2$;
- $(\varphi, \alpha\beta)$ -successful if it is $(\varphi, \alpha\beta)$ -hopeful and x and y are not $(\text{Shift}(\varphi, F), \alpha\beta)$ -related.

Fact 4.7. Let F be a fan and set $x := \text{Pivot}(F)$ and $y := \text{vEnd}(F)$. Define $\psi := \text{Shift}(\varphi, F)$. If F is $(\varphi, \alpha\beta)$ -hopeful (resp. $(\varphi, \alpha\beta)$ -successful), then the edge xy is $(\psi, \alpha\beta)$ -hopeful (resp. $(\psi, \alpha\beta)$ -successful).

Proof. Note that $M(\psi, x) = M(\varphi, x)$ and $M(\psi, y) \supseteq M(\varphi, y)$ (the inclusion is, in fact, strict unless F is of length 1). This implies that if F is $(\varphi, \alpha\beta)$ -hopeful, then $\deg(x, \psi, \alpha\beta) = \deg(x, \varphi, \alpha\beta) < 2$ and $\deg(y, \psi, \alpha\beta) \leq \deg(y, \varphi, \alpha\beta) < 2$, and hence the edge xy is $(\psi, \alpha\beta)$ -hopeful, as desired. The statement for $(\varphi, \alpha\beta)$ -successful fans follows. ■

Next we establish two lemmas that allow us to find successful or, at least, hopeful fans. Modulo our choice of terminology, the first of these lemmas was already known to Vizing [20] (see [6] for a textbook presentation), while the second one comes from the work of Grebík and Pikhurko [11].

Lemma 4.8 (*First fan lemma*). *Let φ be a partial proper coloring and let $xy \in E \setminus \text{dom}(\varphi)$ be an uncolored edge. Then there exists a φ -shiftable fan F with pivot x and $\text{Start}(F) = xy$ such that:*

- either F is φ -happy,
- or, for some distinct $\alpha, \beta \in [\Delta + 1]$, F is $(\varphi, \alpha\beta)$ -successful.

Proof. Assume that there is no φ -happy fan with pivot x and starting edge xy . For each neighbor z of x , fix an arbitrary color $\beta(z) \in M(\varphi, z)$. We shall recursively build a sequence $y_0, y_1, \dots$ of pairwise distinct neighbors of x such that the fan $(xy_0, xy_1, \dots)$ is φ -shiftable. Start by setting $y_0 := y$. Once the vertices $y_0, \dots, y_i$ have been determined, consider the color $\beta(y_i)$. If $\beta(y_i) \in M(\varphi, x)$, then the fan $(xy_0, \dots, xy_i)$ is φ -happy, which contradicts our assumption. Thus, $\beta(y_i) \notin M(\varphi, x)$, which means that x has a (unique) neighbor z with $\varphi(xz) = \beta(y_i)$. Note that $z \neq y_i$, because the color $\beta(y_i)$ is missing at y_i but not at z . Similarly, $z \neq y_0$ because the edge xy_0 is uncolored, while xz is not. Now, if $z \in \{y_1, \dots, y_{i-1}\}$, then we stop the construction, and otherwise we set $y_{i+1} := z$ and proceed to the next step. Since x has only finitely many neighbors, this construction must terminate. That is, we shall eventually obtain a sequence $y_0, \dots, y_k$ of distinct neighbors of x such that:

- (i) $y_0 = y$;
- (ii) for all $0 \leq i \leq k$, $\beta(y_i) \notin M(\varphi, x)$;
- (iii) for all $0 \leq i \leq k-1$, $\varphi(xy_{i+1}) = \beta(y_i)$;
- (iv) $\beta(y_k) = \varphi(xy_j)$ for some $1 \leq j \leq k-1$.

Set $\beta := \beta(y_k)$. If j is the index from (iv), then (iii) yields $\beta = \varphi(xy_j) = \beta(y_{j-1})$; that is, β is missing both at y_k and at y_{j-1} . Now let α be any color in $M(\varphi, x)$. By (ii), $\alpha \neq \beta$. Define

$$F := (xy_0, \dots, xy_k) \quad \text{and} \quad F' := (xy_0, \dots, xy_{j-1}) = F|j.$$

By construction, both F and F' are $(\varphi, \alpha\beta)$ -hopeful fans. We claim that at least one of F, F' is $(\varphi, \alpha\beta)$ -successful. For brevity, let

$$\psi := \text{Shift}(\varphi, F) \quad \text{and} \quad \psi' := \text{Shift}(\varphi, F') = \text{Shift}_{j-1}(\varphi, F).$$

Suppose that F is not $(\varphi, \alpha\beta)$ -successful. This means that the vertices x and y_k are $(\psi, \alpha\beta)$ -related. By Fact 4.7, the edge xy_k is $(\psi, \alpha\beta)$ -hopeful, so $G_{\alpha\beta}(\psi, x)$ must be a

path with endpoints x and y_k . Furthermore, $\psi(xy_{j-1}) = \beta$, so the (unique) edge of this path incident to x is xy_{j-1} . All the other edges of this path avoid the vertex x , which means that their colors cannot be changed by shifting a fan with pivot x . In particular, they have the same colors in φ , ψ , and ψ' , which implies that the vertices y_{j-1} and y_k are $(\psi', \alpha\beta)$ -related. Since both these vertices miss the color β in ψ' , the graph $G_{\alpha\beta}(\psi', y_{j-1})$ is a path with endpoints y_{j-1} and y_k . This path cannot include x (because x is missing α), and thus y_{j-1} and x are not $(\psi', \alpha\beta)$ -related, as desired. ■

Lemma 4.9 (Second fan lemma). *Let φ be a partial proper coloring and let $xy \in E \setminus \text{dom}(\varphi)$ be an uncolored edge. Fix a pair of distinct colors $\alpha, \beta \in [\Delta + 1]$ such that $\deg(x, \varphi, \alpha\beta) = 1$. Then there exists a φ -shiftable fan F with pivot x and $\text{Start}(F) = xy$ such that:*

- either F is φ -happy,
- or, for some distinct $\gamma, \delta \in [\Delta + 1]$ with $\{\gamma, \delta\} \cap \{\alpha, \beta\} = \emptyset$, F is $(\varphi, \gamma\delta)$ -successful,
- or F is $(\varphi, \alpha\beta)$ -hopeful, $\text{vEnd}(F) \neq y$, and no edge in F is colored α or β by φ .

Proof. Again, we may assume that there is no φ -happy fan with pivot x and starting edge xy . Since $\deg(x, \varphi, \alpha\beta) = 1$, precisely one of the colors α, β is missing at x . For concreteness, suppose that $\alpha \in M(\varphi, x)$ and $\beta \notin M(\varphi, x)$. For each neighbor z of x , fix a color $\delta(z) \in M(\varphi, z)$; it does not matter which color from $M(\varphi, z)$ we pick, except that we make sure that $\delta(y) \neq \beta$ (this is possible since y is incident to an uncolored edge and thus $|M(\varphi, y)| \geq 2$). As in the proof of Lemma 4.8, we build a sequence $y_0, y_1, \dots$ of pairwise distinct neighbors of x such that the fan $(xy_0, xy_1, \dots)$ is φ -shiftable. Start by setting $y_0 := y$. Once the vertices $y_0, \dots, y_i$ have been determined, consider the color $\delta(y_i)$. If $\delta(y_i) \in M(\varphi, x)$, then the fan $(xy_0, \dots, xy_i)$ is φ -happy, which contradicts our assumption. Thus, $\delta(y_i) \notin M(\varphi, x)$. At this point, if $\delta(y_i) = \beta$, then we stop the construction. Otherwise we proceed as in the proof of Lemma 4.8. In other words, if $\delta(y_i) \neq \beta$, then we let z be the (unique) neighbor of x such that $\varphi(xz) = \delta(y_i)$. If $z \in \{y_1, \dots, y_{i-1}\}$, then we stop, and otherwise, we let $y_{i+1} := z$ and move on to the next step. Since x has only finitely many neighbors, this construction must terminate. Hence, we eventually obtain a sequence $y_0, \dots, y_k$ of distinct neighbors of x such that:

- (i) $y_0 = y$;
- (ii) for all $0 \leq i \leq k$, $\delta(y_i) \notin M(\varphi, x)$;
- (iii) for all $0 \leq i \leq k - 1$, $\varphi(xy_{i+1}) = \delta(y_i)$;
- (iv) for all $0 \leq i \leq k - 1$, $\delta(y_i) \neq \beta$;
- (v) either $\delta(y_k) = \beta$, or else, $\delta(y_k) = \varphi(xy_j)$ for some $1 \leq j \leq k - 1$.

Now we consider two cases.

Case 1. $\delta(y_k) = \beta$.

Then the fan $F := (xy_0, \dots, xy_k)$ is, by construction, $(\varphi, \alpha\beta)$ -hopeful. Furthermore, since $\delta(y) \neq \beta$, we have $\mathbf{vEnd}(F) = y_k \neq y$. Finally, for each $1 \leq i \leq k$, we have $\varphi(xy_i) = \delta(y_{i-1})$, and since $\delta(y_{i-1}) \neq \beta$ by (iv) and $\delta(y_{i-1}) \notin M(\varphi, x)$ by (ii), we conclude that no edge in F is colored α or β .

Case 2. $\delta(y_k) \neq \beta$ and $\delta(y_k) = \varphi(xy_j)$ for some $1 \leq j \leq k-1$.

Set $\delta := \delta(y_k)$ and let j be the index from (v), so that the color δ is missing both at y_k and at y_{j-1} . Since x is incident to an uncolored edge, we can pick a color $\gamma \in M(x, \varphi)$ distinct from α . Now $\{\gamma, \delta\} \cap \{\alpha, \beta\} = \emptyset$ and the argument from the proof of Lemma 4.8 shows that at least one of the fans

$$F := (xy_0, \dots, xy_k) \quad \text{and} \quad F' := (xy_0, \dots, xy_{j-1})$$

is $(\varphi, \gamma\delta)$ -successful, as desired. ■

As an illustration, we can now use Lemma 4.8 to complete the proof of Vizing's theorem:

Corollary 4.10 (to Lemma 4.8). *For every proper partial coloring φ and every uncolored edge $e = xy$, there is a connected augmenting subgraph $H \subseteq G$ whose only uncolored edge is e .*

Proof. By Fact 4.3, it suffices to find a φ -happy chain C with $\mathbf{Start}(C) = e$. To this end, let F be a fan satisfying the conclusion of Lemma 4.8. If F is φ -happy, then we are done, so assume that instead there are distinct colors $\alpha, \beta \in [\Delta + 1]$ such that F is $(\varphi, \alpha\beta)$ -successful. Let $z := \mathbf{vEnd}(F)$ and set $\psi := \mathbf{Shift}(\varphi, F)$. By Fact 4.7, the edge xz is $(\psi, \alpha\beta)$ -successful, so, by Fact 4.5, the path-chain $P := P(\psi, x, z, \alpha\beta)$ is ψ -happy. Therefore, the chain $F + P$ is as desired. ■

5. The Multi-Step Vizing Algorithm

5.1. The sets $R^+(x, \varphi)$ and $R^-(y, \varphi)$

Before we begin describing our construction of multi-step Vizing chains, we need to introduce some notation that will be useful for keeping track of intersections between different alternating paths. Let φ be a proper partial coloring and let $\alpha, \beta \in [\Delta + 1]$ be distinct colors. Given $x \in V$, we let $R^+(x, \varphi, \alpha\beta) \subseteq V$ denote the set of all vertices y that satisfy at least one of the following two conditions:

- $y = x$ or y is adjacent to x ,
- or there is a neighbor z of x such that $\deg(z, \varphi, \alpha\beta) < 2$ and z and y are $(\varphi, \alpha\beta)$ -related.

We also define $R^-(y, \varphi, \alpha\beta) := \{x \in V : y \in R^+(x, \varphi, \alpha\beta)\}$ and set

$$R^+(x, \varphi) := \bigcup_{\alpha \neq \beta} R^+(x, \varphi, \alpha\beta) \quad \text{and} \quad R^-(y, \varphi) := \bigcup_{\alpha \neq \beta} R^-(y, \varphi, \alpha\beta).$$

Lemma 5.1. *For each proper partial coloring φ and every vertex $y \in V$, we have $|R^-(y, \varphi)| \leq (\Delta + 1)^3$.*

Proof. At most Δ elements of $R^-(y, \varphi)$ are neighbors of y , and for each (unordered) pair of colors α, β , there are at most two vertices z that are $(\varphi, \alpha\beta)$ -related to y and satisfy $\deg(z, \varphi, \alpha\beta) < 2$. Each such vertex z has at most Δ neighbors, so

$$|R^-(y, \varphi)| \leq 1 + \Delta + 2\Delta \binom{\Delta + 1}{2} = \Delta^3 + \Delta^2 + \Delta + 1 \leq (\Delta + 1)^3. \quad \blacksquare$$

Given a chain C , we write $R^-(C, \varphi) := \bigcup_{y \in V(C)} R^-(y, \varphi)$.

Lemma 5.2. *Let φ be a proper partial coloring and let C be a φ -shiftable chain. Set $\psi := \text{Shift}(\varphi, C)$. If x is a vertex that is not in $R^-(C, \varphi) \cap R^-(C, \psi)$, then $R^+(x, \varphi) = R^+(x, \psi)$.*

Proof. It is enough to show that if $x \notin R^-(C, \varphi)$, then $R^+(x, \varphi) = R^+(x, \psi)$; the conclusion for $x \notin R^-(C, \psi)$ then follows by exchanging the roles of φ and ψ and replacing C by its reverse C^* . So, suppose that $x \notin R^-(C, \varphi)$. We first prove that $R^+(x, \varphi) \subseteq R^+(x, \psi)$. Take any vertex $y \in R^+(x, \varphi)$. If $y = x$ or y is adjacent to x , then $y \in R^+(x, \psi)$ by definition, so suppose instead that there exist colors α, β and a neighbor z of x such that $\deg(z, \varphi, \alpha\beta) < 2$ and z and y are $(\varphi, \alpha\beta)$ -related. In other words, $G(z, \varphi, \alpha\beta)$ is a path starting at z and containing y . Since $x \notin R^-(C, \varphi)$, no vertex on this path is incident to an edge in C . This implies that $G(z, \psi, \alpha\beta) = G(z, \varphi, \alpha\beta)$, so $G(z, \psi, \alpha\beta)$ is still a path starting at z and passing through y , and hence $y \in R^+(x, \psi)$, as desired.

Now we show $R^+(x, \psi) \subseteq R^+(x, \varphi)$. Take a vertex $y \in R^+(x, \psi)$. Again, if $y = x$ or y is adjacent to x , then we are done, so assume that there exist colors α, β and a neighbor z of x such that $\deg(z, \psi, \alpha\beta) < 2$ and z and y are $(\psi, \alpha\beta)$ -related. Note that $\deg(z, \varphi, \alpha\beta) < 2$ as well; otherwise, some edge incident to z would be colored differently in φ and ψ , which would imply that $z \in V(C)$ and thus $x \in R^-(C, \varphi)$. Therefore, the graph $G(z, \varphi, \alpha\beta)$ is a path starting at z . As before, we observe that no vertex on this path can be incident to an edge in C , since $x \notin R^-(C, \varphi)$. But this means that $G(z, \psi, \alpha\beta) = G(z, \varphi, \alpha\beta)$, so if z and y are $(\psi, \alpha\beta)$ -related, then they are also $(\varphi, \alpha\beta)$ -related, as desired. $\blacksquare$

5.2. The algorithm

For the remainder of §5, we fix the following data:

- a proper partial coloring φ ;
- an uncolored edge $e = xy \in E \setminus \text{dom}(\varphi)$.

We also fix a numerical parameter $\ell \in \mathbb{N}^+$. (We will eventually make ℓ be a value of order $\text{poly}(\Delta) \log n$, but for now this does not matter.) We shall describe a randomized procedure, called the **Multi-Step Vizing Algorithm**, that attempts to build a φ -happy chain C with starting edge e . The chain C will have the form $C = F_0 + P_0 + F_1 + P_1 + \dots$, where each F_i is a fan and each P_i is a path-chain. The construction proceeds in steps, indexed by natural numbers (the total number of steps may be finite or infinite). On Step i , we build the fan F_i and the path P_i .

5.2.1. Initial set-up

Initially, we set $\varphi_0 := \varphi$, $e_0 := e$, $x_0 := x$, and $y_0 := y$.

5.2.2. Input for Step i

At the start of Step i , we have already constructed a number of interrelated structures. However, in order to describe how Step i operates, only the following objects will be needed:

- (In1) a sequence of proper partial colorings $\varphi_0, \psi_0, \dots, \varphi_{i-1}, \psi_{i-1}, \varphi_i$;
- (In2) an edge $e_i = x_i y_i$ with $\varphi_i(e_i) = \text{blank}$;
- (In3) if $i > 0$, a pair of distinct colors α_{i-1} and β_{i-1} such that $\deg(x_i, \varphi_i, \alpha_{i-1} \beta_{i-1}) = 1$;
- (In4) a sequence $F_0, P_0, \dots, F_{i-1}, P_{i-1}$, where each F_j is a fan and each P_j is a path-chain.

5.2.3. Step i

We start by constructing a fan F_i . To this end, we consider two cases.

Case 1. *There is a φ_i -happy fan F with $\text{Pivot}(F) = x_i$ and $\text{Start}(F) = e_i$.*

Then we make F_i be any such fan and stop the construction. In this case, the algorithm succeeds.

Case 2. *There is no φ_i -happy fan F with $\text{Pivot}(F) = x_i$ and $\text{Start}(F) = e_i$.*

If $i = 0$, then we use Lemma 4.8 to obtain a φ_0 -shiftable fan F_0 with pivot x_0 and $\text{Start}(F_0) = e_0$ and a pair of distinct colors α_0, β_0 such that F_0 is $(\varphi_0, \alpha_0 \beta_0)$ -successful. If, on the other hand, $i > 0$, then we apply Lemma 4.9 instead with $\alpha_{i-1}, \beta_{i-1}$ in place of α, β (this is possible by (In3)) and obtain a φ_i -shiftable fan F_i with pivot x_i and $\text{Start}(F_i) = e_i$ and a pair of distinct colors α_i, β_i such that:

- (F1) either $\{\alpha_i, \beta_i\} \cap \{\alpha_{i-1}, \beta_{i-1}\} = \emptyset$ and F_i is $(\varphi_i, \alpha_i \beta_i)$ -successful,

(F2) or the following statements are valid:

- (F2a) $\{\alpha_i, \beta_i\} = \{\alpha_{i-1}, \beta_{i-1}\}$;
- (F2b) F_i is $(\varphi_i, \alpha_i\beta_i)$ -hopeful;
- (F2c) $\text{vEnd}(F_i) \neq y_i$; and
- (F2d) no edge in F_i is colored α_i or β_i by φ_i .

Note that regardless of whether (F1) or (F2) holds, F_i is $(\varphi_i, \alpha_i\beta_i)$ -hopeful. Let

$$z_i := \text{vEnd}(F_i) \quad \text{and} \quad \psi_i := \text{Shift}(\varphi_i, F_i).$$

By Fact 4.7, the edge $x_i z_i$ is $(\psi_i, \alpha_i\beta_i)$ -hopeful, so we can define

$$Q_i := P(x_i, z_i, \psi_i, \alpha_i\beta_i).$$

Set $u_i(0) := x_i$, $u_i(1) := z_i$ and write $Q_i = (u_i(0)u_i(1), u_i(1)u_i(2), u_i(2)u_i(3), \dots)$. Next we proceed to define a path P_i , which will be an initial segment of Q_i . There are again a few cases to consider.

Subcase 2.1. *The fan F_i is $(\varphi_i, \alpha_i\beta_i)$ -successful and $\text{length}(Q_i) \leq \ell$.*

Then we let $P_i := Q_i$ and stop the construction. In this case, the algorithm **succeeds**.

Subcase 2.2. *The fan F_i is $(\varphi_i, \alpha_i\beta_i)$ -successful and $\text{length}(Q_i) \geq \ell + 1$.*

Note that since $\text{length}(Q_i) \geq \ell + 1$, the vertex $u_i(\ell + 1)$ is well-defined. Pick an integer $1 \leq \ell_i \leq \ell$ uniformly at random and set

$$x_{i+1} := u_i(\ell_i), \quad y_{i+1} := u_i(\ell_i + 1), \quad \text{and} \quad P_i := Q_i|(\ell_i + 1).$$

By construction, P_i is a path-chain of length $\ell_i + 1 \leq \ell + 1$ and we have

$$\text{vEnd}(P_i) = y_{i+1} \quad \text{and} \quad e_{i+1} := x_{i+1}y_{i+1} = \text{End}(P_i).$$

There are now two more subcases.

Subsubcase 2.2.1. *At least one of the following holds:*

- $x_{i+1} \in R^{\leftarrow}(F_j, \psi_j)$ for some $j \leq i$,
- or $x_{i+1} \in R^{\leftarrow}(P_j, \psi_j)$ for some $j \leq i - 1$.

In this case we stop the construction and declare that the algorithm **fails**.

Subsubcase 2.2.2. *The following statements are valid:*




Fig. 9. The colorings φ_i , ψ_i , and φ_{i+1} . In this case, $\ell_i = 5$.

- $x_{i+1} \notin R^+(F_j, \psi_j)$ for all $j \leq i$; and
- $x_{i+1} \notin R^+(P_j, \psi_j)$ for all $j \leq i - 1$.

Then we let $\varphi_{i+1} := \text{Shift}(\psi_i, P_i)$. Note that the unique edge of color α_i or β_i in φ_{i+1} incident to x_{i+1} is $u_i(\ell_i - 1)x_{i+1}$, so $\deg(x_{i+1}, \varphi_{i+1}, \alpha_i\beta_i) = 1$ (see Fig. 9c). Thus, (In3) holds with $i + 1$ in place of i , and hence we may continue to Step $i + 1$. This finishes Subcase 2.2.

Subcase 2.3. The fan F_i is not $(\varphi_i, \alpha_i\beta_i)$ -successful.

We will show that this situation cannot occur (see Lemma 5.5). For now let us say that in this case the procedure stops and the algorithm **fails**.

5.2.4. Output of Step i

Here is a brief résumé of the possible outcomes for Step i :

- (Out1) In Case 1, the algorithm successfully outputs a φ -happy chain $F_0 + P_0 + \cdots + P_{i-1} + F_i$.
- (Out2) In Subcase 2.1, the algorithm successfully outputs a φ -happy chain $F_0 + P_0 + \cdots + F_i + P_i$.
- (Out3) In Subsubcase 2.2.1, the algorithm fails.
- (Out4) In Subsubcase 2.2.2, the algorithm moves on to Step $i + 1$.
- (Out5) In Subcase 2.3, the algorithm fails, but we will prove in Lemma 5.5 that this cannot happen.

Note that the only case in which the algorithm moves on to the next step is Subsubcase 2.2.2.

5.3. The backtracking lemma

In this subsection we prove a lemma that plays a central role in our analysis of the Multi-Step Vizing Algorithm. In essence, it formalizes the idea of tracing back alternating paths that we sketched in §3.

Lemma 5.3 (*Backtracking*). *Suppose that on Step i of the Multi-Step Vizing Algorithm, Case 2 occurs. Fix a vertex $u \in V(Q_i)$. Then $x_i \in R^-(u, \psi_j) \cap R^-(u, \varphi_j)$ for all $j \leq i$, and, in particular, $x_i \in R^-(u, \varphi)$.*

Proof. Let $\alpha := \alpha_i$ and $\beta := \beta_i$. If $u = x_i$, then we are done, so assume that $u \neq x_i$. Then $x_i \in R^-(u, \psi_i)$ since z_i is a neighbor of x_i with $\deg(z_i, \psi_i, \alpha\beta) < 2$ and u is $(\psi_i, \alpha\beta)$ -related to z_i . Next we show that

$$x_i \in R^-(u, \varphi_i, \alpha\beta). \quad (5.4)$$

Since the fan F_i is $(\varphi_i, \alpha\beta)$ -hopeful, we have $\deg(z_i, \varphi_i, \alpha\beta) < 2$. Suppose, towards a contradiction, that u and z_i are not $(\varphi_i, \alpha\beta)$ -related. By construction, u and z_i are joined by a path in $G(\psi_i, \alpha\beta)$. If u and z_i are not $(\varphi_i, \alpha\beta)$ -related, then at least one edge on this path must have a color different from α and β in φ_i , which means that it must belong to F_i . But every edge in F_i contains x_i , so x_i is an internal vertex on some path in $G(\psi_i, \alpha\beta)$, which is impossible as $\deg(x_i, \psi_i, \alpha\beta) = \deg(x_i, \varphi_i, \alpha\beta) < 2$.

If $i = 0$, then we are done, so assume that $i > 0$. On Step $i - 1$, the algorithm must have entered Subsubcase 2.2.2 (otherwise it would have stopped without reaching Step i), which means that the following statements hold:

- $x_i \notin R^-(F_j, \psi_j)$ for all $j \leq i - 1$; and
- $x_i \notin R^-(P_j, \psi_j)$ for all $j \leq i - 2$.

By iteratively applying Lemma 5.2, we conclude that

$$R^+(x_i, \varphi_0) = R^+(x_i, \psi_0) = R^+(x_i, \varphi_1) = R^+(x_i, \psi_1) = \cdots = R^+(x_i, \psi_{i-1}).$$

Thus, it remains to argue that $x_i \in R^-(u, \psi_{i-1})$. We will in fact show that

$$x_i \in R^-(u, \psi_{i-1}, \alpha\beta).$$

There are two possibilities to consider. If on Step i , situation (F1) occurs, then $\{\alpha, \beta\} \cap \{\alpha_{i-1}, \beta_{i-1}\} = \emptyset$. The only colors that are involved in the change from ψ_{i-1} to φ_i are α_{i-1} and β_{i-1} , so in this case $G(\psi_{i-1}, \alpha\beta) = G(\varphi_i, \alpha\beta)$. In particular, (5.4) yields

$x_i \in R^-(u, \psi_{i-1}, \alpha\beta)$, as desired. Now assume that on Step i , situation (F2) occurs and, in particular, $\{\alpha_{i-1}, \beta_{i-1}\} = \{\alpha, \beta\}$.

Claim. *Under the above assumptions, $\deg(z_i, \psi_{i-1}, \alpha\beta) = \deg(z_i, \varphi_i, \alpha\beta) < 2$.*

Proof. When the path P_{i-1} is shifted, the only vertices whose degrees in the $\alpha\beta$ -colored graph change are the first and the last two vertices of P_{i-1} , i.e., x_{i-1} , z_{i-1} , x_i , and y_i . Since x_i and z_i are adjacent, $z_i \neq x_i$. Also, $z_i \neq y_i$ by (F2c). The only remaining options are $z_i = x_{i-1}$ and $z_i = z_{i-1}$, but both of them imply $x_i \in R^-(F_{i-1}, \psi_{i-1})$, which is a contradiction. $\dashv$

It remains to check that u and z_i are $(\psi_{i-1}, \alpha\beta)$ -related. Suppose, towards a contradiction, that they are not. This means that the subgraphs $G(z_i, \psi_{i-1}, \alpha\beta)$ and $G(z_i, \varphi_i, \alpha\beta)$ are distinct, and hence at least one vertex in $G(z_i, \psi_{i-1}, \alpha\beta)$ is incident to an edge that is colored differently in ψ_{i-1} and in φ_i . All such vertices, except x_{i-1} , are $(\psi_{i-1}, \alpha\beta)$ -related to z_{i-1} (see Fig. 9b for an illustration). Thus, z_i is $(\psi_{i-1}, \alpha\beta)$ -related either to x_{i-1} or to z_{i-1} . In any case, $x_i \in R^-(F_{i-1}, \psi_{i-1})$; a contradiction. $\blacksquare$

Next we apply Lemma 5.3 to show that Subcase 2.3 cannot occur:

Lemma 5.5. *On Step i of the Multi-Step Vizing Algorithm, Subcase 2.3 cannot occur.*

Proof. Let $\alpha := \alpha_i$ and $\beta := \beta_i$ and suppose, towards a contradiction, that the fan F_i is not $(\varphi_i, \alpha\beta)$ -successful. This implies that $i > 0$ and (F2) holds; in particular, $\{\alpha_{i-1}, \beta_{i-1}\} = \{\alpha, \beta\}$. Since F_i is $(\varphi_i, \alpha\beta)$ -hopeful but not $(\varphi_i, \alpha\beta)$ -successful, the vertices x_i and z_i are $(\psi_i, \alpha\beta)$ -related. From (F2d), we conclude that $G(\psi_i, \alpha\beta) = G(\varphi_i, \alpha\beta)$, so x_i and z_i are also $(\varphi_i, \alpha\beta)$ -related. Since $\{\alpha_{i-1}, \beta_{i-1}\} = \{\alpha, \beta\}$, x_i is $(\varphi_i, \alpha\beta)$ -related to x_{i-1} (see Fig. 9c for an illustration). Therefore, z_i and x_{i-1} are $(\varphi_i, \alpha\beta)$ -related, and, by (F2d) again, z_i and x_{i-1} are $(\psi_i, \alpha\beta)$ -related. This implies that $x_{i-1} \in V(Q_i)$, so, by Lemma 5.3, $x_i \in R^-(x_{i-1}, \psi_{i-1})$ and thus $x_i \in R^-(F_{i-1}, \psi_{i-1})$. But if this were the case, then the algorithm would have entered Subsubcase 2.2.1 and stopped on Step $i-1$; a contradiction. $\blacksquare$

6. Probabilistic analysis of the Multi-Step Vizing Algorithm

6.1. The algorithm is likely to succeed

We use the set-up and the notation of §5.2. Our goal is to obtain a lower bound on the probability that the Multi-Step Vizing Algorithm terminates successfully within a limited number of steps.

Lemma 6.1. Fix $T \in \mathbb{N}^+$ and set $\lambda := \ell/(\Delta + 1)^3$. Assuming that $\lambda > 1$, the probability that the Multi-Step Vizing Algorithm successfully terminates on Step i for some $i < T$ is at least

$$1 - \frac{n}{\lambda^T} - \frac{3T(\Delta + 1)^3}{\lambda - 1}.$$

Proof. First we bound the probability that the algorithm enters Step T . Fix a vertex $u \in V$ and consider the probability $\mathbb{P}[x_T = u]$. Notice that if $x_T = u$, then, by Lemma 5.3,

$$x_T = u, \quad x_{T-1} \in R^-(x_T, \varphi), \quad x_{T-2} \in R^-(x_{T-1}, \varphi), \quad \dots, \quad \text{and} \quad x_1 \in R^-(x_2, \varphi). \quad (6.2)$$

By Lemma 5.1, the number of sequences $(x_1, \dots, x_T)$ satisfying (6.2) is at most $(\Delta + 1)^{3(T-1)}$. Since each of the vertices $x_1, \dots, x_T$ is chosen uniformly at random from a set of ℓ candidates, we conclude that

$$\mathbb{P}[x_T = u] \leq \frac{(\Delta + 1)^{3(T-1)}}{\ell^T} \leq \frac{1}{\lambda^T}.$$

Since the vertex x_T must be defined for the algorithm to enter Step T , we have

$$\mathbb{P}[\text{the algorithm enters Step } T] \leq \sum_{u \in V} \mathbb{P}[x_T = u] \leq \frac{n}{\lambda^T}.$$

Next we bound the probability that the algorithm fails on Step i for some $i < T$. Thanks to Lemma 5.5, the only situation in which the algorithm fails on Step i is Subsubcase 2.2.1, i.e.,

- $x_{i+1} \in R^-(F_j, \psi_j)$ for some $j \leq i$,
- or $x_{i+1} \in R^-(P_j, \psi_j)$ for some $j \leq i - 1$.

Claim. For all $j \leq i$, we have

$$\mathbb{P}[x_{i+1} \in R^-(F_j, \psi_j)] \leq \frac{\Delta + 1}{\lambda^{i-j+1}} \quad \text{and} \quad \mathbb{P}[x_{i+1} \in R^-(P_j, \psi_j)] \leq \frac{2(\Delta + 1)^3}{\lambda^{i-j}}.$$

Proof. We will show that the desired upper bounds hold even if we allow an adversary to specify the outcome of the algorithm's execution until it is time to pick ℓ_j (if the algorithm terminates earlier, or if it does not enter Subcase 2.2 on Step j , then the events $x_{i+1} \in R^-(F_j, \psi_j)$ and $x_{i+1} \in R^-(P_j, \psi_j)$ cannot happen, so we may ignore that possibility). So, fix the outcome of the algorithm's execution up to the stage when ℓ_j is supposed to be chosen and let $\mathbb{P}^*$ denote conditional probability subject to this outcome. At this point, we have already constructed the fan F_j , the coloring ψ_j , and the path $Q_j = (u_j(0)u_j(1), u_j(1)u_j(2), \dots)$ of length at least $\ell + 1$, but we have not yet picked

the vertex x_{j+1} . Fix any $u \in V$ and consider the (conditional) probability $\mathbb{P}^*[x_{i+1} = u]$. If $x_{i+1} = u$, then, by Lemma 5.3,

$$x_{i+1} = u, \quad x_i \in R^{\leftarrow}(x_{i+1}, \varphi), \quad x_{i-1} \in R^{\leftarrow}(x_i, \varphi), \quad \dots, \quad \text{and} \quad x_{j+1} \in R^{\leftarrow}(x_{j+2}, \varphi). \quad (6.3)$$

By Lemma 5.1, the number of sequences $(x_{j+1}, \dots, x_{i+1})$ satisfying (6.3) is at most $(\Delta + 1)^{3(i-j)}$. Since each of the vertices $x_{j+1}, \dots, x_{i+1}$ is chosen uniformly at random from a set of ℓ candidates, we get

$$\mathbb{P}^*[x_{i+1} = u] \leq \frac{(\Delta + 1)^{3(i-j)}}{\ell^{i-j+1}} = \frac{1}{\lambda^{i-j}\ell}.$$

Set $U := R^{\leftarrow}(F_j, \psi_j)$. By Lemma 5.1, we have $|U| \leq (\Delta + 1)^3 |V(F_j)| \leq (\Delta + 1)^4$, so

$$\mathbb{P}^*[x_{i+1} \in R^{\leftarrow}(F_j, \psi_j)] = \sum_{u \in U} \mathbb{P}^*[x_{i+1} = u] \leq \frac{(\Delta + 1)^4}{\lambda^{i-j}\ell} = \frac{\Delta + 1}{\lambda^{i-j+1}}.$$

Regardless of the value ℓ_j , the set $R^{\leftarrow}(P_j, \psi_j)$ is contained in $W := \bigcup_{k=0}^{\ell+1} R^{\leftarrow}(u_j(k), \psi_j)$. By Lemma 5.1, $|W| \leq (\Delta + 1)^3(\ell + 2) \leq 2(\Delta + 1)^3\ell$ (where we are using that $\ell > (\Delta + 1)^3 \geq 1$, so $\ell + 2 \leq 2\ell$), and hence

$$\mathbb{P}^*[x_{i+1} \in R^{\leftarrow}(P_j, \psi_j)] \leq \sum_{u \in W} \mathbb{P}^*[x_{i+1} = u] \leq \frac{2(\Delta + 1)^3}{\lambda^{i-j}}. \quad \dashv$$

Using the above claim and the union bound, we obtain

$$\begin{aligned} & \mathbb{P}[\text{the algorithm fails on Step } i \text{ for some } i < T] \\ & \leq \sum_{j=0}^{T-1} \sum_{i=j}^{\infty} \mathbb{P}[x_{i+1} \in R^{\leftarrow}(F_j, \psi_j)] + \sum_{j=0}^{T-2} \sum_{i=j+1}^{\infty} \mathbb{P}[x_{i+1} \in R^{\leftarrow}(P_j, \psi_j)] \\ \text{[by the claim]} \quad & \leq \sum_{j=0}^{T-1} \sum_{i=j}^{\infty} \frac{\Delta + 1}{\lambda^{i-j+1}} + \sum_{j=0}^{T-2} \sum_{i=j+1}^{\infty} \frac{2(\Delta + 1)^3}{\lambda^{i-j}} \\ & = \sum_{j=0}^{T-1} \sum_{k=1}^{\infty} \frac{\Delta + 1}{\lambda^k} + \sum_{j=0}^{T-2} \sum_{k=1}^{\infty} \frac{2(\Delta + 1)^3}{\lambda^k} \\ \text{[since } \lambda > 1] \quad & = \frac{T(\Delta + 1)}{\lambda - 1} + \frac{2(T-1)(\Delta + 1)^3}{\lambda - 1} \leq \frac{3T(\Delta + 1)^3}{\lambda - 1}. \end{aligned}$$

Putting everything together, we conclude that

$$\begin{aligned} & \mathbb{P}[\text{the algorithm succeeds on Step } i \text{ for some } i < T] \\ & = 1 - \mathbb{P}[\text{the algorithm enters Step } T] \end{aligned}$$

$$\begin{aligned}
& - \mathbb{P}[\text{the algorithm fails on Step } i \text{ for some } i < T] \\
& \leq 1 - \frac{n}{\lambda^T} - \frac{3T(\Delta + 1)^3}{\lambda - 1}. \quad \blacksquare
\end{aligned}$$

Lemma 6.1 immediately yields the existence of “small” augmenting subgraphs:

Corollary 6.4 (to Lemma 6.1). *Assuming n is sufficiently large, the following statement holds: For every proper partial coloring φ and every uncolored edge $e = xy$, there is a connected augmenting subgraph H whose only uncolored edge is e with $|E(H)| \leq (\Delta + 1)^6(\log n)^2$.*

Proof. By Fact 4.3, it suffices to find a φ -happy chain C with $\text{Start}(C) = e$ and

$$\text{length}(C) \leq (\Delta + 1)^6(\log n)^2.$$

To this end, fix a small positive constant ε (any $\varepsilon < 1/3$ would work), let $T := \lfloor \varepsilon \log n \rfloor$, and run the Multi-Step Vizing Algorithm with $e_0 = e$ and

$$\ell := (\Delta + 1)^6 \lfloor \log n \rfloor.$$

Since $\lambda := \ell/(\Delta + 1)^3 = (\Delta + 1)^3 \lfloor \log n \rfloor > 1$, we can use Lemma 6.1 to conclude that the algorithm successfully terminates on Step i for some $i < T$ with probability at least

$$1 - \frac{n}{\lambda^T} - \frac{3T(\Delta + 1)^3}{\lambda - 1} = 1 - n \left(\frac{1}{(\Delta + 1)^3 \lfloor \log n \rfloor} \right)^{\lfloor \varepsilon \log n \rfloor} - 3\varepsilon - o(1) = 1 - 3\varepsilon - o(1) > 0.$$

Whenever the algorithm terminates successfully on Step $i < T$, it produces a φ -happy chain C with $\text{Start}(C) = e$ of the form $C = F_0 + P_0 + \cdots + P_{i-1} + F_i$ or $C = F_0 + P_0 + \cdots + F_i + P_i$. Since $\text{length}(F_i) \leq \Delta$ and $\text{length}(P_i) \leq \ell + 1$ for all i , we obtain

$$\text{length}(C) \leq T(\ell + \Delta) \leq (\varepsilon + o(1))(\Delta + 1)^6(\log n)^2,$$

which yields the desired bound. $\blacksquare$

6.2. Proof of Theorem 1.3

For the reader’s convenience, we recall the statement of Theorem 1.3.

Theorem 1.3. *There is $n_0 \in \mathbb{N}$ such that for all $n > n_0$, the following holds. Let G be an n -vertex graph of maximum degree Δ . Fix a partition $E(G) = X \sqcup U$ and let $\varphi: X \rightarrow [\Delta + 1]$ be a proper partial edge-coloring. Then there exists a subset $W \subseteq U$ of size $|W| \geq |U|/((\Delta + 1)^{10}(\log n)^2)$ such that it is possible to assign to each edge $e \in W$ a connected augmenting subgraph $H^e \subseteq G$ with the following properties:*

- for each $e \in W$, $E(H^e) \cap U = \{e\}$, that is, e is the unique uncolored edge in H^e ;
- the vertex sets of the graphs H^e , $e \in W$, are pairwise disjoint;
- for each $e \in W$, $|E(H^e)| \leq (\Delta + 1)^6 (\log n)^2$.

Proof. Throughout the proof, we will be assuming that n is sufficiently large. Let the endpoints of each uncolored edge e be x^e and y^e (it does not matter which endpoint we call x^e and which y^e). Set the value of the parameter ℓ to

$$\ell := (\Delta + 1)^6 \lfloor \log n \rfloor$$

and run the Multi-Step Vizing Algorithm with $e_0 = e$, $x_0 = x^e$, and $y_0 = y^e$ for every uncolored edge e , making the processes for different uncolored edges independent of each other. We use variables with “ e ” as a superscript, such as F_i^e and P_i^e , to denote the structures obtained using the execution of the algorithm with $e_0 = e$. In particular, for each $e \in U$, we have a sequence of vertices $x^e = x_0^e, x_1^e, x_2^e, \dots$.

Let $\lambda := \ell/(\Delta + 1)^3 = (\Delta + 1)^3 \lfloor \log n \rfloor$. Fix a small positive constant ε (any $\varepsilon < 1/12$ would work) and set $T := \lfloor \varepsilon \log n \rfloor$. Say that an uncolored edge e is **lucky** if the execution of the Multi-Step Vizing Algorithm with $e_0 = e$ successfully terminated on Step i for some $i < T$. Let $L \subseteq U$ denote the set of all lucky edges. For $e \in L$, we let C^e be the φ -happy chain with $\text{Start}(C^e) = e$ produced by the algorithm with $e_0 = e$. In other words, if the algorithm terminated on Step i due to Case 1, then $C^e = F_0^e + P_0^e + \dots + P_{i-1}^e + F_i^e$, and if it terminated due to Subcase 2.1, then $C^e = F_0^e + P_0^e + \dots + F_i^e + P_i^e$. If $e \in L$, then we let $H^e := H(C^e)$. By construction, H^e is a connected φ -augmenting subgraph whose only uncolored edge is e . Furthermore, since e is lucky,

$$|E(H^e)|, |V(H^e)| \leq \text{length}(C^e) + 1 \leq T(\ell + \Delta) + 1 \leq (\varepsilon + o(1))(\Delta + 1)^6 (\log n)^2. \quad (6.5)$$

Our goal now is to show that, with positive probability, there is a set $W \subseteq L$ of lucky edges of size at least $|W| \geq |U|/((\Delta + 1)^{10} (\log n)^2)$ such that the graphs H^e , $e \in W$, are pairwise vertex-disjoint. To this end, define an auxiliary graph Γ with $V(\Gamma) := U$ by making two distinct edges e, h adjacent in Γ if and only if e and h are both lucky and $V(H^e) \cap V(H^h) \neq \emptyset$.

Claim. For each uncolored edge e , $\mathbb{E}[\deg_\Gamma(e)] \leq (\varepsilon + o(1))(\Delta + 1)^{10} (\log n)^2$.

Proof. We will show that the desired upper bound holds even if we allow an adversary to specify the outcome of the algorithm for $e_0 = e$. So, fix the outcome of the algorithm’s execution for $e_0 = e$ and let $\mathbb{E}^*$ denote conditional expectation subject to this outcome. We may assume that e is lucky, since otherwise $\deg_\Gamma(e) = 0$. Then

$$\mathbb{E}^*[\deg_\Gamma(e)] \leq \sum_{u \in V(H^e)} \mathbb{E}^* \left[\left| \{h \in U \setminus \{e\} : h \in L \text{ and } u \in V(H^h)\} \right| \right]. \quad (6.6)$$

We may write “ $\mathbb{E}$ ” instead of “ $\mathbb{E}^*$ ” on the right-hand side of (6.6) because the construction of H^h for $h \neq e$ is independent from the construction of H^e . Fix an arbitrary vertex $u \in V(H^e)$. Observe that if $u \in V(H^h)$ for some h , then there is some i such that $u \in V(F_i^h)$ or $u \in V(P_i^h)$, so

$$\begin{aligned} \mathbb{E} [|\{h \in U \setminus \{e\} : h \in L \text{ and } u \in V(H^h)\}|] \\ \leq \sum_{i=0}^{\infty} \mathbb{E} [|\{h \in U : u \in V(F_i^h) \text{ or } u \in V(P_i^h)\}|]. \end{aligned}$$

If $u \in V(F_i^h)$ or $u \in V(P_i^h)$, then, by Lemma 5.3,

$$x_i^h \in R^-(u, \varphi), \quad x_{i-1}^h \in R^-(x_i^h, \varphi), \quad \dots, \quad x_1^h \in R^-(x_2^h, \varphi), \quad \text{and} \quad x^h = x_0^h \in R^-(x_1^h, \varphi). \quad (6.7)$$

Since for every $x \in V$, there are at most Δ edges h such that $x^h = x$, we conclude that, by Lemma 5.1, there are at most $\Delta(\Delta+1)^{3(i+1)} \leq (\Delta+1)^{3i+4}$ possible sequences $(h, x_0^h, x_1^h, \dots, x_i^h)$ satisfying (6.7). Since the vertices $x_1^h, \dots, x_i^h$ are chosen uniformly at random from sets of ℓ candidates, we obtain

$$\mathbb{E} [|\{h \in U : u \in V(F_i^h) \text{ or } u \in V(P_i^h)\}|] \leq \frac{(\Delta+1)^{3i+4}}{\ell^i} = \frac{(\Delta+1)^4}{\lambda^i}.$$

Thus we can write

$$\begin{aligned} \mathbb{E} [|\{h \in U \setminus \{e\} : h \in L \text{ and } u \in V(H^h)\}|] &\leq \sum_{i=0}^{\infty} \frac{(\Delta+1)^4}{\lambda^i} = \frac{\lambda}{\lambda-1} (\Delta+1)^4 \\ &= (1+o(1))(\Delta+1)^4. \end{aligned}$$

Plugging this into (6.6) and using (6.5) gives

$$\mathbb{E}^* [\deg_{\Gamma}(e)] \leq (1+o(1))(\Delta+1)^4 |V(H^e)| \leq (\varepsilon+o(1))(\Delta+1)^{10} (\log n)^2. \quad \dashv$$

Using the above claim and Markov’s inequality, we see that for each uncolored edge e ,

$$\mathbb{P} [\deg_{\Gamma}(e) > 2\varepsilon(\Delta+1)^{10} (\log n)^2 - 1] \leq \frac{1}{2} + o(1). \quad (6.8)$$

Let $W' \subseteq L$ be the set of all lucky edges e satisfying $\deg_{\Gamma}(e) \leq 2\varepsilon(\Delta+1)^{10} (\log n)^2 - 1$. Using Lemma 6.1 and (6.8), we conclude that for each $e \in U$,

$$\mathbb{P} [e \in W'] \geq 1 - \frac{n}{\lambda T} - \frac{3T(\Delta+1)^3}{\lambda-1} - \mathbb{P} [\deg_{\Gamma}(e) > 2\varepsilon(\Delta+1)^{10} (\log n)^2 - 1]$$

$$\begin{aligned}
&\geq 1 - n \left(\frac{1}{(\Delta + 1)^3 \lfloor \log n \rfloor} \right)^{\lfloor \varepsilon \log n \rfloor} - 3\varepsilon - \frac{1}{2} - o(1) \\
&= \frac{1}{2} - 3\varepsilon - o(1) > \frac{1}{4}.
\end{aligned}$$

Therefore, with positive probability, $|W'| \geq |U|/4$. Note that the maximum degree of the induced subgraph of Γ on W' is, by definition, at most $2\varepsilon(\Delta + 1)^{10}(\log n)^2 - 1$, which means that it contains an independent set $W \subseteq W'$ of size at least $|W| \geq |W'|/(2\varepsilon(\Delta + 1)^{10}(\log n)^2)$. Since W is an independent set in Γ , the subgraphs H^e for $e \in W$ are pairwise vertex-disjoint, and $|W| \geq |U|/(8\varepsilon(\Delta + 1)^{10}(\log n)^2) \geq |U|/((\Delta + 1)^{10}(\log n)^2)$ with positive probability. This completes the proof of Theorem 1.3. ■

References

- [1] N. Alon, L. Babai, A. Itai, A fast and simple randomized parallel algorithm for the maximal independent set problem, *J. Algorithms* 7 (4) (1986) 567–583.
- [2] L. Barenboim, M. Elkin, *Distributed Graph Coloring: Fundamentals and Recent Developments*, Synthesis Lectures on Distributed Computing Theory, 2013.
- [3] A. Bernshteyn, Distributed algorithms, the Lovász Local Lemma, and descriptive combinatorics, <https://arxiv.org/abs/2004.04905>, 2020 (preprint).
- [4] A. Bernshteyn, Probabilistic constructions in continuous combinatorics and a bridge to distributed algorithms, <https://arxiv.org/abs/2102.08797>, 2021 (preprint).
- [5] Y.-J. Chang, Q. He, W. Li, S. Pettie, J. Uitto, The complexity of distributed edge coloring with small palettes, in: *Proc. 29th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2018, pp. 2633–2652, full version: <https://arxiv.org/abs/1708.04290>.
- [6] R. Diestel, *Graph Theory*, 5th ed., Springer-Verlag, Berlin, Heidelberg, 2017.
- [7] M. Fischer, M. Ghaffari, F. Kuhn, Deterministic distributed edge-coloring via hypergraph maximal matching, in: *IEEE Symposium on Foundations of Computer Science (FOCS)*, 2017, pp. 180–191, full version: <https://arxiv.org/abs/1704.02767>.
- [8] M. Ghaffari, D. Harris, F. Kuhn, On derandomizing local distributed algorithms, in: *IEEE Symposium on Foundations of Computer Science (FOCS)*, 2018, pp. 662–673, full version: <https://arxiv.org/abs/1711.02194>.
- [9] M. Ghaffari, F. Kuhn, Y. Maus, J. Uitto, Deterministic distributed edge-coloring with fewer colors, in: *Proc. 50th Annual ACM SIGACT Symposium on Theory of Computing*, 2018, pp. 418–430, <https://arxiv.org/abs/1711.05469> (arXiv version).
- [10] A. Goldberg, S. Plotkin, G. Shannon, Parallel symmetry-breaking in sparse graphs, *SIAM J. Discrete Math.* 1 (4) (1988) 434–446.
- [11] J. Grebík, O. Pikhurko, Measurable versions of Vizing’s theorem, *Adv. Math.* 374 (18) (2020).
- [12] D.G. Harris, Distributed local approximation algorithms for maximum matching in graphs and hypergraphs, in: *IEEE Symposium on Foundations of Computer Science (FOCS)*, 2019, pp. 700–724, full version: <https://arxiv.org/abs/1807.07645>.
- [13] A.S. Kechris, A.S. Marks, Descriptive graph combinatorics, <http://math.caltech.edu/~kechris/papers/combinatorics20.pdf>, 2019 (preprint).
- [14] N. Linial, Locality in distributed graph algorithms, *SIAM J. Comput.* 21 (1) (1992) 193–201.
- [15] M. Luby, A simple parallel algorithm for the maximal independent set problem, *SIAM J. Comput.* 15 (4) (1986) 1036–1053.
- [16] O. Pikhurko, Borel combinatorics of locally finite graphs, <https://arxiv.org/abs/2009.09113>, 2020 (preprint).
- [17] V. Rozhoň, M. Ghaffari, Polylogarithmic-time deterministic network decomposition and distributed derandomization, in: *ACM Symposium on Theory of Computing (STOC)*, 2020, pp. 350–363, full version: <https://arxiv.org/abs/1907.10937>.
- [18] M. Stiebitz, D. Scheide, B. Toft, L.M. Favrholdt, *Graph Edge Coloring*, John Wiley & Sons, Hoboken, NJ, 2012.

- [19] H.-H. Su, H.T. Vu, Towards the locality of Vizing's theorem, in: Proc. 51st Annual ACM SIGACT Symposium on Theory of Computing, 2019, pp. 355–364, full version: <https://arxiv.org/abs/1901.00479>.
- [20] V.G. Vizing, Об оценке хроматического класса p -графа (Russian) (On an estimate of the chromatic class of a p -graph), *Diskretn. Anal.* 3 (1964) 25–30.