

An Exponential Separation Between Randomized and Deterministic Complexity in the LOCAL Model*

Yi-Jun Chang
University of Michigan

Tsvi Kopelowitz
University of Michigan

Seth Pettie
University of Michigan

Abstract

Over the past 30 years numerous algorithms have been designed for symmetry breaking problems in the LOCAL model, such as maximal matching, MIS, vertex coloring, and edge-coloring. For most problems the best randomized algorithm is at least exponentially faster than the best deterministic algorithm. In this paper we prove that these exponential gaps are *necessary* and establish numerous connections between the deterministic and randomized complexities in the LOCAL model. Each of our results has a very compelling take-away message:

Fast Δ -coloring of trees requires random bits. Building on the recent randomized lower bounds of Brandt et al. [11], we prove that the randomized complexity of Δ -coloring a tree with maximum degree Δ is $\Theta(\log_{\Delta} \log n)$, for any $\Delta \geq 55$, whereas its deterministic complexity is $\Theta(\log_{\Delta} n)$ for any $\Delta \geq 3$.¹ This also establishes a large separation between the deterministic complexity of Δ -coloring and $(\Delta + 1)$ -coloring trees.

Randomized lower bounds imply deterministic lower bounds. We prove that any deterministic algorithm for a natural class of problems that runs in $O(1) + o(\log_{\Delta} n)$ rounds can be transformed to run in $O(\log^* n - \log^* \Delta + 1)$ rounds. If the transformed algorithm violates a lower bound (even allowing randomization), then one can conclude that the problem requires $\Omega(\log_{\Delta} n)$ time deterministically. (This gives an alternate proof that deterministically Δ -coloring a tree with small Δ takes $\Omega(\log_{\Delta} n)$ rounds.)

Deterministic lower bounds imply randomized lower bounds. We prove that the randomized complexity of any natural problem on instances of size n is at least its deterministic complexity on instances of size $\sqrt{\log n}$. This shows that a *deterministic* $\Omega(\log_{\Delta} n)$ lower bound for any problem (Δ -coloring a tree, for example) implies a *randomized* $\Omega(\log_{\Delta} \log n)$ lower bound. It also illustrates that the *graph shattering* technique employed in recent randomized symmetry breaking algorithms is *absolutely essential* to the LOCAL model. For example, it is provably impossible to improve the $2^{O(\sqrt{\log \log n})}$ terms in the complexities of the best MIS and $(\Delta + 1)$ -coloring algorithms without *also* improving the $2^{O(\sqrt{\log n})}$ -round Panconesi-Srinivasan algorithms.

*This work is supported by NSF grants CCF-1217338, CNS-1318294, and CCF-1514383. Email: {cyijun,tsvi,pettie}@umich.edu

¹For simplicity, here we suppress any $\log^* n$ additive term in $\Theta(\cdot)$.

1 Introduction

One of the central problems of theoretical computer science is to determine the value of *random bits*. If the distinction is between computable vs. incomputable functions, random bits are provably useless in centralized models (Turing machines) [32]. However, this is not true in the distributed world! The celebrated Fischer-Lynch-Patterson theorem [15] states that asynchronous deterministic agreement is impossible with one unannounced failure, yet it is possible to accomplish with probability 1 using randomization. See Ben-Or [6] and [10, 30, 19]. There are also a number of basic symmetry breaking tasks that are trivially impossible to solve by identical, synchronized, deterministic processes, for example, medium access control to an Ethernet-like channel.

In this paper we examine the value of random bits in Linial’s [24] LOCAL model, which, for the sake of clarity, we bifurcate into two models RandLOCAL and DetLOCAL. In both models the input graph $G = (V, E)$ and communications network are identical. Each vertex hosts a processor and all vertices run the same algorithm. Each edge supports communication in both directions. The computation proceeds in synchronized *rounds*. In a round, each processor performs some computation and sends a message along each incident edge, which is delivered before the beginning of the next round. Each vertex v is initially aware of its degree $\deg(v)$ and certain global parameters such as $n \stackrel{\text{def}}{=} |V|$, $\Delta = \Delta(G) \stackrel{\text{def}}{=} \max_{v \in V} \deg(v)$, and possibly others.² In the LOCAL model the only measure of efficiency is the number of rounds. *All local computation is free and the size of messages is unbounded.* Henceforth “time” refers to the number of rounds.

DetLOCAL: In order to avoid trivial impossibilities, all vertices are assumed to hold unique $\Theta(\log n)$ -bit IDs. Except for the registers holding $\deg(v)$ and $\text{ID}(v)$, the initial state of v is identical to every other vertex. The algorithm executed at each vertex is deterministic.

RandLOCAL: In this model each vertex may locally generate an unbounded number of independent truly random bits. (There are no globally shared random bits.) Except for the register holding $\deg(v)$, the initial state of v is identical to every other vertex. Algorithms in this model operate for a specified number of rounds and have some probability of *failure*, the definition of which is problem specific. We usually only consider algorithms whose global probability of failure is at most $1/\text{poly}(n)$.

Observe that the lack of IDs in RandLOCAL is not a practical limitation. Before the first round each vertex can locally generate a random $\Theta(\log n)$ -bit ID, which is unique with probability $1 - 1/\text{poly}(n)$. For technical reasons it is convenient to assume that vertices are not initially differentiated by IDs.³

Early work in the LOCAL models suggested that randomness is of limited help. Naor [25] showed that Linial’s $\Omega(\log^* n)$ lower bound [24] for 3-coloring the ring holds even in RandLOCAL, and Naor and Stockmeyer [26] proved that the class of problems solvable by $O(1)$ -round algorithms is the same in RandLOCAL and DetLOCAL. However, in the intervening decades we have seen dozens of examples of symmetry breaking algorithms for RandLOCAL that are substantially faster than their counterparts in DetLOCAL; see [5] for an extensive survey or Table 1 for a glimpse at three archetypal problems: maximal independent set (MIS), maximal matching, and $(\Delta + 1)$ -coloring.

Graph Shattering. A little pattern matching in the bounds of Table 1 shows that the randomized symmetry breaking algorithms are exponentially faster *in two ways*. Their dependence on Δ is exponentially faster and their dependence on n is usually identical to the best deterministic complexity, but for $\text{poly}(\log n)$ -size instances, for example, $2^{O(\sqrt{\log n})}$ becomes $2^{O(\sqrt{\log \log n})}$. This second phenomenon is no coincidence! It is a direct result of the *graph shattering* approach to symmetry breaking used in [5] and further in [12, 14, 16, 18, 7, 21, 29]. The idea is to apply some randomized procedure that fixes some fragment of the output (e.g., part of the MIS is fixed, part of the coloring is fixed, etc.), thereby effectively removing a large fraction of the vertices from further consideration. If it can be shown that the connected components in the subgraph still under consideration have size $\text{poly}(\log n)$, one can revert to the best available *deterministic* algorithm and solve the problem on each component of the “shattered” graph in parallel.

²The assumption that global parameters are common knowledge can sometimes be removed; see Korman, Sereni, and Viennot [20].

³Notice that the role of “ n ” is different in the two LOCAL models: in DetLOCAL it only affects the ID length whereas in RandLOCAL it only affects the failure probability.

Problem	Model and Result	Citation
MIS:	DetLOCAL: $O\left(\min\left\{\Delta + \log^* n, 2^{O(\sqrt{\log n})}\right\}\right)$	[4, 28]
	RandLOCAL: $O\left(\log \Delta + 2^{O(\sqrt{\log \log n})}\right)$	[16]
	Lower Bound: $\Omega\left(\min\left\{\sqrt{\log n / \log \log n}, \log \Delta / \log \log \Delta + \log^* n\right\}\right)$	[22, 24, 25]
Maximal Matching:	DetLOCAL: $O\left(\min\left\{\Delta + \log^* n, \log^4 n\right\}\right)$	[27, 17]
	RandLOCAL: $O\left(\log \Delta + \log^4 \log n\right)$	[5]
	Lower Bound: $\Omega\left(\min\left\{\sqrt{\log n / \log \log n}, \log \Delta / \log \log \Delta + \log^* n\right\}\right)$	[22, 24, 25]
$(\Delta + 1)$-Coloring	DetLOCAL: $O\left(\min\left\{\Delta^{3/4} \log \Delta + \log^* n, 2^{O(\sqrt{\log n})}\right\}\right)$	[2, 28]
	RandLOCAL: $O\left(\sqrt{\log \Delta} + 2^{O(\sqrt{\log \log n})}\right)$	[18]
	Lower Bound: $\Omega(\log^* n)$	[24, 25]

Table 1: A sample of symmetry breaking results for three problems.

Lower Bounds in the LOCAL Model. Until recently, the main principle used to prove lower bounds in the LOCAL model was *indistinguishability*. The first application of this principle was by Linial [24] himself, who argued that any algorithm for coloring degree- Δ trees either uses $\Omega(\Delta / \log \Delta)$ colors or takes $\Omega(\log_{\Delta} n)$ time. The proof is as follows (i) in $o(\log_{\Delta} n)$ time, a vertex *cannot always distinguish* whether the input graph G is a tree or a graph with girth $\Omega(\log_{\Delta} n)$, (ii) for all Δ and all n , there exists a degree- Δ graph with girth $\Omega(\log_{\Delta} n)$ and chromatic number $\chi = \Omega(\Delta / \log \Delta)$, hence⁴ (iii) any $o(\log_{\Delta} n)$ -time algorithm for coloring trees could also color such a graph, and therefore must use at least χ colors. A significantly more subtle indistinguishability argument was used by Kuhn, Moscibroda, and Wattenhofer [22], who showed that $O(1)$ -approximate vertex cover, maximal matching, MIS, and several other problems have $\Omega(\min\{\log \Delta / \log \log \Delta, \sqrt{\log n / \log \log n}\})$ lower bounds⁵. By its nature, indistinguishability is not very good at separating randomized and deterministic complexities. Very recently, Brandt et al. [11] developed a lower bound technique that explicitly incorporates error probabilities and proved that several problems on graphs with constant Δ take $\Omega(\log \log n)$ time in RandLOCAL (with error probability $1/\text{poly}(n)$) such as *sinkless orientation*, *sinkless coloring*, and Δ -coloring. Refer to Section 2 for definitions of these problems. Since the existence of a sinkless orientation can be proved with the Lovász local lemma (LLL), this gave $\Omega(\log \log n)$ lower bounds on distributed algorithms for the constructive LLL. See [12, 16] for upper bounds on the distributed LLL.

1.1 New Results

In this paper we exhibit an exponential separation between RandLOCAL and DetLOCAL for several *specific* symmetry breaking problems. More generally, we give new connections between the randomized and deterministic complexities of all *locally checkable labeling* problems (refer to Section 2 for a definition of LCLs), a class that includes essentially any natural symmetry breaking problem.

Separation of RandLOCAL and DetLOCAL. We extend Brandt et al.’s [11] randomized lower bound as follows: on degree- Δ graphs, Δ -coloring takes $\Omega(\log_{\Delta} \log n)$ time in RandLOCAL and $\Omega(\log_{\Delta} n)$ time in DetLOCAL. The hard graphs in this lower bound have girth $\Omega(\log_{\Delta} n)$, so by the indistinguishability principle, these lower bounds also apply to Δ -coloring trees. On the upper bound side, Barenboim and Elkin [3] showed that for $\Delta \geq 3$, Δ -coloring trees takes $O(\log_{\Delta} n + \log^* n)$ time in DetLOCAL. We give an elementary proof that for $\Delta \geq 55$, Δ -coloring trees can be done in $O(\log_{\Delta} \log n + \log^* n)$ time in RandLOCAL, matching Brandt et al.’s [11] lower bound up to the $\log^* n$. A more complicated algorithm for Δ -coloring trees could be derived from [29], for $\Delta > \Delta_0$ and some very large constant

⁴Linial [24] actually only used the existence of Δ -regular graphs with high girth and chromatic number $\Omega(\sqrt{\Delta})$. See [8] for constructions with chromatic number $\Omega(\Delta / \log \Delta)$.

⁵In [23] the same authors argue that these problems have a lower bound of $\Omega(\min\{\log \Delta, \sqrt{\log n}\})$. However, recently Bar-Yehuda, Censor-Hillel, and Schwartzman [1] pointed out an error in their proof.

Δ_0 .⁶

Randomized lower bounds imply deterministic lower bounds. We give a second, more generic proof that Δ -coloring trees takes $\Omega(\log_\Delta n)$ time. The proof shows that any $f(\Delta) + o(\log_\Delta n)$ time algorithm for an LCL problem can be transformed in a black box way to run in $O((1 + f(\Delta))(\log^* n - \log^* \Delta + 1))$ time. Thus, on bounded-degree graphs, there are no “natural” deterministic time bounds between $\omega(\log^* n)$ and $o(\log n)$. Any $\omega(\log^* n)$ lower bound for bounded degree graphs (in either RandLOCAL or DetLOCAL) *immediately* implies an $\Omega(\log n)$ lower bound in DetLOCAL. This reduction can be parameterized in many different ways. Under a different parametrization it shows that any $O(\log^{1-\frac{1}{k+1}} n)$ -time DetLOCAL algorithm for an LCL problem can be transformed to run in $O(\log^k \Delta (\log^* n - \log^* \Delta + 1))$ -time. For example, if one were to develop a deterministic $O(\sqrt{\log n})$ -time MIS or maximal matching algorithm—almost matching one of the KMW [22] lower bounds—it would immediately imply an $O(\log \Delta (\log^* n - \log^* \Delta + 1))$ -time MIS/maximal matching algorithm, which almost matches the *other* KMW lower bound, for $\Delta > \log^{O(1)} n$. By some strange coincidence, [5] gave an analogous reduction for MIS/maximal matching in bounded arboricity graphs, but for RandLOCAL and in the *reverse direction*. Specifically, any $O(\log^k \Delta + f(n))$ -time RandLOCAL MIS/maximal matching algorithm can be transformed into an $O(\log^{1-\frac{1}{k+1}} n + f(n))$ -time RandLOCAL algorithm for bounded arboricity graphs.

Deterministic lower bounds imply randomized lower bounds. We prove that for any LCL problem, its RandLOCAL complexity on instances of size n is at least its DetLOCAL complexity on instances of size $\sqrt{\log n}$. This *reverses* the implication proved above. For example, if we begin with a proof that Δ -coloring takes $\Omega(\log_\Delta n)$ time in DetLOCAL, then we conclude that it must take $\Omega(\log_\Delta \log n)$ time in RandLOCAL. This result has a very clear take-away message: the *graph shattering* technique applied by recent randomized symmetry breaking algorithms [5, 16] is *inherent* to the RandLOCAL model and every optimal RandLOCAL algorithm for instances of size n must, in some way, encode an optimal DetLOCAL algorithm on $\text{poly}(\log n)$ -size instances. It is therefore impossible to improve the $2^{O(\sqrt{\log \log n})}$ terms in the RandLOCAL MIS and coloring algorithms of [5, 16, 18, 14] without also improving the $2^{O(\sqrt{\log n})}$ -time DetLOCAL algorithms of Panconesi and Srinivasan [28], and it is impossible to improve the $O(\log^4 \log n)$ term in the RandLOCAL maximal matching algorithm of [5] without also improving the $O(\log^4 n)$ DetLOCAL maximal matching algorithm of [17].

2 Preliminaries

Graph Notation. For a graph $G = (V, E)$ and for $u, v \in V$, let $\text{dist}_G(v, u)$ be the distance between v and u in G . Let $N(v) = \{u \mid (v, u) \in E\}$ be the neighborhood of v and let $N^r(v) = \{u \mid \text{dist}_G(v, u) \leq r\}$ be the set of all vertices within distance r of v .

Locally Checkable Labeling. The class of *Locally Checkable Labeling* (LCL) [26] problems are intuitively those graph problems whose solutions can be verified in $O(1)$ rounds, given a suitable labeling of the graph. Formally, an LCL problem is defined by a fixed radius r , a finite set Σ of vertex labels, and a set \mathcal{C} of acceptable labeled subgraphs. For any legal solution I to the problem there is a labeling $\lambda_I : V \rightarrow \Sigma$ that encodes I (plus possibly other information) such that for each $v \in V$, the labeled subgraph induced by $N^r(v)$ lies in \mathcal{C} . Moreover, for any *non-solution* I' to the problem, there is no labeling $\lambda_{I'}$ with this property. The following symmetry breaking problems are LCLs for $r = 1$.

- MAXIMAL INDEPENDENT SET (MIS). Given a graph $G = (V, E)$, find a set $I \subseteq V$ such that for any vertex $v \in V$, we have $N(v) \cap I = \emptyset$ iff $v \in I$.
- k -COLORING. Given a graph $G = (V, E)$, find an assignment $V \rightarrow \{1, 2, \dots, k\}$ such that for each edge $\{u, v\} \in E$, u and v are assigned to different numbers (also called colors).

⁶The reason we are interested in minimizing the $\Delta_0 \leq \Delta$ for which the algorithm works is somewhat technical. It seems as if Δ -coloring trees is a problem whose character makes a *qualitative* transition when Δ is a small enough constant. Using our technique (graph shattering) we may be able to replace 55 with a smaller constant, *but not too small*. Any algorithm that 3-colors 3-regular trees, for example, will need to be qualitatively very different in its design.

For MIS it suffices to label vertices with $\Sigma = \{0, 1\}$ indicating whether they are in the MIS. For k -Coloring we use $\Sigma = \{1, \dots, k\}$. The definition of LCLs is easily generalized to the case where the input graph G is supplemented with some labeling (e.g., an edge-coloring) or where λ labels both vertices and edges. Brandt et al. [11] considered the following problems.

- Δ -SINKLESS COLORING. Given a Δ -regular graph $G = (V, E)$ and a proper Δ -edge coloring of E using colors in $\{1, 2, \dots, \Delta\}$, find a Δ -coloring of V using colors in $\{1, 2, \dots, \Delta\}$ such that there is no edge $\{u, v\} \in E$ for which u, v and $\{u, v\}$ all have the same color.
- Δ -SINKLESS ORIENTATION. Given a Δ -regular graph $G = (V, E)$ and a proper Δ -edge coloring of E , find an orientation of the edges such that all edges have out-degree ≥ 1 .

Observe that both Δ -Sinkless Coloring and Δ -Sinkless Orientation are LCL graph problems with $r = 1$. For Sinkless Orientation $\Sigma = \{\rightarrow, \leftarrow\}^\Delta$ encodes the directions of all edges incident to a vertex, and the radius $r = 1$ is necessary and sufficient to verify that the orientations declared by both endpoints of an edge are consistent.

Linial's coloring. In the DetLOCAL model the initial $\Theta(\log n)$ -bit IDs can be viewed as an $n^{O(1)}$ -coloring of the graph. Our algorithms make frequent use of Linial's [24] coloring algorithm, which recolors the vertices using a smaller palette.

Theorem 1 ([24]). *Let G be a graph which has been k -colored. Then it is possible to deterministically re-color G using $5\Delta^2 \log k$ colors in one round.*

Theorem 2 ([24]). *There exists a universal constant $\beta > 0$ such that there is a DetLOCAL algorithm that computes a $\beta \cdot \Delta^2$ -coloring of a graph in $\mathcal{O}(\log^* n - \log^* \Delta + 1)$ time.*

3 The Necessity of Graph Shattering

Theorem 3 establishes that the graph shattering technique [5] is optimal and unavoidable in RandLOCAL. In particular, the randomized complexity of any symmetry breaking problem always hinges on its deterministic complexity.

Theorem 3. *Let \mathcal{P} be an LCL problem. Define $\text{Det}_{\mathcal{P}}(n, \Delta)$ to be the complexity of the optimal deterministic algorithm for \mathcal{P} in the DetLOCAL model and define $\text{Rand}_{\mathcal{P}}(n, \Delta)$ to be its complexity in the RandLOCAL model, with global error probability $1/n$. Then*

$$\text{Det}_{\mathcal{P}}(n, \Delta) \leq \text{Rand}_{\mathcal{P}}(2^{n^2}, \Delta).$$

Proof. Let $\mathcal{A}_{\text{Rand}}$ be a randomized algorithm for \mathcal{P} . Each vertex running $\mathcal{A}_{\text{Rand}}$ generates a string of $r(n, \Delta)$ random bits and proceeds for $t(n, \Delta)$ rounds, where r and t are two arbitrary functions. The probability that the algorithm fails in any way is at most $1/n$. Our goal is to convert $\mathcal{A}_{\text{Rand}}$ into a deterministic algorithm \mathcal{A}_{Det} in the DetLOCAL model. Let $G = (V, E)$ be the network on which \mathcal{A}_{Det} runs. Initially each $v \in V$ knows $n = |V|$, Δ , and a unique $\text{ID}(v) \in \{0, 1\}^{c \log n}$. Let $\mathcal{G}_{n, \Delta}$ be the set of all n -vertex graphs with unique vertex IDs in $\{0, 1\}^{c \log n}$ and maximum degree at most Δ . Regardless of Δ ,

$$|\mathcal{G}_{n, \Delta}| \leq 2^{\binom{n}{2} + cn \log n} \ll 2^{n^2} \stackrel{\text{def}}{=} N$$

Imagine simulating $\mathcal{A}_{\text{Rand}}$ on a graph $G' \in \mathcal{G}_{n, \Delta}$ whose vertices are given input parameters (N, Δ) , that is, we imagine G' is disconnected from the remaining $N - n$ vertices. The probability that $\mathcal{A}_{\text{Rand}}$ fails on an N -vertex graph is at most $1/N$, so the probability that any vertex in G' witnesses a failure is also certainly at most $1/N$.

Suppose we select a function $\phi : \{0, 1\}^{c \log n} \rightarrow \{0, 1\}^{r(N, \Delta)}$ uniformly at random from the space of all such functions. Define $\mathcal{A}_{\text{Det}}[\phi]$ to be the *deterministic* algorithm that simulates $\mathcal{A}_{\text{Rand}}$ for $t(N, \Delta)$ steps,

where the string of random bits generated by v is fixed to be $\phi(\text{ID}(v))$. We shall call ϕ a *bad* function if $\mathcal{A}_{\text{Det}}[\phi]$ fails to compute the correct answer on some member of $\mathcal{G}_{n,\Delta}$. By the union bound,

$$\begin{aligned} \Pr_{\phi}(\phi \text{ is bad}) &\leq \sum_{G' \in \mathcal{G}_{n,\Delta}} \Pr_{\phi}(\mathcal{A}_{\text{Det}}[\phi] \text{ errs on } G') \\ &= \sum_{G' \in \mathcal{G}_{n,\Delta}} \Pr(\mathcal{A}_{\text{Rand}} \text{ errs on } G', \text{ with input parameters } (N, \Delta)) \\ &\leq |\mathcal{G}_{n,\Delta}|/N < 1. \end{aligned}$$

Thus, there exists some good ϕ . Any ϕ can be encoded as a long bit-string $\langle \phi \rangle \stackrel{\text{def}}{=} \phi(0)\phi(1)\cdots\phi(2^{c \log n} - 1)$. Define ϕ^* to be the good function for which $\langle \phi^* \rangle$ is lexicographically first.

The algorithm \mathcal{A}_{Det} is as follows. Each vertex v , given input parameters (n, Δ) , first computes $N = 2^{n^2}, t(N, \Delta), r(N, \Delta)$, then performs the simulations of $\mathcal{A}_{\text{Rand}}$ necessary to compute ϕ^* . Once ϕ^* is computed it executes $\mathcal{A}_{\text{Det}}[\phi^*]$ for $t(N, \Delta)$ rounds. By definition, $\mathcal{A}_{\text{Det}}[\phi^*]$ never errs when run on any member of $\mathcal{G}_{n,\Delta}$. \square

Remark 1. *Theorem 3 works equally well when t and r are functions of n, Δ , and possibly other quantitative global graph parameters. For example, the time may depend on measures of local sparsity (as in [14]), arboricity/degeneracy (as in [3, 5]), or neighborhood growth (as in [31]).*

Naor and Stockmeyer [26] proved that the class of truly local ($O(1)$ -time) problems in RandLOCAL and DetLOCAL is identical for bounded Δ . Theorem 3 implies something slightly stronger, since $\log^* n$ and $\log^*(\sqrt{\log n})$ differ by a constant.

Corollary 1. *Any RandLOCAL algorithm for an LCL running in $t(n) = 2^{O(\log^* n)}$ time can be derandomized without asymptotic penalty. The corresponding DetLOCAL algorithm runs in $O(t(n))$ time.*

4 Lower bounds for Δ -coloring Δ -regular Trees

In this section we prove that on degree- Δ graphs with girth $\Omega(\log_{\Delta} n)$, Δ -coloring takes $\Omega(\log_{\Delta} \log n)$ time in RandLOCAL and $\Omega(\log_{\Delta} n)$ time in DetLOCAL. Since the girth of the graphs used to prove these lower bounds is $\Omega(\log_{\Delta} n)$, by the indistinguishability principle they also apply to the problem of Δ -coloring trees.

Sinkless coloring and sinkless orientations. Brandt et al. [11] proved $\Omega(\log \log n)$ lower bounds on RandLOCAL algorithms, that have a $1/\text{poly}(n)$ probability of failure, for sinkless coloring and sinkless orientation of 3-regular graphs. We say that a sinkless coloring algorithm \mathcal{A} has failure probability p if, for *each* individual edge $e = \{u, v\}$, the probability that $\text{color}(u) = \text{color}(v) = \text{color}(\{u, v\})$ is at most p . Thus, by the union bound, the *global* probability of failure is at most $p|E|$. We say a that sinkless orientation algorithm \mathcal{A} has failure probability p if, for each $v \in V$, the probability that v is a sink is at most p . We say that monochromatic edges and sinks are *forbidden configurations* for sinkless coloring and sinkless orientation, respectively.

The following two lemmas are proven in [11] for $\Delta = 3$. It is straightforward to go through the details of the proof and track the dependence on Δ .

Lemma 1 ([11]). *Let $G = (V, E, \psi)$ be a Δ -regular graph with girth g that is equipped with a proper Δ -edge coloring ψ . Suppose that there is a RandLOCAL algorithm \mathcal{A} for Δ -sinkless coloring taking $t < \frac{g-1}{2}$ rounds such that $\forall e \in E$, \mathcal{A} outputs a forbidden configuration at e with probability at most p . Then there is a RandLOCAL algorithm \mathcal{A}' for Δ -sinkless orientation taking t rounds such that $\forall v \in V$, \mathcal{A}' outputs a forbidden configuration at v with probability at most $2\Delta p^{1/3}$.*

Lemma 2 ([11]). *Let $G = (V, E, \psi)$ be a Δ -regular graph with girth g that is equipped with a proper Δ -edge coloring ψ . Suppose that there is a RandLOCAL algorithm \mathcal{A}' for sinkless orientation taking $t < \frac{g-1}{2}$ rounds such that $\forall v \in V$, \mathcal{A}' outputs a forbidden configuration at v with probability at most p . Then there is a RandLOCAL algorithm \mathcal{A} for Δ -sinkless coloring taking $t-1$ rounds such that $\forall e \in E$, \mathcal{A} outputs a forbidden configuration at e with probability at most $4p^{1/(\Delta+1)}$.*

The following theorem generalizes Corollary 25 in [11] to allow non-constant Δ and arbitrary failure probability p .

Theorem 4. *Any RandLOCAL algorithm for Δ -coloring a graph with degree at most Δ and error probability p takes at least $t = \min\{\epsilon \log_{3(\Delta+1)} \ln(1/p), \epsilon \log_{\Delta} n\} \geq 1$ rounds for a sufficiently small $\epsilon > 0$.*

Proof. For any $\Delta \geq 3$ there exist a bipartite Δ -regular graphs with girth $\Omega(\log_{\Delta} n)$; see [13, 9]. Such graphs are trivially Δ -edge colorable. Moreover, any Δ -coloring of such a graph is also a valid Δ -sinkless coloring. Applying Lemmas 1 and 2 we conclude that any t -round Δ -sinkless coloring algorithm with error probability p can be transformed into a $(t-1)$ -round Δ -sinkless coloring algorithm with error probability $4(2\Delta)^{\frac{1}{\Delta+1}} p^{\frac{1}{3(\Delta+1)}} < 7p^{\frac{1}{3(\Delta+1)}}$. Iterating this process t times, it follows that there exists a 0-round Δ -sinkless coloring algorithm with (local) error probability $O(p^{(\frac{1}{3(\Delta+1)})^t})$. Note that

$$p^{(\frac{1}{3(\Delta+1)})^t} \geq p^{(\frac{1}{3(\Delta+1)})^{\epsilon \log_{3(\Delta+1)} \ln(1/p)}} = p^{(\ln p)^{-\epsilon}} = \exp(-(\ln(1/p))^{1-\epsilon})$$

Because the graph is Δ -regular and the vertices undifferentiated by IDs, any 0-round RandLOCAL algorithm colors each vertex independently according to the same distribution. The probability that any vertex is involved in a forbidden configuration (a monochromatic edge) is therefore at least $1/\Delta^2$. Since $\epsilon \log_{3(\Delta+1)} \ln(1/p) \geq 1$ we have $\Delta < \ln(1/p)$, but

$$\frac{1}{\Delta^2} \geq \exp(-2 \ln \ln(1/p)) \gg \exp(-(\ln(1/p))^{1-\epsilon}).$$

Thus, there is no RandLOCAL Δ -sinkless coloring algorithm that takes t -rounds and errs with probability p . \square

Corollary 2 is an immediate consequence of Theorem 4.

Corollary 2. *Any RandLOCAL algorithm for Δ -coloring a graph with global error probability $1/\text{poly}(n)$ takes $\Omega(\log_{\Delta} \log n)$ time.*

Theorem 4 does not immediately extend to DetLOCAL. Recall that in the DetLOCAL model vertices are initially endowed with $O(\log n)$ -bit IDs whereas in RandLOCAL they are undifferentiated.

Theorem 5. *Any DetLOCAL algorithm that Δ -colors degree- Δ graphs with girth $\Omega(\log_{\Delta} n)$ or degree- Δ trees requires $\Omega(\log_{\Delta} n)$ time.*

Proof. Let \mathcal{A}_{Det} be a DetLOCAL algorithm that Δ -colors a graph in $t = t(n, \Delta)$ rounds and G be the input graph. We construct a RandLOCAL algorithm $\mathcal{A}_{\text{Rand}}$ taking $O(t)$ rounds as follows. Before the first round each vertex locally generates a random n -bit ID. Assume for the time being that these IDs are unique, and therefore constitute a 2^n -coloring of G . Let $G' = (V, \{(u, v) \mid \text{dist}_G(u, v) \leq 2t + 1\})$. The maximum degree Δ' in G' is clearly less than n . We apply one step of Linial's recoloring algorithm (Theorem 1) to G' and obtain a coloring with palette size $O(\Delta'^2 \log(2^n)) = O(n^3)$. A step of Linial's algorithm in G' is simulated in G using $O(t)$ time. Using these colors as $(3 \log n + O(1))$ -bit IDs, we simulate \mathcal{A}_{Det} in G for t steps. Since no vertex can see two vertices with the same ID, this algorithm necessarily behaves as if all IDs are unique. Observe that because \mathcal{A}_{Det} is deterministic, the only way $\mathcal{A}_{\text{Rand}}$ can err is if the initial n -bit IDs fail to be unique. This occurs with probability $p < n^2/2^n$. By Theorem 4 $\mathcal{A}_{\text{Rand}}$ takes $\Omega(\min\{\log_{\Delta} \log(1/p), \log_{\Delta} n\}) = \Omega(\log_{\Delta} n)$ time. \square

5 Gaps in Deterministic Time Complexity

The Time Hierarchy Theorem informally says that a Turing machine can solve more problems given more time. A similar question can be asked in the setting of distributed computation. For example, does increasing the number of rounds from $\Theta(\log^* n)$ to $\Theta(\log \log n)$ allow one to solve more problems? In this section, we will demonstrate a general technique that allows one to speedup deterministic algorithms in the DetLOCAL model. Based on this technique, we demonstrate the existence of a ‘‘gap’’ in possible DetLOCAL complexities.

A graph class is *hereditary* if it is closed under removing vertices and edges. Examples of hereditary graph classes are general graphs, forests, bounded arboricity graphs, triangle-free graphs, planar graphs, and chordal graphs. We prove that for graphs with constant Δ the time complexity of *any* LCL problem on any graph from a *hereditary* graph class is either $\Omega(\log n)$ or $O(\log^* n)$.

Theorem 6. *Let \mathcal{P} be an LCL graph problem with parameters r , Σ , and \mathcal{C} , and let \mathcal{A} be a DetLOCAL algorithm for solving \mathcal{P} . Let β be the universal constant from Theorem 2. Suppose that the cost of \mathcal{A} on instances of \mathcal{P} with n vertices, where the instance is taken from a hereditary graph class, is at most $f(\Delta) + \epsilon \log_\Delta n$ time, where $f(\Delta) \geq 0$ and $\epsilon = \frac{1}{4+4\log\beta+4r}$ is a constant. Then there exists a DetLOCAL algorithm \mathcal{A}' that solves \mathcal{P} on the same instances in $O((1+f(\Delta))(\log^* n - \log^* \Delta + 1))$ time.*

Proof. Notice that for any instance of \mathcal{P} with n vertices and ID length ℓ , it must be that $\ell \geq \log n$ and so the running time of \mathcal{A} on such instances is bounded by $T(\Delta, \ell) \leq f(\Delta) + \frac{\epsilon \ell}{\log \Delta}$.

Let $G = (V, E)$ be an instance of \mathcal{P} . The algorithm \mathcal{A}' on G works as follows. Let $\tau = 1 + \log \beta$ be a constant. We use Linial's coloring technique to produce short IDs of length ℓ' that are distinct within distance $4f(\Delta) + 2\tau + 2r$. Let $G' = (V, E')$ be the graph $E' = \{\{u, v\} \mid u, v \in V \text{ and } \text{dist}_G(u, v) \leq 4f(\Delta) + 2\tau + 2r\}$. The max degree in G' is at most $\Delta^{4f(\Delta)+2\tau+2r}$ since $\max_{u \in V} \{|N^{4f(\Delta)+2\tau+2r}(u)|\} \leq \Delta^{4f(\Delta)+2\tau+2r}$. Each vertex $u \in V$ simulates G' by collecting $N^{4f(\Delta)+2\tau+2r}(u)$ in $O(f(\Delta) + \tau + r)$ time.

We simulate the algorithm of Theorem 2 on G' by treating each of the ℓ' bit IDs of vertices in V as a color. This produces a $\beta \cdot \Delta^{8f(\Delta)+4\tau+4r}$ -coloring, which is equivalent to identifiers of length $\ell' = (8f(\Delta) + 4\tau + 4r) \log \Delta + \log \beta$. Although these identifiers are not globally unique, they are distinct in $N^{2f(\Delta)+\tau+r}(u)$ for each vertex $u \in V$. The time complexity of this process is

$$(4f(\Delta) + 2\tau + 2r) \cdot O(\log^* n - \log^* \Delta + 1).$$

Finally, we apply \mathcal{A} on G while implicitly assuming that the graph size is $2^{\ell'}$ and using the shorter IDs. The runtime of this execution of \mathcal{A} is:

$$\begin{aligned} f(\Delta) + \frac{\epsilon \ell'}{\log \Delta} &= f(\Delta) + \frac{\epsilon((8f(\Delta) + 4\tau + 4r) \log \Delta + \log \beta)}{\log \Delta} \\ &= (1 + 8\epsilon)f(\Delta) + 1 + \frac{\epsilon \log \beta}{\log \Delta} && \epsilon(4\tau + 4r) = 1 \\ &\leq (1 + 8\epsilon)f(\Delta) + \tau && \log \Delta \geq 1, \epsilon < 1 \\ &\leq 2f(\Delta) + \tau. && 8\epsilon = \frac{2}{\tau + r} \leq 1 \end{aligned}$$

Whether the output labeling of $u \in V$ is legal depends on the labeling of the vertices in $N^r(u)$, which depends on the graph structure and the IDs in $N^{2f(\Delta)+\tau+r}(u)$. Due to the hereditary property of the graph class under consideration, for each $u \in V$, $N^{2f(\Delta)+\tau+r}(u)$ is isomorphic to a subgraph of some $2^{\ell'}$ -vertex graph in the same class. Moreover, the shortened ID in $N^{2f(\Delta)+\tau+r}(u)$ are distinct. Therefore, it is guaranteed that the output of the simulation is a legal labeling.

The total time complexity is

$$(4f(\Delta) + 2\tau + 2r) \cdot O(\log^* n - \log^* \Delta + 1) + 2f(\Delta) + \tau = O((1 + f(\Delta))(\log^* n - \log^* \Delta + 1)).$$

□

Combining Theorem 6 with Corollary 2 and setting $f(\Delta) = O(1)$ provides a new proof of Theorem 5 for small enough Δ . To see this, notice that any lower bound for the RandLOCAL model with error probability $1/\text{poly}(n)$ can be adapted to DetLOCAL since we can randomly pick $O(\log n)$ -bit IDs that are distinct with probability $1 - 1/\text{poly}(n)$. From Corollary 2 any DetLOCAL algorithm that Δ -colors a degree- Δ tree requires $\Omega(\log_\Delta \log n)$ time. However, Theorem 6 states that any DetLOCAL algorithm running in $O(1) + o(\log_\Delta n)$ time can be speeded up to run in $O(\log^* n - \log^* \Delta + 1)$ time. This contradicts the lower bound whenever $\log_\Delta \log n \gg \log^* n - \log^* \Delta + 1$. Hence Δ -coloring a degree- Δ tree takes $\Omega(\log_\Delta n)$ time in DetLOCAL for small enough Δ such that $\log_\Delta \log n \gg \log^* n - \log^* \Delta + 1$.

Another consequence of Theorem 2 is that the deterministic time complexity of a problem can either be solved very efficiently (i.e. in $O((1 + f(\Delta))(\log^* n - \log^* \Delta + 1))$ time) or requires $\Omega(f(\Delta) + \log_{\Delta} n)$ time (which is at least the order of the diameter when the underlying graph is a complete regular tree). Such a consequence is the strongest when Δ is small. For example, if Δ is a constant, Theorem 2 implies the following corollary:

Corollary 3. *The time complexity of any LCL problem on any hereditary graph class that has constant Δ in the DetLOCAL model is either $\Omega(\log n)$ or $O(\log^* n)$.*

A simple adaptation of the proof of Theorem 6 shows an even stronger dichotomy when $\Delta = 2$.

Theorem 7. *The DetLOCAL time complexity of any LCL problem on any hereditary graph class with $\Delta = 2$ is either $\Omega(n)$ or $O(\log^* n)$.*

We remark that an interpretation of the time complexity requirement in Theorems 6 and 7 is that the diameter of a graph with maximum degree Δ is at least $\Omega(\log_{\Delta} n)$ for $\Delta \geq 3$ and $\Omega(n)$ when $\Delta = 2$. If we allow the possibility for an algorithm to see the entire graph, then the algorithm can solve the problem globally.

Given a $O(\sqrt{\log n})$ -time deterministic algorithm, one may feel that it is possible to use Theorem 6 to improve the time complexity to $O(\log^* n)$ since $\sqrt{\log n} = o(\log_{\Delta} n)$ for the case $\Delta = \exp(o(\sqrt{\log n}))$. However, the class of graphs with $\Delta = \exp(o(\sqrt{\log n}))$ is not hereditary, and so Theorem 6 does not apply. Nonetheless, Linial's coloring technique can be made to speed up algorithms with time complexity of the form $f(\Delta) + g(n)$.

Theorem 8. *Let \mathcal{P} be an LCL graph problem with parameters r , Σ , and \mathcal{C} , and let \mathcal{A} be a DetLOCAL algorithm for solving \mathcal{P} . Suppose that the runtime of the algorithm \mathcal{A} on instances of \mathcal{P} from a hereditary graph class is $O(\log^k \Delta + \log^{\frac{k}{k+1}} n)$. Then there exists a deterministic algorithm \mathcal{A}' that solves \mathcal{P} on the same instances in $O(\log^k \Delta (\log^* n - \log^* \Delta + 1))$ time.*

Proof. Notice that for any instance of \mathcal{P} with n vertices and ID length ℓ , it must be that $\ell \geq \log n$ and so the running time of \mathcal{A} on such instances is bounded by $\epsilon_1 \log^k \Delta + \epsilon_2 \ell^{\frac{k}{k+1}}$, for some constant ϵ_1, ϵ_2 .

We set $\tau = \epsilon \log^k \Delta$, with the parameter ϵ to be determined. Similar to the proof of Theorem 6, the algorithm \mathcal{A}' first produces shortened ID that are distinct for vertices within distance $2\tau + 2r$, and then simulates \mathcal{A} on the shortened ID in τ rounds.

Let $G' = (V, E')$ be the graph $E' = \{\{u, v\} | u, v \in V \text{ and } \text{dist}_G(u, v) \leq 2\tau + 2r\}$. The maximum degree in G' is at most $\Delta^{2\tau+2r}$. Each vertex $u \in V$ simulates G' by collecting $N^{2\tau+2r}(u)$ in $O(\tau + r)$ time.

We simulate the algorithm of Theorem 2 on G' by treating each of the ℓ bit IDs of vertices in V as a color. This produces a $\beta \cdot \Delta^{4\tau+4r}$ -coloring, which is equivalent to identifiers of length $\ell' = (4\tau + 4r) \log \Delta + \log \beta$. Although these identifiers are not globally unique, they are distinct in $N^{\tau+r}(u)$ for each vertex $u \in V$. The time complexity of this process is

$$(2\tau + 2r) \cdot O(\log^* n - \log^* \Delta + 1).$$

Finally, we apply \mathcal{A} on G while implicitly assuming that the graph size is $2^{\ell'}$ and using the shorter IDs. By setting ϵ as a large enough number such that $\epsilon_1 + \epsilon_2 (4(\epsilon + r + \log \beta))^{\frac{k}{k+1}} \leq \epsilon$, the runtime of this execution of \mathcal{A} is:

$$\begin{aligned} \epsilon_1 \log^k \Delta + \epsilon_2 (\ell')^{\frac{k}{k+1}} &= \epsilon_1 \log^k \Delta + \epsilon_2 ((4\tau + 4r) \log \Delta + \log \beta)^{\frac{k}{k+1}} \\ &\leq \epsilon_1 \log^k \Delta + \epsilon_2 \left(4(\epsilon \log^k \Delta + r + \log \beta) \log \Delta\right)^{\frac{k}{k+1}} \\ &\leq \epsilon_1 \log^k \Delta + \epsilon_2 \left(4(\epsilon + r + \log \beta) \log^{k+1} \Delta\right)^{\frac{k}{k+1}} \\ &= \left(\epsilon_1 + \epsilon_2 (4(\epsilon + r + \log \beta))^{\frac{k}{k+1}}\right) \log^k \Delta \\ &\leq \epsilon \log^k \Delta \\ &= \tau \end{aligned}$$

Whether the output labeling of $u \in V$ is legal depends on the labeling of the vertices in $N^r(u)$, which depends on the graph structure and the IDs in $N^{\tau+r}(u)$. Due to the hereditary property of the graph class under consideration, for each $u \in V$, $N^{\tau+r}(u)$ is isomorphic to a subgraph of some $2^{\ell'}$ -vertex graph in the same class. Moreover, the shortened ID in $N^{\tau+r}(u)$ are distinct. Therefore, it is guaranteed that the output of the simulation is a legal labeling.

The total time complexity is at most

$$(2\tau + 2r) \cdot O(\log^* n - \log^* \Delta + 1) + \tau = O(\log^k \Delta (\log^* n - \log^* \Delta + 1)).$$

□

A note about MIS lower bounds. Kuhn, Moscibroda, and Wattenhofer [22] showed that for a variety of problems (including MIS) there is a lower bound of $\min(\log \Delta / \log \log \Delta, \sqrt{\log n / \log \log n})$ rounds. The lower bound graph they used to prove such these result has $\log \Delta / \log \log \Delta = O(\sqrt{\log n / \log \log n})$. By Theorem 8, setting $k = 1$ implies that if there is a deterministic algorithm for MIS that runs in $O(\sqrt{\log n})$ time, then there is another deterministic algorithm running in $O(\log \Delta (\log^* n - \log^* \Delta + 1))$ time. Interestingly, Barenboim, Elkin, Pettie, and Schneider [5] showed that an MIS algorithm in RandLOCAL running in $O(\log^k \Delta + f(n))$ -time implied another RandLOCAL algorithm running in $O(\log^k \lambda + \log^{1-\frac{1}{k+1}} n + f(n))$ time on graphs of arboricity λ .

6 Algorithms for Δ -coloring Trees

In Section 4, we showed that the problem of Δ -coloring on trees has an $\Omega(\log_{\Delta} n)$ deterministic lower bound and an $\Omega(\log_{\Delta} \log n)$ randomized lower bound. These lower bounds have matching upper bounds by merely a $\log^* n$ additive term.

The algorithm of Barenboim and Elkin [3] demonstrates that the deterministic bound is essentially tight. They proved that Δ -coloring unoriented trees, where $\Delta \geq 3$, takes $O(\log_{\Delta} n + \log^* n)$ time. This is actually a special case of their algorithm, which applies to graphs of bounded arboricity λ .

Theorem 9 ([3]). *For $q \geq 3$, there is a DetLOCAL algorithm for q -coloring trees in $O(\log_q n + \log^* n)$ time, independent of Δ .*

Pettie and Su [29] gave randomized algorithms for $(4 + o(1))\Delta / \ln \Delta$ -coloring triangle-free graphs. Their algorithm makes extensive use of the distributed Lovász local lemma [12] and runs in $\Omega(\log n)$ time. Pettie and Su sketched a proof that Δ -coloring trees takes $O(\log_{\Delta} \log n + \log^* n)$ time, at least for sufficiently large Δ .

Theorem 10 ([29]). *There exists a large constant Δ_0 such that when $\Delta \geq \Delta_0$, there is a RandLOCAL algorithm for Δ -coloring trees in $O(\log_{\Delta} \log n + \log^* n)$ time.*

The nature of the proof of Theorem 10 makes it difficult to calculate a specific Δ_0 for which the theorem applies. Moreover, the proof is only sketched, hidden inside more complicated ideas. We address both of these issues. Firstly, we provide a simple algorithm and elementary proof of Theorem 10. Secondly, we prove Theorem 11, which combines Theorem 10 with a new technique for constant $\Delta \geq 55$, thereby providing a randomized algorithm for Δ -coloring a tree that runs in $O(\log_{\Delta} \log n + \log^* n)$ time for any constant $\Delta \geq 55$.

6.1 A simple proof of Theorem 10.

For a graph $G = (V, E)$ we say that a subset $S \subseteq V$ is a *distance- k* set if the following two conditions are met:

1. For any two distinct vertices $u, v \in S$, we have $u \notin N^{k-1}(v)$.
2. Let $G^k = (V, E^k)$, where there is an edge $(u, v) \in E^k$ if and only if $\text{dist}_G(u, v) = k$. Then S is connected in G^k .

We make use of the following lemmas in the proof of Theorem 10. While the proof of this lemma is implicit in [5] we reproduce it here for sake of clarity.

Lemma 3 ([5]). *The number of distinct distance- k sets of size t is less than $4^t \cdot n \cdot \Delta^{k(t-1)}$.*

Proof. A distance- k set is spanned by a tree in G^k . There are less than 4^t distinct unlabeled trees of t vertices, and there are less than $n\Delta^{k(t-1)}$ ways to embed a t -vertex tree in G^k . The lemma follows since there is an injective mapping from the family of distance- k sets of size t to subtrees of t vertices in G^k . \square

Lemma 4 (Chernoff bound). *Let X be the sum of n i.i.d. random 0/1 variables. For any $0 < \delta < 1$, we have:*

$$\begin{aligned} \text{For } 0 < \delta < 1, \quad \Pr[X \leq (1 - \delta)\mathbb{E}[X]] &< \exp(-\delta^2\mathbb{E}[X]/2). \\ \text{For } 0 < \delta < 1, \quad \Pr[X \geq (1 + \delta)\mathbb{E}[X]] &< \exp(-\delta^2\mathbb{E}[X]/3). \\ \text{For } \delta \geq 1, \quad \Pr[X \geq (1 + \delta)\mathbb{E}[X]] &< \exp(-\delta\mathbb{E}[X]/3). \end{aligned}$$

Proof of Theorem 10. Our algorithm has two phases. The first phase, which takes $t = O(\log^* \Delta)$ rounds, partially colors the graph using colors in $\{1, 2, \dots, \Delta - \sqrt{\Delta}\}$. The second phase, which takes $O(\log_{\Delta} \log n + \log^* n)$ rounds, applies a deterministic algorithm to $\sqrt{\Delta}$ -color the remaining uncolored vertices using colors in $\{\Delta - \sqrt{\Delta} + 1, \dots, \Delta\}$. We assume throughout the proof that Δ is at least a large enough constant.

Phase 1. The first phase of the algorithm takes $t = O(\log^* \Delta)$ rounds. In each round, the algorithm attempts to color some uncolored vertices. We will explain soon how uncolored vertices decide if they participate in a given round. In the beginning of round i , for each vertex $v \in V$, let $\Psi_i(v)$ denote v 's available palette (i.e. the set of colors that v can choose in round i), and let $N_i(v)$ denote the set of uncolored vertices adjacent to v that are trying to color themselves in this round. Initially, we set $N_1(v) = N(v)$, and $\Psi_1(v) = \{1, 2, \dots, \Delta - \sqrt{\Delta}\}$, for all v . That is, in the first round all vertices attempt to color themselves, and they all have the full palette of this phase available for choices of a color.

We maintain the following two properties for each vertex v that is attempting to color itself at round i :

- **Large Palette Property.** $\mathcal{P}_1(v) : |\Psi_i(v)| \geq \frac{\Delta}{200}$.
- **Small Degree Property.** $\mathcal{P}_2(v) : |N_i(v)| \leq \frac{\Delta}{c_i}$, where c_i is defined as: $c_1 = 1$, $c_2 = 1 - \frac{1}{200}$, and $c_i = \min\{\Delta^{0.1}, c_{i-1} \cdot \exp(\frac{c_{i-1}}{3 \cdot 200 \cdot e^{200}})\}$ for $i > 2$.

Notice that c_i is a constant, for all i . Let t be the smallest number i such that $c_i = \Delta^{0.1}$. Notice that $t = O(\log^* \Delta)$.

The intuition behind the two properties $\mathcal{P}_1(v)$ and $\mathcal{P}_2(v)$ is that they ensure that (i) participating vertices always have a large enough palette to use, and (ii) there is a large separation between the palette size and the degree so that we can color a large fraction of vertices in each round.

For each $1 \leq i \leq t$, the i^{th} round consists of two constant time sub-routines **ColorBidding**(i) and **Filtering**(i). In **ColorBidding**(i), each participating vertex v selects a random subset of colors S_v . If there is a color in S_v that does not belong to $\bigcup_{u \in N_i(v)} S_u$, the vertex v *succeeds* and colors itself with any such color. If such a color is chosen, denote it by **Color**(v). After **ColorBidding**(i), we execute **Filtering**(i) which filters out some vertices and thereby preventing \mathcal{P}_1 and \mathcal{P}_2 from being violated. Such vertices are called *bad* vertices, and they will no longer participate in the remaining rounds of Phase 1.

ColorBidding(i).

Do the following steps in parallel for each uncolored vertex v that is not bad:

1. If $c_i = 1$, then choose one color uniformly at random from $\Psi_i(v)$, and S_v contains only this color. Otherwise ($c_i > 1$), construct the set S_v by independently including each color of $\Psi_i(v)$ with probability $\frac{c_i}{|\Psi_i(v)|}$.

2. If $S_v \setminus \bigcup_{u \in N_i(v)} S_u \neq \emptyset$, then permanently color v by picking an arbitrary color in $S_v \setminus \bigcup_{u \in N_i(v)} S_u$ for $\text{Color}(v)$.
3. $\Psi_{i+1}(v) \leftarrow \Psi_i(v) \setminus \{\text{Color}(u) \mid u \in N_i(v) \text{ is permanently colored}\}$.

We define $N'_i(v)$ as the set of participating vertices *after* $\text{ColorBidding}(i-1)$ and *before* $\text{Filtering}(i-1)$ that are adjacent to v . In other words, $N'_i(v) = N_{i-1}(v) \setminus \{u \mid u \text{ is permanently colored in } \text{ColorBidding}(i-1)\}$.

Filtering(i).

For each uncolored vertex v that is not bad:

1. If $i = 1$ and $|\Psi_2(v)| - |N'_2(v)| < \frac{\Delta}{200}$, then mark v as a bad vertex.
2. If $1 < i < t$ and $|N'_{i+1}(v)| > \frac{\Delta}{c_{i+1}}$, then mark v as a bad vertex.
3. If $i = t$ then mark v as a bad vertex.

Phase 2. By the filtering rule for $i = t$, all the remaining uncolored vertices after the Phase 1 are bad vertices. We color the bad vertices in Phase 2. We will later prove that after phase 1, with high probability a connected component induced by bad vertices has size at most $\Delta^4 \log n$. Hence we use Theorem 9 to $\sqrt{\Delta}$ -color such connected components using the $\sqrt{\Delta}$ reserved colors. For simplicity, if this phase lasts for too long (which may happen with low probability) the algorithm just stops and fails.

Runtime. The runtime of phase 1 is $t = O(\log^* \Delta)$ rounds. The runtime of phase 2 is $O(\log_{\sqrt{\Delta}} (\Delta^4 \log n) + \log^* (\Delta^4 \log n)) = O(\log_{\Delta} \log n + \log^* n)$. Thus, the total runtime is $O(\log_{\Delta} \log n + \log^* n)$ rounds.

Analysis. The analysis of phase 2 relies only on proving that, with high probability, all of the connected components induced by bad vertices after phase 1 are of size at most $\Delta^4 \log n$. So we focus on analyzing phase 1.

The decision of whether a vertex v that participates in round i becomes marked as bad in this round depends on the vertices in $N^2(v)$ that participate in this round and the random bits used by these vertices during this round. Our analysis will work for any such arbitrary setting (of participating vertices in $N^2(v)$ and their random bits). In particular, we prove that in any round of Phase 1, for any vertex v participating in this round and any arbitrary choice of vertices in $N^2(v)$ that participate in this round with their random bits (that are used in this round), v becomes a bad vertex with probability at most $\exp(-\text{poly}(\Delta))$. Such a proof means that the choice of whether v becomes bad or not does not depend on any arbitrary behaviour of vertices participating in this round that are not in $N^2(v)$. This proof is covered by the following claims, whose proofs are given in the Appendix (we partition the cases since each case requires a different proof).

Claim 1. *The probability that a vertex v is marked as bad in round $i = 1$ is at most $\exp(-\Omega(\Delta))$, regardless of the random bits used by vertices in $N^2(v)$.*

Claim 2. *The probability that a vertex v that participates in round $1 < i < t$ is marked as bad in round i is at most $\exp(-\Omega(\Delta^{0.1}))$, regardless of the random bits used by vertices in $N^2(v)$.*

Claim 3. *The probability that a vertex v that participates in round $i = t$ is marked as bad in round i is at most $\exp(-\Omega(\Delta^{0.1}))$, regardless of the random bits used by vertices in $N^2(v)$.*

By the union bound for all rounds in Phase 1, the probability that any vertex v becomes a bad vertex after Phase 1 is $O(\log^* \Delta) \exp(-\text{poly}(\Delta)) = \exp(-\text{poly}(\Delta))$, regardless of the choice of random bits for all vertices not in $N^2(v)$. Therefore, just before Phase 2, for any distant-5 set T of size s , the probability that all vertices in T are bad is at most $\exp(-s \cdot \text{poly}(\Delta))$. By Lemma 3, there are at most $4^s \cdot n \cdot \Delta^{4(s-1)}$ distinct

distant-5 set T of size s . By the union bound, with probability at least $(4^s \cdot n \cdot \Delta^{4(s-1)}) \cdot \exp(-s \cdot \text{poly}(\Delta))$, there is no distant-5 set of size s that contains only bad vertices. This probability can be upper bounded by n^{-c} for any c when $s = \log n$.

This concludes the proof of Theorem 10. \square

6.2 Algorithm for $\Delta \geq 55$.

Inherently, the above proof (and also the proof in [29]) is hard to analyze quantitatively without the aid of $O(\cdot)$ notation. It seems to require a very large Δ for the proof to work well, since in each round several Chernoff bounds are applied to make sure that some requirements are met (see the proof of Claims 1, 2, and 3 in the Appendix), and we need a large enough Δ to make these Chernoff bounds work. In what follows we present a different algorithm with a significantly simpler analysis for Δ -coloring trees.

Theorem 11. *For $\Delta \geq 55$, there exists a RandLOCAL algorithm Δ -coloring of a tree can be computed in $O(\log_\Delta \log n + \log^* n)$ time.*

Proof. We assume that $\Delta = O(1)$ is constant. If it is sufficiently large, apply Theorem 10. Our algorithm has three phases:

Phase 1. We execute the following procedure to partially color the graph with colors in $\{4, 5, \dots, \Delta\}$.

Initially $U \leftarrow V$.

For i from Δ down to 4, do the following steps in parallel for each vertex $v \in U$:

1. Choose a real number $x(v) \in [0, 1]$ uniformly at random.
2. Let $K = \left\{ v \mid x(v) < \min_{u \in N(v) \cap U} x(u) \right\}$ be the set of all vertices holding local minima.
3. Find any MIS $I \supseteq K$ of U . All vertices in I are colored i .
4. Set $U \leftarrow U \setminus I$ (remove all colored vertices).

The above procedure ensures that the number of uncolored neighbors of a vertex $v \in U$ is at most $i - 1$ after step 4. Therefore, at the end of the Phase 1, we have $|N(v) \cap U| \leq 3$ for any uncolored vertex v .

The MIS required in Step 3 can be computed in $O(\Delta + \log^* n) = O(\log^* n)$ time [4], or in $O(\Delta^2 + \log^* n) = O(\log^* n)$ time via Theorem 2.

Phase 2. We will later show that the set of the vertices $S = \{v \in U \text{ s.t. } |N(v) \cap U| = 3\}$ form connected components of size at most $O(\log n)$ with probability $\geq 1 - n^{-c}$. Hence we apply Theorem 9 to 3-color the set S (using the colors 1, 2, 3) in $O(\log \log n)$ time. We then update $U = U \setminus S$ after coloring the vertices in S .

Phase 3. For each vertex v that remains uncolored, the number of its available colors (i.e. $\{1, \dots, \Delta\} \setminus \{\text{color}(u) \mid u \in N(v) \text{ is colored}\}$) is greater than the number of its uncolored neighbors (i.e. $|N(v) \cap U|$). We apply an $O(\log^* n)$ -time MIS algorithm twice to get a 3-coloring of vertices in U (with three colors $1', 2', 3'$). For $i = 1', 2', 3'$, we recolor the vertices that are colored with i' using any of its available colors.

In view of the above, to prove the theorem, it suffices to show that the set $S = \{v \in U \text{ s.t. } |N(v) \cap U| = 3\}$ (which is defined in Phase 2) form connected components of size at most $O(\log n)$ with probability $\geq 1 - n^{-c}$.

Given any distant-3 set T of size t , we select any vertex in $V \setminus T$ as a root to make the tree rooted, and for each $v_i \in T$, we define w_i as the parent of v_i . We also define $D_i = \bigcup_{u \in N(v_i) \setminus \{w_i\}} N(u)$, for each $v_i \in T$. We observe that $\forall v_i, v_j \in T, D_i \cap D_j = \emptyset$. We prove the the following claim (The proof is given in the Appendix):

Claim 4. *There is some constant ϵ such that the probability for a vertex v_i to be in S conditioned on arbitrary behavior of vertices not in D_i is at most $\frac{1-\epsilon}{4\Delta^3}$.*

Thus, the probability that all vertices in T belong to S is at most $(\frac{1-\epsilon}{4\Delta^3})^t$. By Lemma 3, as long as

$$t \geq \frac{(c+1)\log n - \log(\Delta^3)}{\log \frac{1}{1-\epsilon}},$$

with probability $\geq 1 - n^{-c}$ we have:

$$4^t n \Delta^{3(t-1)} \cdot \left(\frac{1-\epsilon}{4\Delta^3}\right)^t \leq n^{-c}.$$

In words, with high probability there is no distant-3 set of size t whose vertices are all in S .

Since (for $\Delta \geq 2$) any connected subgraph with number of vertices at least $\Delta^2 t$ must contain a distant-3 set of size t , we conclude that, with high probability, S forms connected components of size at most $O(\Delta^2 \log n) = O(\log n)$. \square

References

- [1] R. Bar-Yehuda, K. Censor-Hillel, and G. Schwartzman. A distributed $(2 + \epsilon)$ -approximation for vertex cover in $O(\log \Delta / \epsilon \log \log \Delta)$ rounds. *CoRR*, abs/1602.03713, 2016.
- [2] L. 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 (PODC)*, pages 345–354, 2015.
- [3] L. Barenboim and M. Elkin. Sublogarithmic distributed MIS algorithm for sparse graphs using Nash-Williams decomposition. *Distributed Computing*, 22(5-6):363–379, 2010.
- [4] L. Barenboim, M. Elkin, and F. Kuhn. Distributed $(\Delta + 1)$ -coloring in linear (in Δ) time. *SIAM J. Comput.*, 43(1):72–95, 2014.
- [5] L. Barenboim, M. Elkin, S. Pettie, and J. Schneider. The locality of distributed symmetry breaking. *J. ACM*, 2016. to appear.
- [6] M. Ben-Or. Another advantage of free choice: Completely asynchronous agreement protocols. In *Proceedings of the 2nd ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC)*, pages 27–30, 1983.
- [7] T. Bisht, K. Kothapalli, and S. V. Pemmaraju. Brief announcement: Super-fast t -ruling sets. In *Proceedings 33rd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 379–381, 2014.
- [8] B. Bollobás. Chromatic number, girth and maximal degree. *Discrete Mathematics*, 24(3):311–314, 1978.
- [9] B. Bollobás. *Extremal graph theory*, volume 11 of *London Mathematical Society Monographs*. Academic Press Inc. [Harcourt Brace Jovanovich Publishers], London, 1978.
- [10] G. Bracha. An asynchronous $(n - 1)/3$ -resilient consensus protocol. In *Proceedings of the 3rd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 154–162, 1984.
- [11] S. Brandt, O. Fischer, J. Hirvonen, B. Keller, T. Lempiäinen, J. Rybicki, J. Suomela, and J. Uitto. A lower bound for the distributed Lovász local lemma. In *Proceedings 48th ACM Symposium on the Theory of Computing*, page ?, 2016.
- [12] K.-M. Chung, S. Pettie, and H.-H. Su. Distributed algorithms for the Lovász local lemma and graph coloring. In *Proceedings 33rd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 134–143, 2014.

- [13] X. Dahan. Regular graphs of large girth and arbitrary degree. *Combinatorica*, 34(4):407–426, 2014.
- [14] M. Elkin, S. Pettie, and H. H. Su. $(2\Delta - 1)$ -edge coloring is much easier than maximal matching in the distributed setting. In *Proceedings 26th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 355–370, 2015.
- [15] M. J. Fischer, N. A. Lynch, and M. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, 1985.
- [16] M. Ghaffari. An improved distributed algorithm for maximal independent set. In *Proceedings 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 270–277, 2016.
- [17] M. Hańćkowiak, M. Karoński, and A. Panconesi. On the distributed complexity of computing maximal matchings. *SIAM J. Discrete Mathematics*, 15(1):41–57 (electronic), 2001.
- [18] D. Harris, J. Schneider, and H.-H. Su. Distributed $(\Delta + 1)$ -coloring in sublogarithmic rounds. In *Proceedings 48th ACM Symposium on Theory of Computing (STOC)*, 2016.
- [19] V. King and J. Saia. Byzantine agreement in polynomial expected time. In *Proceedings of the 45th ACM Symposium on Theory of Computing (STOC)*, pages 401–410, 2013.
- [20] A. Korman, J.-S. Sereni, and L. Viennot. Toward more localized local algorithms: removing assumptions concerning global knowledge. *Distributed Computing*, 26(5–6):289–308, 2013.
- [21] K. Kothapalli and S. V. Pemmaraju. Super-fast 3-ruling sets. In *Proceedings IARCS Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 18 of *LIPICs*, pages 136–147. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2012.
- [22] F. Kuhn, T. Moscibroda, and R. Wattenhofer. What cannot be computed locally! In *Proceedings 23rd Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 300–309, 2004.
- [23] F. Kuhn, T. Moscibroda, and R. Wattenhofer. Local computation: Lower and upper bounds. *CoRR*, abs/1011.5470, 2010.
- [24] N. Linial. Locality in distributed graph algorithms. *SIAM J. Comput.*, 21(1):193–201, 1992.
- [25] M. Naor. A lower bound on probabilistic algorithms for distributive ring coloring. *SIAM J. Discrete Mathematics*, 4(3):409–412, 1991.
- [26] M. Naor and L. J. Stockmeyer. What can be computed locally? *SIAM J. Comput.*, 24(6):1259–1277, 1995.
- [27] A. Panconesi and R. Rizzi. Some simple distributed algorithms for sparse networks. *Distributed Computing*, 14(2):97–100, 2001.
- [28] A. Panconesi and A. Srinivasan. On the complexity of distributed network decomposition. *J. Algor.*, 20(2):356–374, 1996.
- [29] S. Pettie and H.-H. Su. Distributed algorithms for coloring triangle-free graphs. *Information and Computation*, 243:263–280, 2015.
- [30] M. O. Rabin. Randomized Byzantine generals. In *Proceedings of the 24th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 403–409, 1983.
- [31] J. Schneider and R. Wattenhofer. An optimal maximal independent set algorithm for bounded-independence graphs. *Distributed Computing*, 22(5-6):349–361, 2010.
- [32] M. Sipser. *Introduction to the Theory of Computation*. International Thomson Publishing, 3rd edition, 2012.

Appendix

A Proofs of Claims 1, 2, and 3

The proofs for these claims rely on the Large Palette Property $\mathcal{P}_1(v)$ and the Small Degree Property $\mathcal{P}_2(v)$. First, we observe that our filtering rules imply that these two properties hold after each round:

- The filtering rule for $i = 1$ guarantees that the Large Palette Property $\mathcal{P}_1(v)$ is met for all vertices that remain after the filtering. Notice that $|\Psi_2(v)| - |N'_2(v)| \geq \frac{\Delta}{200}$ implies $|\Psi_i(v)| \geq \frac{\Delta}{200}$, for all i .
- The filtering rules for $i = 1$ and $1 < i < t$ ensures that the Small Degree Property $\mathcal{P}_2(v)$ holds for all i , since $|N'_{i+1}(v)| \leq \frac{\Delta}{c_{i+1}}$ implies that $|N_{i+1}(v)| \leq \frac{\Delta}{c_{i+1}}$.

Proof of Claim 1. Recall that a vertex v is marked as bad in round 1 only if $|\Psi_2(v)| - |N'_2(v)| < \frac{\Delta}{200}$. Thus, we assume $|N(v)| \geq \frac{199 \cdot \Delta}{200}$, since otherwise v will not be marked as bad.

For each neighbor u of v , let E_u denote the event of u being colored in the first round. Since the graph is a tree, for all $u \in N(v)$ the events E_u are independent. Assuming sufficiently large Δ , we have:

$$\Pr[E_u] \geq \left(1 - \frac{1}{\Delta - \sqrt{\Delta}}\right)^{|N(u)|} \geq \left(1 - \frac{1}{\Delta - \sqrt{\Delta}}\right)^\Delta \geq \frac{1}{3}.$$

By a Chernoff bound (with $\delta = \frac{79}{199}$, and the expected number of colored neighbors being at least $\frac{199}{200} \cdot \frac{\Delta}{3}$), the number of colored neighbors of v in the first round is at least $\frac{\Delta}{5} = (1 - \delta) \cdot \frac{199}{200} \cdot \frac{\Delta}{3}$ with probability at least

$$1 - \exp\left(\frac{-\left(\frac{79}{199}\right)^2 \cdot \left(\frac{199}{200} \cdot \frac{\Delta}{3}\right)}{2}\right).$$

Let $S = \{u_1, u_2, \dots\}$ be the colored subset of $N(v)$. In what follows, we assume that $|S| \geq \frac{\Delta}{5}$. Conditioned on the color selected by v and the event that S is the colored subset of $N(v)$, each u_j independently selects a color uniformly at random from $\{1, \dots, \Delta - \sqrt{\Delta}\} \setminus \{\text{Color}(v)\}$.

If $\frac{\Delta}{10} - |\bigcup_{j=1}^{\Delta/10} \{\text{Color}(u_j)\}| \geq \frac{\Delta}{200}$, then $|\Psi_2(v)| - |N_2(v)| \geq \frac{\Delta}{200}$, and v is not marked as bad. Otherwise, each u_j , $j > \frac{\Delta}{10}$, chooses a color that is already chosen by some u_k , $k < j$, with probability at least

$$\frac{\frac{\Delta}{10} - \frac{\Delta}{200}}{\Delta - \sqrt{\Delta} - 1} \geq \frac{1}{11}.$$

As a result, the expected value of $|S| - |\bigcup_{j=1}^{|S|} \{\text{Color}(u_j)\}|$ is at least $\frac{\Delta}{10} \cdot \frac{1}{11} = \frac{\Delta}{110}$. By a Chernoff bound (with $\delta = \frac{9}{20}$), this value is at least $\frac{\Delta}{200} = (1 - \delta) \cdot \frac{\Delta}{110}$ with probability at least

$$1 - \exp\left(\frac{-\left(\frac{9}{20}\right)^2 \cdot \frac{\Delta}{110}}{2}\right) = 1 - \exp\left(\frac{-162\Delta}{11}\right).$$

By definition, $|\Psi_2(v)| = |\Psi_1(v)| - |\bigcup_{j=1}^{|S|} \{\text{Color}(u_j)\}|$, and $|N'_2(v)| = |N_1(v) \setminus S| = |N_1(v)| - |S|$. In view of the above, $|\Psi_2(v)| - |N_2(v)| \geq \frac{\Delta}{200}$ (and v is not marked as bad) with probability at least

$$1 - \exp\left(\frac{-\left(\frac{79}{199}\right)^2 \cdot \left(\frac{199}{200} \cdot \frac{\Delta}{3}\right)}{2}\right) - \exp\left(\frac{-162\Delta}{11}\right) = 1 - \exp(-\Omega(\Delta)).$$

□

Proof of Claim 2. Our goal is to show that $|N'_{i+1}(v)| \leq \frac{\Delta}{c_i \cdot \exp(\frac{c_i}{3 \cdot 200 \cdot e^{200}})} = \frac{\Delta}{c_{i+1}}$ with high probability.

We include each available color in S_v independently with probability $\frac{c_i}{|\Psi_i(u)|}$, so the expected value of $|S_v|$ is c_i . By a Chernoff bound, the event of $|S_v| \leq \frac{\Delta}{2 \cdot 200} = (1 + \delta)c_i$, where $\delta = \left(\frac{\Delta}{2 \cdot 200 \cdot c_i} - 1\right)$, happens with probability at least:

$$1 - \exp\left(\frac{-\left(\frac{\Delta}{2 \cdot 200 \cdot c_i} - 1\right) \cdot c_i}{2}\right) \geq 1 - \exp\left(\frac{-\Delta}{5 \cdot 200}\right).$$

For any $u \in N_i(v)$, a color in $\Psi_i(u) \setminus S_v$ belongs to $S_u \setminus \bigcup_{w \in N_i(u)} S_w$ with probability at least

$$\frac{c_i}{|\Psi_i(u)|} \cdot \prod_{w \in N_i(u) \setminus \{v\}} \left(1 - \frac{c_i}{|\Psi_i(w)|}\right) \geq \frac{c_i}{|\Psi_i(u)|} \cdot \left(1 - \frac{c_i \cdot 200}{\Delta}\right)^{\frac{\Delta}{c_i} - 1} \geq \frac{c_i}{e^{200} |\Psi_i(u)|}.$$

Notice that the term $\frac{c_i}{|\Psi_i(u)|}$ is the probability that a color in $\Psi_i(u) \setminus S_v$ is chosen to be in S_u , and the term $\prod_{w \in N_i(u) \setminus \{v\}} \left(1 - \frac{c_i}{|\Psi_i(w)|}\right)$ is the probability that a color in $\Psi_i(u) \setminus S_v$ does not belong to $\bigcup_{w \in N_i(u)} S_w$.

Under the condition that $|S_v| \leq \frac{\Delta}{2 \cdot 200}$, we have $|\Psi_i(u) \setminus S_v| \geq \frac{\Delta}{200} - \frac{\Delta}{2 \cdot 200} = \frac{\Delta}{2 \cdot 200}$. Then the set $S_u \setminus \bigcup_{w \in N_i(u)} S_w$ is empty with probability at most:

$$\left(1 - \frac{c_i}{e^{200} |\Psi_i(u)|}\right)^{\frac{\Delta}{2 \cdot 200}} \leq \left(1 - \frac{c_i}{e^{200} \Delta}\right)^{\frac{\Delta}{2 \cdot 200}} \leq \exp\left(-\frac{c_i}{2 \cdot 200 \cdot e^{200}}\right).$$

Hence $u \in N_i(v)$ remains uncolored with probability at most $\exp\left(-\frac{c_i}{2 \cdot 200 \cdot e^{200}}\right)$.

[Case 1. $\exp\left(-\frac{c_i}{2 \cdot 200 \cdot e^{200}}\right) \cdot |N_i(v)| \geq \Delta^{0.1}$] We choose δ to be the number such that $(1 + \delta) \exp\left(-\frac{c_i}{2 \cdot 200 \cdot e^{200}}\right) = -\frac{c_i}{3 \cdot 200 \cdot e^{200}}$. Notice that we always have $\delta \geq \exp\left(\frac{1}{6 \cdot 200 \cdot e^{200}}\right) - 1$, which is a positive constant. By a Chernoff bound on all vertices in $N_i(v)$, we have $|N'_{i+1}(v)| \leq \exp\left(-\frac{c_i}{3 \cdot 200 \cdot e^{200}}\right) \cdot |N_i(v)|$ with probability at least

$$1 - \exp\left(\frac{-\max\{1, \delta^2\} \cdot \exp\left(-\frac{c_i}{2 \cdot 200 \cdot e^{200}}\right) \cdot |N_i(v)|}{3}\right) \geq 1 - \exp(-\Omega(\Delta^{0.1})).$$

[Case 2. $\exp\left(-\frac{c_i}{2 \cdot 200 \cdot e^{200}}\right) \cdot |N_i(v)| < \Delta^{0.1}$] By a Chernoff bound on all vertices in $N_i(v)$, we have $|N'_{i+1}(v)| \leq \Delta^{0.8} \cdot \exp\left(-\frac{c_i}{2 \cdot 200 \cdot e^{200}}\right) \cdot |N_i(v)| \leq \Delta^{0.9}$ with probability at least

$$1 - \exp\left(\frac{(\Delta^{0.8} - 1) \cdot \exp\left(-\frac{c_i}{2 \cdot 200 \cdot e^{200}}\right) \cdot |N_i(v)|}{3}\right) \geq 1 - \exp\left(\frac{-\Delta^{0.9} + \Delta^{0.1}}{3}\right).$$

In any case, we have $|N'_{i+1}(v)| \leq \max\left\{\frac{|N_i(v)|}{\exp\left(\frac{c_i}{3 \cdot 200 \cdot e^{200}}\right)}, \Delta^{0.9}\right\} = \frac{\Delta}{c_{i+1}}$ (and so v is not marked as bad) with probability at least $1 - \exp(-\Omega(\Delta^{0.1}))$. \square

Proof of Claim 3. A color in $\Psi_i(v)$ belongs to $S_v \setminus \bigcup_{u \in N_i(v)} S_u$ with probability at least

$$\frac{c_i}{|\Psi_i(v)|} \cdot \prod_{u \in N_i(v)} \left(1 - \frac{c_i}{|\Psi_i(u)|}\right) \geq \frac{c_i}{|\Psi_i(v)|} \cdot \left(1 - \frac{c_i \cdot 200}{\Delta}\right)^{\frac{\Delta}{c_i}} \geq \frac{c_i}{1.1 e^{200} |\Psi_i(v)|}.$$

Therefore, v remains uncolored (and is marked as bad) with probability at most

$$\left(1 - \frac{c_i}{1.1 e^{200} |\Psi_i(v)|}\right)^{|\Psi_i(v)|} \leq \exp\left(-\frac{\Delta^{0.1}}{1.1 e^{200}}\right).$$

\square

B Proof of Claim 4

Proof of Claim 4. For notational simplicity, we write $v \stackrel{\text{def}}{=} v_i$ and $w \stackrel{\text{def}}{=} w_i$.

We observe that $v \in S$ only when for all $i = \Delta$ to 4, at most one neighbor of v is colored with i in Step 2. In addition, in the beginning of Step 1, we must have $|N(v) \cap U| = i$, for all $i = \Delta$ to 4.

Now, assume that we are at the beginning of Step 1, the vertex v still has no neighbors of repeated colors, and $|N(v) \cap U| = i$. The probability that a neighbor $u \in N(v) \setminus \{w\}$ is colored i in Step 2 when $x(v) = z \in [0, 1]$ is at least:

$$\int_{y=0}^{y=z} \Pr[\forall r \in (N(u) \cap U) \setminus \{v\}, x(r) \leq y] dy \geq \int_{y=0}^{y=z} (1-y)^{i-1} dy = \frac{1 - (1-z)^i}{i}.$$

Note that the variable y represents the random variable $x(u)$, and u is colored when (i) $y \in [0, z)$, and (ii) $x(r) \in (y, 1]$ for all $r \in (N(u) \cap U) \setminus \{v\}$. Also note that $|(N(u) \cap U) \setminus \{v\}| \leq i - 1$.

We write $p_i(z) = \frac{1 - (1-z)^i}{i}$. Then the probability that at most one neighbor of v is colored with i in Step 2 can be upper bounded by:

$$\begin{aligned} P_i &= \int_{x=0}^{x=1} \Pr[\text{binom}(|(N(v) \cap U) \setminus \{w\}|, p_i(x)) \leq 1] dx \\ &\leq \int_{x=0}^{x=1} [(1 - p_i(x))^{i-1} + (i-1)p_i(x)(1 - p_i(x))^{i-2}] dx. \end{aligned}$$

Notice that $|(N(v) \cap U) \setminus \{w\}| \geq i - 1$.

When $i = 4$, the P_4 is about 0.88718. As i increases, P_i decreases monotonically, approaching 0.73576. By a numerical calculation, so long as $\Delta \geq 55$, the probability that v is in S conditioned on arbitrary behavior of vertices not in $\bigcup_{u \in N(v) \setminus \{w\}} N(u)$ is at most

$$\prod_{i=4}^{i=\Delta} P_i < \frac{1}{4\Delta^3},$$

as desired. □