

# Distributed $\Delta$ -Coloring Plays Hide-and-Seek

**Alkida Balliu** · alkida.balliu@cs.uni-freiburg.de · University of Freiburg

**Sebastian Brandt** · brandt@cispa.de · CISPA Helmholtz Center for Information Security

**Fabian Kuhn** · kuhn@cs.uni-freiburg.de · University of Freiburg

**Dennis Olivetti** · dennis.olivetti@cs.uni-freiburg.de · University of Freiburg

## Abstract

We prove several new tight or near-tight distributed lower bounds for classic symmetry breaking problems in graphs. As a basic tool, we first provide a new insightful proof that any deterministic distributed algorithm that computes a  $\Delta$ -coloring on  $\Delta$ -regular trees requires  $\Omega(\log_{\Delta} n)$  rounds and any randomized such algorithm requires  $\Omega(\log_{\Delta} \log n)$  rounds. We prove this by showing that a natural relaxation of the  $\Delta$ -coloring problem is a fixed point in the round elimination framework.

As a first application, we show that our  $\Delta$ -coloring lower bound proof directly extends to arbdefective colorings. An arbdefective  $c$ -coloring of a graph  $G = (V, E)$  is given by a  $c$ -coloring of  $V$  and an orientation of  $E$ , where the arbdefect of a color  $i$  is the maximum number of monochromatic outgoing edges of any node of color  $i$ . We exactly characterize which variants of the arbdefective coloring problem can be solved in  $O(f(\Delta) + \log^* n)$  rounds, for some function  $f$ , and which of them instead require  $\Omega(\log_{\Delta} n)$  rounds for deterministic algorithms and  $\Omega(\log_{\Delta} \log n)$  rounds for randomized ones.

As a second application, which we see as our main contribution, we use the structure of the fixed point as a building block to prove lower bounds as a function of  $\Delta$  for problems that, in some sense, are *much easier* than  $\Delta$ -coloring, as they can be solved in  $O(\log^* n)$  deterministic rounds in bounded-degree graphs. More specifically, we prove lower bounds as a function of  $\Delta$  for a large class of distributed symmetry breaking problems, which can all be solved by a simple sequential greedy algorithm. For example, we obtain novel results for the fundamental problem of computing a  $(2, \beta)$ -ruling set, i.e., for computing an independent set  $S \subseteq V$  such that every node  $v \in V$  is within distance  $\leq \beta$  of some node in  $S$ . We in particular show that  $\Omega(\beta\Delta^{1/\beta})$  rounds are needed even if initially an  $O(\Delta)$ -coloring of the graph is given. With an initial  $O(\Delta)$ -coloring, this lower bound is tight and without, it still nearly matches the existing  $O(\beta\Delta^{2/(\beta+1)} + \log^* n)$  upper bound. The new  $(2, \beta)$ -ruling set lower bound is an exponential improvement over the best existing lower bound for the problem, which was proven in [FOCS '20]. As a special case of the lower bound, we also obtain a tight linear-in- $\Delta$  lower bound for computing a maximal independent set (MIS) in trees. While such an MIS lower bound was known for general graphs, the best previous MIS lower bounds for trees was  $\Omega(\log \Delta)$ . Our lower bound even applies to a much more general family of problems that allows for almost arbitrary combinations of natural constraints from coloring problems, orientation problems, and independent set problems, and provides a single unified proof for known and new lower bound results for these types of problems.

All of our lower bounds as a function of  $\Delta$  also imply substantial lower bounds as a function of  $n$ . For instance, we obtain that the maximal independent set problem, on trees, requires  $\Omega(\log n / \log \log n)$  rounds for deterministic algorithms, which is tight.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Problem Definitions . . . . .	5
1.2	Our Results . . . . .	6
1.3	Additional Related Work . . . . .	8
<b>2</b>	<b>Road Map</b>	<b>11</b>
<b>3</b>	<b>Preliminaries</b>	<b>12</b>
3.1	<b>LOCAL</b> model . . . . .	12
3.2	Black-White Formalism . . . . .	14
3.3	Round Elimination . . . . .	15
<b>4</b>	<b><math>\Delta</math>-Coloring Fixed Point and Its Applications</b>	<b>19</b>
4.1	A Fixed Point For $\Delta$ -Coloring . . . . .	20
<b>5</b>	<b>The Problem Family</b>	<b>23</b>
5.1	Problem Definition . . . . .	23
5.2	Intuition Behind These Problems . . . . .	24
5.3	Relation with Natural Problems. . . . .	26
<b>6</b>	<b>Lower Bounds</b>	<b>26</b>
6.1	Computing $\mathcal{R}(\Pi_\Delta(z))$ . . . . .	27
6.2	Proof of Lemma 6.4 . . . . .	30
6.3	Relaxing $\overline{\mathcal{R}}(\mathcal{R}(\Pi_\Delta(z)))$ . . . . .	36
6.4	Zero Round Solvability . . . . .	39
6.5	A Problem Sequence . . . . .	39
<b>7</b>	<b>Lifting Results to the LOCAL Model</b>	<b>40</b>
<b>8</b>	<b>Lower Bound Results</b>	<b>40</b>
<b>9</b>	<b>Upper Bounds</b>	<b>44</b>
9.1	Algorithms for Arbdefective Colored Ruling Sets . . . . .	45
<b>10</b>	<b>Open Problems</b>	<b>46</b>
<b>A</b>	<b>Proof of Theorem 7.1</b>	<b>52</b>

# 1 Introduction

In the area of distributed graph algorithms, we usually assume that the graph  $G = (V, E)$  on which we want to solve some graph problem also represents a communication network. The nodes of  $G$  are autonomous agents that communicate over the edges of  $G$ . Each node  $v \in V$  is equipped with a unique  $O(\log n)$ -bit identifier (where  $n = |V|$ ). More precisely, the nodes typically interact in synchronous rounds and in each round, every node  $v \in V$  can send a message to each of its neighbors in  $G$ . In the standard LOCAL model [Lin87, Pel00], those messages can be of arbitrary size. Except for information that we assume to be publicly known (such as, e.g., the family of graphs from which  $G$  is chosen), at the beginning, nodes do not know anything about the graph  $G$ . At the end of an algorithm, every node  $v \in V$  has to output its local part of the solution of the considered graph problem. The time or round complexity of such a distributed algorithm is defined as the number of rounds that are necessary until all nodes terminate and output their part of the solution.

Since the early works on distributed and parallel graph algorithms [ABI86, Lin87, Lub86], two problems that have been at the very core of the attention are the problems of computing a vertex coloring of  $G$  and of computing a maximal independent set (MIS) of  $G$ . Both problems, as well as their many variants, have been studied intensively throughout the last more than 30 years. A well-studied variant of MIS is the problem of computing an  $(\alpha, \beta)$ -ruling set, introduced in the late 1980s [AGLP89]. The problem asks for a subset  $S$  of the nodes satisfying that nodes in  $S$  are at distance at least  $\alpha$  from each other, and nodes not in  $S$  are at distance at most  $\beta$  from a node in  $S$ . Note that  $(2, 1)$ -ruling sets are equivalent to MIS, and by increasing  $\beta$  we obtain relaxed variants of MIS. Interestingly, it is often the case that whenever the computation of an MIS can be used as a subroutine to solve some problem of interest, it is actually enough to compute a  $(2, \beta)$ -ruling set for some  $\beta > 1$ , which is easier to compute. In fact, the fastest known distributed algorithms for several fundamental distributed problems use the computation of ruling sets as subroutines and the complexity of those algorithms critically depends on having fast  $(2, \beta)$ -ruling set algorithms for some values  $\beta > 1$ . Examples are deterministic and randomized algorithms for computing  $\Delta$ -colorings<sup>1</sup> [GHKM18, PS95], randomized distributed algorithms for the Lovász Local Lemma [FG17, GHK18], as well as randomized distributed algorithms for computing an MIS [Gha16, BEPS16]. In some cases, improved algorithms for ruling sets would directly lead to improved algorithms for such problems. Hence, understanding if it is possible to obtain faster algorithms for ruling sets is of great importance. The first lower bounds on computing ruling sets (except MIS) were proven in [BBO20] and as one of the main contributions of our work, we improve those lower bounds exponentially and we in fact show that some existing ruling set algorithms are optimal. This for example implies that in order to improve the current best algorithm for  $\Delta$ -coloring [GHKM18], we need to find a genuinely different algorithm, which is not based on ruling sets.

**Round Elimination.** As many recent lower bounds for local distributed graph problems, the lower bounds of [BBO20, BBKO21] are based on the *round elimination* technique [BFH<sup>+</sup>16, Bra19]. The high-level idea of the round elimination technique is as follows. Assume that the time needed to solve some desired problem, such as computing an MIS, is  $T$  rounds. Then, every algorithm that solves the problem in  $T$  rounds must already have computed something non-trivial after  $T - 1$  rounds. And it turns out that one can exactly characterize the weakest problem that needs to be solved after  $T - 1$  rounds if the desired problem has to be computed after  $T$  rounds. Starting from our target problem, we can then construct a sequence of problems where each problem in the sequence

---

<sup>1</sup>As common when considering the distributed  $\Delta$ -coloring problem, we will always implicitly assume that  $\Delta \geq 3$  and that the input graph does not contain a connected component that is  $\Delta$ -regular (which in particular implies, by Brooks' Theorem [Bro41], that the input graph is  $\Delta$ -colorable).

is exactly one round easier than the previous one. And if we can construct a sequence of length  $R$  such that the last problem of the sequence still cannot be solved in 0 rounds, then we know that our target problem cannot be solved in  $R$  rounds.

In the last few years, round elimination has been used to prove an impressive number of strong new distributed lower bounds [BFH<sup>+</sup>16, CKP19, CHL<sup>+</sup>20, Bra19, BBH<sup>+</sup>19, BHOS19, BO20, BBE<sup>+</sup>20, BBO20, BBKO21]. A highlight of this line of work certainly are the lower bounds proven in [BBH<sup>+</sup>19], where it is shown that the number of rounds required to compute a maximal matching is at least  $\Omega(\min\{\Delta, \log n / \log \log n\})$  deterministically and  $\Omega(\min\{\Delta, \log \log n / \log \log \log n\})$  in the randomized LOCAL model, where  $\Delta$  is the largest degree of  $G$ . Because a maximal matching of a graph  $G$  is an MIS of the line graph of  $G$ , the same bounds also hold for computing an MIS. Since an MIS can be computed deterministically in  $O(\Delta + \log^* n)$  rounds [BEK14], both bounds are tight as a function of  $\Delta$ . In this paper we prove new tight lower bounds for a large family of distributed graph problems. As one special case of our general lower bound result, we obtain that the same lower bounds that were proven in [BBH<sup>+</sup>19] for MIS in general graphs also hold for computing an MIS in regular trees, which is an exponential improvement over the existing lower bounds [BBO20, BBKO21].

**New Technique.** There are some problems that behave in a special way in the round elimination framework: if we compute the problem that is exactly one round easier, we obtain the problem itself. We call these problems *fixed points*. Such a behavior is not contradictory and such problems actually exist; in fact, proving that a problem is a fixed point is a powerful tool: we know that being a fixed point implies that the problem requires  $\Omega(\log_{\Delta} n)$  rounds for deterministic algorithms, and  $\Omega(\log_{\Delta} \log n)$  rounds for randomized ones, in the LOCAL model. More generally, fixed points provide a conceptually simple method to prove lower bounds: if we can show that a relaxed version of a given problem<sup>2</sup> is a fixed point, then we directly obtain a lower bound for the original problem. Understanding whether all problems that have a deterministic complexity of  $\Omega(\log_{\Delta} n)$  rounds admit a relaxation to some fixed point is a major open question. Clearly, fixed point relaxations cannot exist for problems that can be solved deterministically in  $o(\log_{\Delta} n)$  rounds, and in particular they cannot exist for problems that, for some function  $f$ , can be solved in  $O(f(\Delta) + \log^* n)$  rounds.

Despite this apparent incompatibility, we show in this work that, perhaps surprisingly, fixed points can also be used for proving lower bounds for problems that can be solved in  $O(f(\Delta) + \log^* n)$  rounds. Problems of this kind are difficult to understand in the context of round elimination: typically, when we apply a round elimination step to such a problem, we obtain a problem with a description that is exponentially larger than the description of the original problem. Hence, if we apply multiple round elimination steps, we obtain a problem sequence where the description size grows like a power tower, and this makes it extremely hard to characterize this sequence, or even merely determine which problems in the sequence are 0-round solvable. We show that we can sometimes avoid this growth by relaxing the problem of interest by *embedding a fixed point* in it, and this allows us to prove lower bounds for a large family of problems whose behavior was too complex to deal with previously. In the following we provide a high-level explanation of this concept where, for simplicity, we restrict attention to one specific problem of this family—the MIS problem.

By applying round elimination to the MIS problem for  $k - 1$  times, we obtain a problem  $\Pi_k$  that can be described by two main parts,  $\Pi_k^{\text{good}}$  and  $\Pi_k^{\text{bad}}$ . Part  $\Pi_k^{\text{good}}$  allows two types of nodes, colored ones and uncolored ones. The requirements are that uncolored nodes must have a colored neighbor, colored nodes have a color from a palette of size  $k$ , and the graph induced by colored nodes must be colored properly (note that the MIS problem is equivalent to  $\Pi_1 = \Pi_1^{\text{good}}$ ). Part  $\Pi_k^{\text{bad}}$  allows some

---

<sup>2</sup>A problem  $\Pi'$  is a relaxation of  $\Pi$  if any solution for  $\Pi$  is also a solution for  $\Pi'$ .

other types of outputs that cannot be easily understood. The number of “types of outputs” of this form is roughly a power tower of height  $k$ .

The above directly implies an  $O(\Delta + \log^* n)$ -round upper bound for MIS: a  $(\Delta + 1)$ -coloring of the nodes (which can be found in  $o(\Delta) + O(\log^* n)$  rounds [FHK16]) is a solution to  $\Pi_{\Delta+1}^{\text{good}}$  and therefore also to the more relaxed  $\Pi_{\Delta+1}$ , which in turn has a complexity of precisely  $\Delta$  rounds less than  $\Pi_1 = \text{MIS}$ . In order to prove that this complexity is tight by providing a matching lower bound, we need to show that none of the outputs allowed by  $\Pi_k^{\text{bad}}$  actually contribute to making the problem easier, but this is extremely challenging, since the number of outputs allowed by  $\Pi_k^{\text{bad}}$  is a power tower of height  $k$ .

In this work, we show how to prevent  $\Pi_k^{\text{bad}}$  from appearing altogether in the sequence of problems obtained by round elimination. Let  $\Pi_k^{\text{rel}}$  be the problem obtained by replacing the  $k$ -coloring part of  $\Pi_k^{\text{good}}$  with a fixed point relaxation of the  $k$ -coloring problem (for which we prove the existence). Note that  $\Pi_k^{\text{rel}}$  is at least as easy as  $\Pi_k^{\text{good}}$ . We show that, by applying the round elimination technique to  $\Pi_k^{\text{rel}}$ , we obtain a problem that is at least as hard as  $\Pi_{k+1}^{\text{rel}}$ . This allows us to obtain a sequence of problems of length  $\Omega(\Delta)$ , where each problem is at least one round easier than the previous one, the last problem (and all previous ones) is not 0-round solvable, and  $\Pi_k^{\text{bad}}$  does not appear at all.

In other words, the reason why we obtain  $\Pi_k^{\text{bad}}$  is that the  $\Delta$ -coloring problem is hidden in the sequence of problems obtained by starting from the MIS problem, and we find it fascinating that the hardness of a hard problem (the relaxed  $\Delta$ -coloring problem) that requires  $\Omega(\log n)$  rounds on bounded-degree graphs can be used to prove lower bounds for problems that can be solved in time  $O(f(\Delta) + \log^* n)$  for some function  $f$ , i.e., with a much better  $n$ -dependency. More generally, we believe that our new technique of embedding the  $\Omega(\log n)$ -round hardness in the form of a suitable fixed point into a problem of interest provides a promising (perhaps even generic) way to tackle the main challenge of round elimination—finding a method to make the obtained problem sequences understandable. After all, on an intuitive level, the particular property of fixed points—being invariant, and in particular non-growing, under round elimination—perfectly fits the objective of keeping the problem sequence in check by restricting the growth of the problem descriptions, and it seems plausible to us that this is precisely why we obtain several tight and close-to-tight lower bounds.

**New  $\Delta$ -Coloring Fixed Point.** At the core of our argument is a new relaxation of  $k$ -coloring, which we prove to be a fixed point in the round elimination framework in  $\Delta$ -regular trees as long as  $k \leq \Delta$ . That is, if we perform a round elimination step for this relaxed  $k$ -coloring problem, then we arrive at the same relaxed  $k$ -coloring problem. In the relaxed  $k$ -coloring problem, every node  $v$  is allowed to be colored with a subset  $A_v$  of the  $k$  colors. Nodes that choose more than one color can mark some of their edges as “*don’t care*” edges. Specifically, a node  $v$  that chooses  $|A_v| = a_v$  different colors is allowed to mark  $a_v - 1$  of its edges as “*don’t care*” edges. For every edge  $\{u, v\}$  of the graph, if  $u$  or  $v$  mark  $\{u, v\}$  as “*don’t care*” edge, then there are no constraints for this edge. Otherwise, the color sets  $A_u$  and  $A_v$  of  $u$  and  $v$  must be disjoint. Since this problem is a round elimination fixed point for  $k = \Delta$ , we directly get an alternative proof for the result of [BFH<sup>+</sup>16, CKP19] that  $\Delta$ -coloring  $\Delta$ -regular trees requires  $\Omega(\log_{\Delta} n)$  rounds for deterministic and  $\Omega(\log_{\Delta} \log n)$  rounds for randomized distributed algorithms. Further, apart from proper  $\Delta$ -coloring, the described relaxed  $\Delta$ -coloring problem is also a relaxation of many arbdefective coloring problems.

**Arbdefective Coloring.** Arbdefective colorings were introduced in [BE11] and they have since then become an indispensable tool in the development of efficient distributed coloring algorithms, e.g., [BE11, Bar16, FHK16, Kuh20, MT20, GK21]. An arbdefective  $c$ -coloring of a graph  $G = (V, E)$

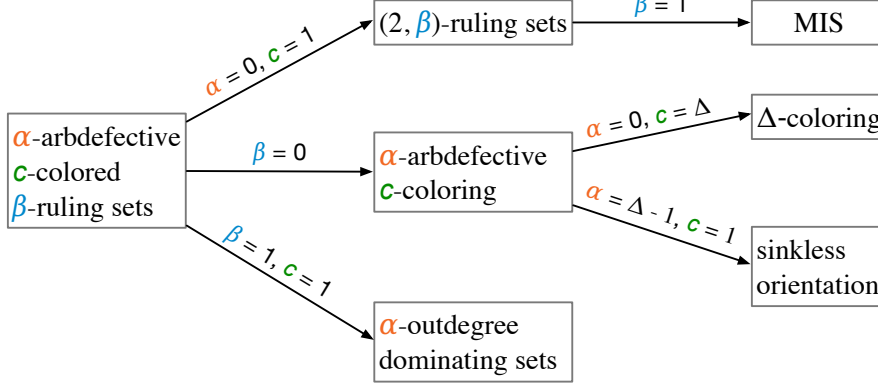


Figure 1: Relation between (a simplified version of) our problem family and concrete problems.

is a coloring of the nodes  $V$  with  $c$  colors together with an orientation of the edges  $E$ . The arbdefect of a node  $v$  of color  $x$  is the number of outneighbors of  $v$  of color  $x$ . Consider the set of nodes that are colored with some set  $A$  of colors in the above relaxed  $k$ -coloring problem. The specification for the nodes with color set  $A$  corresponds exactly to the specification of a color class with arbdefect at most  $|A| - 1$ . If we want to compute a  $c$ -coloring of a  $\Delta$ -regular tree such that for each of the colors  $x \in \{1, \dots, c\}$ , we have a bound  $d_x \geq 0$  on the arbdefect of the nodes with color  $x$ , then our lower bound on the relaxed  $\Delta$  coloring problem directly shows that as long as  $\sum_{x=1}^c (d_x + 1) \leq \Delta$ , this arbdefective coloring problem requires  $\Omega(\log_{\Delta} n)$  rounds deterministically and  $\Omega(\log_{\Delta} \log n)$  rounds with randomization. We also show that as soon as  $\sum_{x=1}^c (d_x + 1) \geq \Delta + 1$ , the problem can be solved in  $O(\Delta + \log^* n)$  rounds by a simple distributed greedy algorithm. In other words, we *exactly* characterize when the arbdefective coloring problem is “easy” and when it is “hard”.

**A General Problem Family.** In this work, we show that our new technique can be used to prove lower bounds for a large family of problems that contains many natural relaxations of MIS and coloring problems. For these problems we prove tight lower bounds as a function of  $\Delta$  (modulo assuming an appropriate initial coloring of the nodes) and improved lower bounds as a function of  $n$ . In this way, we obtain a unified lower bound proof for many problems for which (possibly) worse lower bounds were already known [BFH<sup>+</sup>16, CKP19, BBO20, BBKO21], as well as lower bounds for new problems. For a relation between (a simplified version of) our problem family and various well-known problems, please see Figure 1.

We next describe a simplified variant of our problem family. A full and formal definition of the problem family appears in Section 5. In the simplified variant, a problem has three integer parameters  $\alpha \geq 0$ ,  $\beta \geq 0$ , and  $c \geq 1$ , and we call these problems  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling sets. For a graph  $G = (V, E)$ , the objective is to compute a set  $S \subseteq V$  of nodes together with an  $\alpha$ -arbdefective  $c$ -coloring of the induced subgraph  $G[S]$ . Moreover, all nodes  $v \in V \setminus S$  must be within distance at most  $\beta$  from a node in  $S$  (see Figure 2 for an example).

For  $\beta = 0$ , we obtain the  $\alpha$ -arbdefective  $c$ -coloring problem. Starting from arbdefective coloring, if we set  $\alpha = 0$  and  $c = \Delta$ , we obtain the  $\Delta$ -coloring problem. If we instead set  $\alpha = \Delta - 1$  and  $c = 1$ , we obtain the sinkless orientation problem. For  $\alpha = 0$  and  $c = 1$ , we obtain the  $(2, \beta)$ -ruling set problem, and by setting  $\beta = 1$  we obtain the MIS problem. By setting  $\beta = 1$  and  $c = 1$ , we obtain the  $\alpha$ -outdegree dominating set problem, previously defined and discussed in [BBKO21].

**Main results.** While our results apply to a large family of problems, we summarize our main results in the following. For precise formal statements, we refer to Sections 1.2, 8, and 9.

Our first contribution is related to the problem of computing ruling sets, which are used as a subroutine in many fundamental problems. For these problems, we obtain tight lower bounds as a function of  $\Delta$ , improving exponentially on the results presented in [BBO20]. Also, we obtain improved lower bounds as a function of  $n$ . Our lower bounds hold not only on general graphs, but even on trees, and it turns out that our lower bounds as a function of  $n$  are tight on trees for deterministic algorithms. This in particular implies that, for the MIS problem on trees, we now know its exact deterministic complexity as a function of  $n$ .

Our second contribution regards the arbdefective coloring problem, which is a powerful subroutine in many efficient coloring algorithms. We exactly characterize which variants are easy, in the sense that they can be solved in  $O(f(\Delta) + \log^* n)$  rounds, and which of them are hard, in the sense that they require  $\Omega(\log_\Delta n)$  rounds. Our third contribution consists in introducing a promising new technique, i.e., using fixed points for proving lower bounds for problems that are not in  $\Omega(\log_\Delta n)$ .

Finally, our approach provides a unified proof for various results—both new and already present in the literature—by showing lower bounds for a problem family that interpolates between coloring problems, orientation problems, and independent set problems. The tightness of the obtained results shows that considering these types of problems as special cases of the presented broader problem family is a very useful approach for understanding their complexities, and we hope that similar generalizations will lead to tight bounds for other problems.

## 1.1 Problem Definitions

We next formally define some natural problems for which we prove upper and lower bounds. As discussed above, the full general (and somewhat technical) family of problems for which we prove our lower bounds is formally defined in Section 5. In the following, for a positive integer  $k$ , we use  $[k]$  to denote the set  $[k] := \{1, \dots, k\}$ .

**Generalized Arbdefective Colorings.** We define a generalized notion of arbdefective colorings. We first define the notion of an *arbdefect vector*, which specifies the requirements that a node has to satisfy if it is given a certain color. If there are  $C$  colors available (for some integer  $C \geq 1$ ), an arbdefect vector is a vector of length  $C$ , where the  $i^{\text{th}}$  coordinate indicates the allowed arbdefect of a node having color  $i$ . Concretely, an arbdefect vector  $\vec{d}$  of length  $C$  has the following form  $\vec{d} = (d_1, \dots, d_C) \in \mathbb{N}_0^C$ . A *generalized arbdefective coloring* of a graph is now defined as follows.

**Definition 1.1** (Generalized Arbdefective Coloring). Let  $G = (V, E)$  be a graph,  $C \geq 1$  an integer, and  $\vec{d} = (d_1, \dots, d_C) \in \mathbb{N}_0^C$  an arbdefect vector of length  $C$ . An assignment  $\varphi : V \rightarrow [C]$  of colors in  $[C]$  to the nodes  $V$  is called a  $\vec{d}$ -arbdefective  $C$ -coloring w.r.t. a given orientation of the edges  $E$  of  $G$  if for every node  $v \in V$ , if  $v$  has color  $\varphi(v) = x$ , then  $v$  has at most  $d_x$  outneighbors of color  $x$ .

If we have an arbdefect vector  $\vec{d} = (d_1, \dots, d_C)$  with  $d_i = \alpha$  for all  $i \in [C]$ , we call the corresponding coloring an  $\alpha$ -arbdefective  $C$ -coloring. Note that this is the standard notion of arbdefective coloring as defined in [BE11]. When writing that an algorithm computes an arbdefective coloring of a graph  $G$ , we always implicitly mean that the algorithm computes a coloring of the nodes of  $G$  and a corresponding orientation of the edges of  $G$ . We define the *capacity*  $\text{cap}(\vec{d})$  of an arbdefect vector  $\vec{d}$  as  $\text{cap}(\vec{d}) := \sum_{i=1}^C (d_i + 1)$ .

As discussed above, computing a  $\vec{d}$ -arbdefective coloring in a  $\Delta$ -regular tree is as hard as computing a  $\Delta$ -coloring as long as  $\text{cap}(\vec{d}) \leq \Delta$  and it can be computed by a simple greedy algorithm if  $\text{cap}(\vec{d}) > \Delta$ . For a value  $\sigma > 0$  and graphs of maximum degree  $\Delta$ , we say that an arbdefect vector  $\vec{d}$  is  $\sigma$ -relaxed iff  $\text{cap}(\vec{d}) > \sigma \cdot \Delta$ . Hence, a given generalized arbdefective coloring instance can be solved by a greedy algorithm iff the given arbdefect vector is 1-relaxed.




Figure 2: An example of 2-arbdefective 2-colored 2-ruling set. Nodes in blue and orange are in the set, and they all have arbdefect at most 2. White nodes are not in the set, and there is at least one node in their radius-2 neighborhood that is in the set.

**Arbdefective Colored Ruling Sets.** We next define a family of problems that can be seen both as variants of arbdefective colorings and of ruling sets, and they provide a way to interpolate between them. Figure 2 provides an illustration of the following definition.

**Definition 1.2** ( $\alpha$ -Arbdefective  $c$ -Colored  $\beta$ -Ruling Set). Let  $G = (V, E)$  be a graph, and  $\alpha \geq 0$ ,  $c \geq 1$ ,  $\beta \geq 0$  be integers. A set  $S \subseteq V$  together with an orientation of the edges between nodes in  $S$  is an  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling set if it satisfies the following:

- The induced subgraph  $G[S]$  is colored with an  $\alpha$ -arbdefective  $c$ -coloring.
- For all  $v \in V \setminus S$ , there is a node  $u \in S$  at distance  $\leq \beta$ . Note that if  $\beta = 0$ , then  $S = V$ .

## 1.2 Our Results

Our main lower bound result applies to an entire family of problems for which we gave a high level description before Section 1.1 and that is defined formally in Section 5. We now present the most important results that we obtain as corollaries for the most natural variants of our general problem family (see Table 1 for a partial summary of our results and a comparison with the state of the art).

We first consider arbdefective colorings. We prove that we can exactly characterize which variants of arbdefective coloring can be solved in  $O(f(\Delta) + O(\log^* n))$  for some function  $f$ , and which of them require  $\Omega(\log_\Delta n)$  for deterministic algorithms and  $\Omega(\log_\Delta \log n)$  for randomized ones. In particular, we show the following (follows from Theorem 9.2)

**Theorem 1.3.** *The  $\vec{d}$ -arbdefective  $c$ -coloring problem can be solved in  $O(\Delta + \log^* n)$  in the deterministic LOCAL model if  $\vec{d}$  is 1-relaxed, while otherwise it requires  $\Omega(\log_\Delta n)$  rounds for deterministic algorithms and  $\Omega(\log_\Delta \log n)$  rounds for randomized algorithms.*

By allowing the same arbdefect on each color, we obtain the following corollary.

**Corollary 1.4.** *The  $d$ -arbdefective  $c$ -coloring problem can be solved in  $O(\Delta + \log^* n)$  rounds in the deterministic LOCAL model if  $(d + 1)c > \Delta$ , while if  $(d + 1)c \leq \Delta$  then it requires  $\Omega(\log_\Delta n)$  rounds for deterministic algorithms and  $\Omega(\log_\Delta \log n)$  rounds for randomized algorithms.*



Problem		Complexity	Holds on trees	Holds on general graphs	Ref.	
$(2, \beta)$ -ruling set, as a function of $n$ , for $\beta$ small enough	det.	$O(\log^5 n)$	✓	✓	[GGR21]	
		$O\left(\frac{\log n}{\beta \log \log n}\right)$	✓		[BE10]	
		$\Omega\left(\sqrt{\frac{\log n}{\beta \log \log n}}\right)$	✓	✓	[BBO20]	
		$\Omega\left(\frac{\log n}{\beta \log \log n}\right)$	✓	✓	<b>new</b>	
	rand.	$O(\beta \log^{\frac{1}{\beta+1}} n)$	✓		[BEPS16] [Gha16]	
		$\Omega\left(\sqrt{\frac{\log \log n}{\beta \log \log \log n}}\right)$	✓	✓	[BBO20]	
$\Omega\left(\frac{\log \log n}{\beta \log \log \log n}\right)$		✓	✓	<b>new</b>		
$(2, \beta)$ -ruling set, as a function of $\Delta$ , for $\beta$ small enough, given an $O(\Delta)$ -coloring	det.	$O(\beta \Delta^{1/\beta})$	✓	✓	[SEW13]	
	det. &	$\Omega\left(\frac{\log \Delta}{\beta \log \log \Delta}\right)$	✓	✓	[BBO20]	
	rand.	$\Omega(\beta \Delta^{1/\beta})$	✓	✓	<b>new</b>	
$\alpha$ -arbdefective $c$ -coloring	$(\alpha + 1)c > \Delta$	det.	$O(\Delta + \log^* n)$	✓	✓	easy
		det.	$\Omega(\log_{\Delta} n)$	✓	✓	easy
	$(\alpha + 1)^2 c \leq \varepsilon_0 \Delta$	rand.	$\Omega(\log_{\Delta} \log n)$	✓	✓	easy
		det.	$\Omega(\log_{\Delta} n)$		✓	easy
	$(\alpha + 1)c \leq \varepsilon_0 \Delta$	rand.	$\Omega(\log_{\Delta} \log n)$		✓	easy
		det.	$\Omega(\log_{\Delta} n)$	✓	✓	<b>new</b>
$\alpha$ -arbdefective $c$ -colored $\beta$ -ruling set, as a function of $\Delta$ , for $\beta$ small enough, given an $\alpha$ -arbdef. $O(\frac{\Delta}{\alpha+1})$ -coloring	$(\alpha + 1)c \leq \Delta$	rand.	$\Omega(\log_{\Delta} \log n)$	✓	✓	<b>new</b>
		det.	$\Omega(\log_{\Delta} n)$	✓	✓	<b>new</b>
	$(\alpha + 1)c > \Delta$	det.	$O\left(\beta \left(\frac{\Delta}{(\alpha+1)c}\right)^{1/\beta}\right)$	✓	✓	<b>new</b>
		det. &	$\Omega\left(\beta \left(\frac{\Delta}{(\alpha+1)c}\right)^{1/\beta}\right)$	✓	✓	<b>new</b>
		rand.				

Table 1: Some of our results and how they compare with the state of the art. Whenever a condition contains  $\varepsilon_0$ , then  $\varepsilon_0$  has to be chosen as a sufficiently small positive constant. The term “easy” indicates that the results follow by applying standard techniques.

We now focus on ruling sets and some of their variants. Consider a variant of ruling sets where nodes are required to prove that there is a node in the set at distance at most  $\beta$ . For this variant, it was shown in [BBO20] that, given a proper  $c$ -coloring, one can solve the problem in  $t$  rounds where  $t$  is the minimum value such that  $\binom{t+\beta}{\beta} \geq c$ . We prove that this upper bound is tight, assuming that a proper  $(\Delta + 1)$ -coloring is given as input to the nodes. In particular, we show that this problem requires at least  $t + 1$  rounds, where  $t$  is the maximum value such that  $\binom{t+\beta}{\beta} < \Delta$ .

For the standard variants of ruling sets, which do not require to prove that there is a node in the set at distance at most  $\beta$ , we lose an additive  $\beta$  term in the lower bound, but we still obtain tight results for  $\beta \leq \varepsilon_0 \log \Delta$  for some constant  $\varepsilon_0 > 0$ . In particular, these results exponentially improve the lower bounds in [BBO20], and are asymptotically tight in the case in which an  $O(\Delta)$ -coloring is provided to the nodes. Notice that a stronger lower bound for the case where we are not given an  $O(\Delta)$ -coloring, would directly imply an  $\omega(\log^* n)$  lower bound for  $O(\Delta)$ -vertex coloring (and thus also for  $(\Delta + 1)$ -vertex coloring), which is a long-standing open question.

Our lower bound does not only hold for ruling sets, but for the more general  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling sets that we defined in Section 1.1. While the case of  $\beta = 0$  has been handled by Corollary 1.4, in the remaining cases the following holds.

**Theorem 1.5.** *Let  $t$  be the maximum value such that  $\binom{t+\beta}{\beta} < \frac{\Delta}{c(\alpha+1)}$ . The  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling set problem requires  $\Omega(\min\{t, \log_{\Delta} n\} - \beta)$  in the deterministic LOCAL model and  $\Omega(\min\{t, \log_{\Delta} \log n\} - \beta)$  in the randomized LOCAL model.*

As proven in Corollary 8.8 (see also Table 1), Theorem 1.5 implies that for sufficiently small  $\beta$ , computing an  $\alpha$ -arbdefective  $c$ -coloring  $\beta$ -ruling set requires  $\Omega(\min\{\beta(\frac{\Delta}{c(1+\alpha)})^{1/\beta}, \log_{\Delta} n\})$  rounds deterministically and  $\Omega(\min\{\beta(\frac{\Delta}{c(1+\alpha)})^{1/\beta}, \log_{\Delta} \log n\})$  rounds with randomization. For details and additional implications of the above theorem, we refer to Section 8. The theorem in particular also implies that computing an MIS requires  $\Omega(\frac{\log n}{\log \log n})$  rounds in the deterministic LOCAL model, even in trees. Notice that this result is tight, as in [BE10] it has been shown that MIS on trees can be solved in  $O(\frac{\log n}{\log \log n})$  deterministic rounds. The following theorem shows that our general lower bound for arbdefective colored ruling sets is tight if we are initially given an appropriate coloring as input.

**Theorem 1.6.** *Let  $G = (V, E)$  be a graph and let  $\alpha \geq 0$ ,  $c \geq 1$ ,  $\beta \geq 1$ , and  $C \geq c$ . If an  $\alpha$ -arbdefective  $C$ -coloring of  $G$  is given, then an  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling set of  $G$  can be computed deterministically in the minimum time  $t$  such that  $\binom{t+\beta}{\beta} \geq \lceil \frac{C}{c} \rceil$  in the LOCAL model.*

Theorem 1.6 implies that if an  $\alpha$ -arbdefective  $C$ -coloring of  $G$  is given as input, then an  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling set of  $G$  can be computed deterministically in  $O(\beta(\frac{C}{c})^{1/\beta})$  rounds (cf. Corollary 9.4). By combining with the arbdefective coloring algorithm given by Lemma 9.3, this implies that an  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling set can be computed in  $O(\beta(\frac{C}{(\alpha+1)\sqrt{c}})^{2/(\beta+1)})$  rounds in the deterministic LOCAL model (cf. Corollary 9.5). We refer the reader to Sections 8 and 9 for a full list of results that we obtain.

### 1.3 Additional Related Work

In the following, we list additional relevant previous work. In light of the vast amount of related work on the techniques and problems that we study in this paper, we restrict the discussion to the results that are most relevant in the context of the present paper.

**Round Elimination.** Round elimination has proven to be an extremely useful tool for showing lower bounds in the distributed setting, [BFH<sup>+</sup>16, CKP19, CHL<sup>+</sup>20, Bra19, BHOS19, Oli19, BBH<sup>+</sup>19, BO20, BBO20, BBKO21, BBE<sup>+</sup>20, FP20]. Although we started to widely understand the technique only recently, already Linial’s classic  $\Omega(\log^* n)$  lower bound for 3-coloring a cycle [Lin92] can be seen as a round elimination proof. The clean concept of round elimination, as we use it nowadays, has been introduced in 2016 in [BFH<sup>+</sup>16], where the authors prove a randomized lower bound of  $\Omega(\log_\Delta \log n)$  rounds for computing a sinkless orientation or a  $\Delta$ -coloring. Shortly after, Chang, Kopelowitz, and Pettie [CKP19] lifted the lower bound to an  $\Omega(\log_\Delta n)$ -round lower bound for deterministic algorithms. Later, the lower bound was also extended to the  $(2\Delta - 2)$ -edge coloring problem in [CHL<sup>+</sup>20].

In [Bra19], Brandt further refined the technique by giving an *automatic* way to perform round elimination. That is, given a problem  $\Pi$ , he showed how to mechanically construct a problem  $\Pi'$  that is exactly one round easier than  $\Pi$ . The new automatic round elimination technique, together with a tool by Olivetti [Oli19] that implements it, gave rise to a better understanding of several important problems in distributed computing. As a highlight, [BBH<sup>+</sup>19] used the technique to prove asymptotically tight-in- $\Delta$  lower bounds for computing a maximal matching (MM) and as an immediate corollary also for computing an MIS. [BO20] further improved these results by showing that MM on bipartite 2-colored graphs requires exactly  $2\Delta - 1$  rounds. While the lower bounds for MM hold on trees, the MIS lower bounds only hold on more general graphs. Moreover, the above lower bounds cannot be directly extended to ruling sets, since it is known that ruling sets on the line graph are much easier than MIS [KMW18].

Balliu, Brandt, and Olivetti [BBO20] made progress on this question and used the automatic round elimination technique to show the first  $\omega(\log^* n)$  lower bound for ruling sets and the first such lower bound for MIS on trees. Unfortunately, their lower bounds are exponentially worse than the best known upper bounds: they for example show that, on trees, MIS requires  $\Omega(\log \Delta / \log \log \Delta)$  rounds. Subsequently, [BBKO21] showed that by making use of an input edge coloring, one can apply the round elimination technique to improve and simplify the lower bound for MIS on trees, obtaining a lower bound of  $\Omega(\log \Delta)$  rounds. Unfortunately, this simpler proof does not extend to ruling sets, but the authors show that it can be used to prove lower bounds for a class of problems called bounded outdegree dominating sets (that are the same as arbdefective 1-colored 1-ruling sets).

**Maximal Independent Set.** The problem of computing an MIS is one of the most fundamental symmetry breaking problems and it has been extensively studied in the LOCAL model (for example, see [AGLP89, KW85, Lub86, ABI86, Lin92, Nao91, PS96, KMW16, BE10, SW10b, LW11, BEPS16, BE13, BEK14, Gha16, RG20, GGR21, CPS17, Gha19, GP19]). Already in the 80s, the works of Luby [Lub86] and Alon, Babai, and Itai [ABI86] showed that with randomization, an MIS can be computed in  $O(\log n)$  rounds. The first non-trivial deterministic upper bounds for MIS were also designed in the late 80s, when [AGLP89] introduced a  $2^{O(\sqrt{\log n \log \log n})}$ -round LOCAL algorithm for a more general tool, called network decomposition. Later, Panconesi and Srinivasan [PS96] improved this runtime to  $2^{O(\sqrt{\log n})}$ , and very recently the deterministic complexity of computing a network decomposition (and an MIS) has been improved to polylogarithmic in  $n$  by Rozhoň and Ghaffari [RG20]. An improved variant of this algorithm by Ghaffari, Grunau, and Rozhoň [GGR21] provides the current best distributed deterministic MIS algorithm with a round complexity of  $O(\log^5 n)$ . In addition to analyzing the complexity of computing an MIS solely as a function of  $n$ , there is also a large body of work that studies runtimes expressed as  $f(\Delta) + g(n)$ . In 2008, Barenboim, Elkin, and Kuhn [BEK14] showed that MIS can be computed deterministically in  $O(\Delta + \log^* n)$  deterministic rounds in the LOCAL model. Furthermore, the authors presented randomized algorithms with a round complexity of  $O(\log^2 \Delta) + 2^{O(\sqrt{\log \log n})}$ , which was improved to  $O(\log \Delta) + 2^{O(\sqrt{\log \log n})}$

by Ghaffari [Gha16]. As a result of the network decomposition algorithm of [RG20, GGR21], the complexity has been further improved to  $O(\log \Delta + \log^5 \log n)$ .

On the lower bound side, we have known since the late 80s and early 90s that MIS requires  $\Omega(\log^* n)$  rounds, even for randomized algorithms, which is due to the works of Linial [Lin92] and Naor [Nao91]. The first  $\omega(\log^* n)$  lower bound for MIS was shown in 2004 by Kuhn, Moscibroda, and Wattenhofer [KMW16], who proved a lower bound of  $\Omega\left(\min\left\{\frac{\log \Delta}{\log \log \Delta}, \sqrt{\frac{\log n}{\log \log n}}\right\}\right)$  rounds, even for randomized algorithms. Recently, this lower bound has been improved and complemented by the round elimination based maximal matching lower bound of [BBH<sup>+</sup>19], which implies that computing an MIS requires  $\Omega\left(\min\left\{\Delta, \frac{\log n}{\log \log n}\right\}\right)$  rounds for deterministic algorithms and  $\Omega\left(\min\left\{\Delta, \frac{\log \log n}{\log \log \log n}\right\}\right)$  rounds for randomized algorithms. As discussed, these results do not hold in trees and the first lower bounds for MIS in trees were recently proven in [BBO20, BBKO21]. In this paper, we prove that the same MIS lower bounds that were proven in [BBH<sup>+</sup>19] for general graphs also hold for computing MIS in trees.

The MIS problem has also been studied in specific class of graphs and in particular for trees (e.g., [BE10, LW11, BEPS16, Gha16]). For example, on trees, barenboim and Elkin [BE10] showed that MIS can be solved in  $O(\log n / \log \log n)$  deterministic rounds, and Ghaffari [Gha16] showed that it can be solved in  $O(\sqrt{\log n})$  randomized rounds. As a function of  $\Delta$  and  $n$ , MIS on trees can be solved in  $O(\log \Delta + \log \log n / \log \log \log n)$  randomized rounds [BEPS16, Gha16].

**Ruling Sets.** Ruling sets were introduced in the late 80s in [AGLP89], where the authors used them as a subroutine in their network decomposition algorithm. Since then, there have been several works that studied the complexity of ruling sets (see, e.g., [AGLP89, SW10a, GV07, SEW13, BKP14, BEPS16, Gha16, KMW18, KP12, PPP<sup>+</sup>17, HKN16]), and that used ruling sets as a subroutine for solving problems of interest (see, e.g., [BEPS16, Gha16, GHKM18, CLP18]). [AGLP89] in particular showed that a  $(2, O(\log n))$ -ruling set can be computed in time  $O(\log n)$ . Moreover, as shown by Schneider, Elkin, and Wattenhofer [SEW13], the algorithm of [AGLP89] can be adapted to deterministically compute, for general  $\beta \geq 1$ , a  $(2, \beta)$ -ruling set in time  $O(\beta \Delta^{2/\beta} + \log^* n)$  (or in time  $O(\beta C^{1/\beta})$  if an initial  $C$ -coloring is given).

On the randomized side, we know that  $(2, \log \log n)$ -ruling sets can be computed in  $O(\log \log n)$  randomized rounds, which follows by combining the works of Gfeller and Vicari [GV07] and Schneider, Elkin, and Wattenhofer [SEW13]. Further improvements were obtained in [KP12, BKP14, BEPS16, Gha16]. In combination with the network decomposition result of [RG20, GGR21], the algorithm of Ghaffari [Gha16] implies that a  $(2, \beta)$ -ruling set can be computed in  $O(\beta \log^{1/\beta} \Delta) + \text{poly}(\log \log n)$  rounds in the randomized LOCAL model. On trees, a  $(2, \beta)$ -ruling set can be computed deterministically in time  $O\left(\frac{\log n}{\beta \log \log n}\right)$  by combining a coloring algorithm of Barenboim and Elkin [BE10] with the  $O(\beta C^{1/\beta})$ -round algorithm for a given  $C$ -coloring. With randomization, in trees, it is possible to compute a  $(2, \beta)$ -ruling set in time  $O(\beta \log^{\frac{1}{\beta+1}} n)$  by combining techniques of [BEPS16] with techniques of [Gha16].

On the lower bound side, it was known since the late 80s and early 90s that computing a  $(2, \beta)$ -ruling set requires  $\Omega(\log^* n)$  rounds up to some  $\beta \in \Theta(\log^* n)$ , even for randomized algorithms, and this directly follows from the MIS lower bounds of [Lin92] and [Nao91] for paths and rings and it therefore also holds on trees. The only previous  $\omega(\log^* n)$  lower bounds for ruling sets were obtained in [BBO20]. There, the authors showed that for sufficiently small  $\beta$ , any deterministic algorithm that computes a  $(2, \beta)$ -ruling set requires  $\Omega\left(\min\left\{\frac{\log \Delta}{\beta \log \log \Delta}, \log_{\Delta} n\right\}\right)$  rounds, while any randomized algorithm requires  $\Omega\left(\min\left\{\frac{\log \Delta}{\beta \log \log \Delta}, \log_{\Delta} \log n\right\}\right)$  rounds.

**Arbdefective Colorings.** The arbdefective coloring problem was introduced by Barenboim and Elkin [BE11], and it has proved to be a very useful tool that is used in several state-of-the-art distributed algorithms for computing a proper coloring of a graph [BE11, Bar16, FHK16, MT20, GK21, Kuh20]. The aforementioned paper shows that in graphs of arboricity  $a$ , a  $b$ -arbdefective  $O(a/b)$ -coloring can be deterministically computed in  $O(\min\{a, (a/b)^2\} \log n)$  rounds. The paper uses this to recursively decompose a graph into sparser subgraphs, which can then be colored more efficiently. Subsequently, Barenboim [Bar16] gave an algorithm that computes an  $O(\alpha)$ -arbdefective  $O(\Delta \log \Delta/\alpha)$ -coloring in  $O(\Delta \log^2 \Delta/\alpha + \log^* n)$  rounds. Barenboim [Bar16] and Fraigniaud, Heinrich, and Kosowski [FHK16] used this algorithm to obtain deterministic  $(\Delta + 1)$ -list-coloring algorithms with a round complexity that is sublinear in  $\Delta$ . Barenboim, Elkin, and Goldenberg [BEG18] improved on the arbdefective coloring of [Bar16] by achieving an  $\alpha$ -arbdefective coloring with  $O(\Delta/\alpha)$ -coloring in time  $O(\Delta/\alpha + \log^* n)$  (the same result was later also proved in different ways in [GK21] and [Mau21]).

## 2 Road Map

**Preliminaries.** We start, in Section 3, by providing some preliminaries. In particular, we define the model of computing, we define the class of *locally checkable* problems, and we formally define the round elimination framework.

**A Fixed Point for  $\Delta$ -coloring.** In Section 4, we prove an  $\Omega(\log_\Delta n)$ -round deterministic and  $\Omega(\log_\Delta \log n)$ -round randomized lower bound for the  $\Delta$ -coloring problem in the LOCAL model. These results are already known from prior work [CKP19, BFH<sup>+</sup>16]. The novelty of our proof lies on the fact that we apply the round elimination technique directly to a more natural generalization of  $\Delta$ -coloring. The previous proofs are also based on round elimination, however the relaxation of  $\Delta$ -coloring (which is called sinkless coloring) in the proofs of [CKP19, BFH<sup>+</sup>16] is more related to the sinkless orientation problem than to  $\Delta$ -coloring. Note that the primary goal of [BFH<sup>+</sup>16] is to prove a lower bound for the sinkless orientation problem. While the proof of [BFH<sup>+</sup>16] also implies that  $\Delta$ -coloring requires  $\Omega(\log_\Delta \log n)$  randomized rounds, the proof does not reveal enough of the structure of  $\Delta$ -coloring to be used to prove the lower bounds for MIS or ruling sets. As already described in Section 1, we define a relaxation of  $\Delta$ -coloring, which we prove to be a non-trivial fixed point under the round elimination framework. This proof can be seen as a special case of the more general proof that we will show afterwards. We also explain why a fixed point proof for the  $\Delta$ -coloring problem is helpful in proving lower bounds for problems that are *much easier* than the  $\Delta$ -coloring problem. In particular, we show that, one reason of why it is hard to prove lower bounds for, e.g., ruling sets, is because the  $c$ -coloring problem, for  $c \leq \Delta$ , appears as a subproblem of ruling sets when applying the round elimination technique. The  $c$ -coloring problem does not behave nicely under the round elimination framework, and this makes also ruling sets to not behave nicely under this framework. Relaxing  $c$ -coloring to a non-trivial fixed point makes it behave much better, and the same then helps for ruling sets as well.

**A Problem Family.** In Section 5 we formally define the family of problems for which we prove lower bounds. We also show that this family of problems is rich enough to contain relaxations of many interesting natural problems:  $\Delta$ -coloring, arbdefective colorings, maximal independent set, ruling sets, and arbdefective colored ruling sets. We also show that this family of problems contains even more problems: in general, it gives a way to interpolate between colorings and ruling sets.

**Lower Bound in the Port Numbering Model.** In Section 6, we are going to use the round elimination technique to prove lower bounds for the problems of the family in the port numbering model (a model that is strictly weaker than LOCAL). In particular, we are going to show that some problems in this family are strictly easier than others. Then, by showing that there exists a sequence of  $T$  problems, all non-zero round solvable, such that each problem in the sequence is strictly harder than the next problem in the sequence, we directly get that the first problem of the sequence requires at least  $T$  rounds. This would give us a lower bound for the port numbering model, but what we really want is a lower bound for the LOCAL model. An important technical detail that we will need in order to be able to lift such a result from the port numbering model to the LOCAL model is the following: we need to be able to describe each problem of the sequence with a small number of labels. Hence, in this section we also show that the number of labels of each problem is small enough.

**Lifting Port Numbering Lower Bounds to the LOCAL model.** In Section 7, we show that a long sequence of problems satisfying the previously described requirements directly gives a lower bound for the LOCAL model. Here we use known results, but we strengthen them a bit in order to support a larger number of labels, compared to the one used in previous works.

**Lower Bound Results.** In Section 8, we formalize the connection between the problems in the family and some natural problems. We show that the generic lower bound obtained by combining the results of Section 6 and Section 7 implies lower bounds for many natural problems. For example, in this section we show that the  $(2, \beta)$ -ruling set problem requires  $\Omega(\beta\Delta^{1/\beta})$  rounds, unless  $\Delta$  (as a function of  $n$ ) is very large.

**Upper Bounds.** While the fixed point that we provide in Section 4 is a relaxation of the  $\Delta$ -coloring problem, we prove in Section 8 that it is also a relaxation of some variants of arbdefective coloring. We show that this fixed point exactly characterizes which variants of arbdefective coloring require  $\Omega(\log_\Delta n)$  rounds. In particular, in Section 9 we show that all the variants of arbdefective coloring for which we do not obtain a lower bound in Section 8 can be solved in  $O(\Delta + \log^* n)$ . In Section 9 we also provide tight upper bounds for  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling sets.

**Open Problems.** We conclude, in Section 10, with some open problems. As we have seen, fixed point based proofs for hard problems may help in proving lower bounds for much easier problems, so one natural question is whether such a fixed point based proof exists for all problems that require  $\Omega(\log_\Delta n)$  deterministic rounds in the LOCAL model.

## 3 Preliminaries

### 3.1 LOCAL model

The model of distributed computing considered in this paper is the well-known LOCAL model. In this model, the computation proceeds in synchronous rounds, where, at each round, nodes exchange messages with neighbors and perform some local computation. The size of the messages and the computational power of each node are not bounded. Each node in an  $n$ -node graph has a unique identifier (ID) in  $\{1, \dots, \text{poly } n\}$ . Initially, each node knows its own ID, its own degree, the maximum degree  $\Delta$  in the graph, and the total number of nodes in the graph. Nodes execute the same distributed algorithm, and upon termination, each node outputs its local output (e.g., a color). Nodes correctly solve a distributed problem if and only if their local outputs form a correct global solution (e.g., a proper coloring of the graph). The complexity of an algorithm that solves a problem in the

LOCAL model is measured as the number of rounds it takes such that all nodes give their local output and terminate. Since the size of the messages is not bounded, each node can share with its neighbors all what it knows so far, and hence a  $T$  round algorithm in this model can be seen as a mapping of  $T$ -hop neighborhoods into local outputs. In the randomized version of the LOCAL model, in addition, each node has access to a stream of random bits. The randomized algorithms that we consider are Monte Carlo ones, that is, a randomized algorithm of complexity  $T$  must always terminate within  $T$  rounds and the result it produces must be correct with high probability, i.e., with probability at least  $1 - 1/n$ . The LOCAL model is quite a strong model of computation, hence lower bounds in this model widely apply on other weaker models as well (such as, on the well-studied CONGEST model of distributed computation, where the size of the messages is bounded by  $O(\log n)$  bits).

**Port Numbering Model.** For technical reasons, in order to show our lower bounds for the LOCAL model, we first show lower bounds for the weaker *port numbering* (PN) model, and then we lift these results to the LOCAL model. As in the LOCAL model, in the PN model the size of the messages and the computational power of a node are unbounded. Differently from the LOCAL model, in the PN model nodes do not have IDs. Instead, nodes are equipped with a port numbering, that is, each node has assigned a port numbering to all its incident edges. If we denote with  $\deg(v)$  the degree of a node  $v$ , then each incident edge of node  $v$  has a number, or a *port*, in  $\{1, \dots, \deg(v)\}$  assigned in an arbitrary way, such that for any two incident edges it holds that their port number is not the same. For technical reasons, we will also assume that edges are equipped with a port numbering in  $\{1, 2\}$ , that is, for each edge, we have an arbitrary assignment of port numbers to each of its endpoints, such that, endpoints of the same edge have different port numbers. This essentially results in an arbitrary consistent orientation of the edges. Notice that this technical detail makes the PN model only stronger, hence our lower bounds directly apply on the classical PN model as well.

**Special Port Numbering.** There have been some cases in previous works where, in order to show lower bounds, the authors have made use of a specific input given to the graph (note that this would only make the task of finding lower bounds potentially harder). For example, [BBKO21] showed lower bounds for MIS and out-degree dominating sets on trees in the case where we are given a  $\Delta$ -edge coloring in input. Hence, for the sake of being as general as possible, it is convenient to show that some of our main statements (such as [Theorem 7.1](#)) hold also when the ports satisfy some constraints. More precisely, consider a hypergraph where nodes have degree  $\Delta$  and hyperedges have rank  $\delta$ . Let  $C^n$  and  $C^h$  be two sets, where  $C^n \subseteq \{1, \dots, \delta\}^\Delta$  and  $C^h \subseteq \{1, \dots, \Delta\}^\delta$ . The set  $C^n$  is the *node port numbering* constraint, and restricts the possible ports that hyperedges incident to the nodes are allowed to have. The set  $C^h$  is the *hyperedge port numbering* constraint, and restricts the possible ports that nodes incident to the hyperedges are allowed to have. The node port constraint and the hyperedge port constraint are defined as follows.

- Node port constraint  $C^n$ : Consider a node  $v$ , and let  $h_i$  be the hyperedge connected to port  $i$  of node  $v$ , for  $1 \leq i \leq \Delta$ . Let  $p(h_i, v)$  be the hyperedge port that connects  $h_i$  to  $v$ , then  $(p(h_1, v), \dots, p(h_\Delta, v))$  must be contained in  $C^n$ .
- Hyperedge port constraint  $C^h$ : Consider a hyperedge  $h$  and let  $v_i$  be the node connected to port  $i$  of hyperedge  $h$ , for  $1 \leq i \leq \delta$ . Let  $p(v_i, h)$  be the node port that connects  $v_i$  to  $h$ , then  $(p(v_1, h), \dots, p(v_\delta, h))$  must be contained in  $C^h$ .

If  $C^n$  and  $C^h$  contain all possible tuples, we say that the port numbering is *unconstrained*.

As already mentioned, some specific constrained port assignments can encode inputs, and an example is  $\Delta$ -edge coloring. Consider a graph  $G$ . We can define  $C^n$  and  $C^h$  in such a way that they

imply a  $\Delta$ -edge coloring of  $G$ . Let  $C^n$  be unconstrained, and let  $C^h = \{(i, i) \mid 1 \leq i \leq \Delta\}$ . In this case, if port  $i$  of node  $v$  is connected to an edge  $e = \{u, v\}$ , then the port of  $u$  connected to  $e$  is also  $i$ . Hence, we can interpret this port assignment as a coloring, where edge  $e$  has color  $i$ . We may assume that also in the LOCAL model we have a port numbering assignment, and this assignment may satisfy some constraints  $C^{\text{port}}$ . Lower bounds obtained in this setting are even stronger.

### 3.2 Black-White Formalism

In this work, we prove lower bounds on graphs, and we provide a general theorem that is useful for lifting lower bounds from the port numbering model to the LOCAL model. In order to make such a theorem as general as possible, we let it support also problems defined on *hypergraphs*. In particular, we consider problems defined in the black-white formalism, that has already been used in the literature to formally define so-called *locally checkable problems* on these kinds of graphs.

In this formalism, a problem  $\Pi = \Pi_{\Delta, \delta}$  is described by a triple  $(\Sigma_{\Pi}, \mathcal{N}_{\Pi}, \mathcal{E}_{\Pi})$ , where  $\Sigma_{\Pi}$  is a set of labels, and  $\mathcal{N}_{\Pi}$  and  $\mathcal{E}_{\Pi}$  are sets of tuples of size  $\Delta$  and  $\delta$ , respectively, where each element of the tuples is a label from  $\Sigma_{\Pi}$ . That is,  $\mathcal{N}_{\Pi} \subseteq (\Sigma_{\Pi})^{\Delta}$ , and  $\mathcal{E}_{\Pi} \subseteq (\Sigma_{\Pi})^{\delta}$ .  $\mathcal{N}_{\Pi}$  and  $\mathcal{E}_{\Pi}$  are called *node constraint* and *hyperedge constraint*, respectively. In some cases, we may also refer to them as *white constraint* and *black constraint*, respectively.

Given a hypergraph  $G = (V, E)$ , and the set of node-hyperedge pairs  $P = \{(v, e) \in V \times E \mid v \in e\}$ , solving  $\Pi$  means assigning to each element of  $P$  an element from  $\Sigma_{\Pi}$  such that:

- Let  $v$  be any node of degree exactly  $\Delta$ . Let  $L = (\ell_1, \dots, \ell_{\Delta})$  be the labels assigned to the node-hyperedge pairs incident to  $v$ . A permutation of  $L$  must be in  $\mathcal{N}_{\Pi}$ .
- Let  $e$  be any hyperedge of rank exactly  $\delta$ . Let  $L = (\ell_1, \dots, \ell_{\delta})$  be the labels assigned to the node-hyperedge pairs incident to  $e$ . A permutation of  $L$  must be in  $\mathcal{E}_{\Pi}$ .

The complexity of  $\Pi$  is the minimum number of rounds required for nodes to produce a correct output.

A special case of the black-white formalism is called node-edge formalism, and is given by considering  $\delta = 2$ . Note that, in this case, there must be two labels assigned to each edge, one for each endpoint. Since any hypergraph can be seen as a 2-colored bipartite graph, and vice versa, then the black-white formalism allows us to define problems on 2-colored bipartite graphs as well. In this case, we need to assign a label to each edge, and we have constraints on the labeling of the edges incident on white nodes of degree  $\Delta$ , and on edges incident on black nodes of degree  $\delta$ .

We use regular expressions to concisely describe the allowed configurations of a problem. For example, we can write  $(W[XY]Z) \mid Z^3$  to denote the constraint  $\{(W, X, Z), (W, Y, Z), (Z, Z, Z)\}$ . As a shorthand, we may also write this constraint as  $\{WXZ, WYZ, ZZZ\}$ . Recall that the order does not matter, and hence this constraint also allows the configuration  $WZX$ . We call parts of regular expressions of the form  $[XY]$  *disjunctions*, and while expressions of the form  $W[XY]Z$  denote sets of configurations, we may refer to them as *condensed configurations*. If a specific configuration can be obtained by choosing a label in each disjunction, then we say that the configuration is *contained* in the condensed configuration. For example, the configuration  $WXZ$  is contained in the condensed configuration  $W[XY]Z$ .

The black-white formalism is powerful enough to be able to encode many problems for which the solution can be checked locally, i.e., in  $O(1)$  rounds. In this formalism, only nodes of degree  $\Delta$  and hyperedges of degree  $\delta$  are constrained, while other nodes and hyperedges are not required to satisfy the constraints, but note that proving lower bounds in this restricted setting only makes the lower bounds stronger. We may refer to all problems that can be expressed in this formalism as *locally checkable* problems. We now provide some examples of locally checkable problems.



**Example: Sinkless Orientation in 2-Colored Graphs.** Assume we have a 2-colored graph. We require nodes of degree  $\Delta$  to solve the sinkless orientation problem: nodes must orient their edges, and every node of degree  $\Delta$  must have at least one outgoing edge. Nodes of smaller degree are unconstrained. We can define  $\Sigma_{\Pi} = \{\mathbf{B}, \mathbf{W}\}$ . Since the graph is 2-colored, then every edge has a black and a white endpoint. An edge labeled  $\mathbf{B}$  denotes that the edge is oriented towards the black node, while an edge labeled  $\mathbf{W}$  denotes that the edge is oriented towards the white node. Hence, we can express the requirement of having at least one outgoing edge as follows. The white constraint  $\mathcal{N}_{\Pi}$  is defined as  $\mathbf{B}[\mathbf{B}\mathbf{W}]^{\Delta-1}$ , while the black constraint  $\mathcal{E}_{\Pi}$  is defined as  $\mathbf{W}[\mathbf{B}\mathbf{W}]^{\Delta-1}$ . In other words, white nodes must have at least one edge labeled  $\mathbf{B}$ , meaning that it is oriented towards a black neighbor, and then all other edges are unconstrained. Similarly, black nodes must have at least one edge labeled  $\mathbf{W}$ .

**Example: Maximal Independent Set.** Assume we have a graph, where we require nodes of degree  $\Delta$  to solve the maximal independent set problem. In order to encode this problem, one may try to use two labels, one for nodes in the set, and the other for nodes not in the set. Unfortunately, this is not possible, as with only two labels it is not possible to encode the maximality requirement [BBE<sup>+</sup>20]. Hence, we use three labels, and we define  $\Sigma_{\Pi} = \{\mathbf{M}, \mathbf{P}, \mathbf{U}\}$ . We define the node constraint  $\mathcal{N}_{\Pi}$  as  $\mathbf{M}^{\Delta} \mid \mathbf{P}\mathbf{U}^{\Delta-1}$ , and the edge constraint  $\mathcal{E}_{\Pi}$  as  $\mathbf{M}[\mathbf{U}\mathbf{P}] \mid \mathbf{U}\mathbf{U}$ .

In other words, nodes can be part of the independent set and output  $\mathbf{M}^{\Delta}$ , or can be outside of the independent set, and in this case they have to prove that they have at least one neighbor in the independent set, by outputting  $\mathbf{P}\mathbf{U}^{\Delta-1}$ . The edge constraint only allows  $\mathbf{P}$  to be paired with an  $\mathbf{M}$ , and hence nodes outputting  $\mathbf{P}\mathbf{U}^{\Delta-1}$  must have at least one neighbor outputting  $\mathbf{M}^{\Delta}$ . Also, nodes of the independent set cannot be neighbors, because  $\mathbf{M}\mathbf{M}$  is not allowed. Then, every other configuration is allowed by  $\mathcal{E}_{\Pi}$ .

### 3.3 Round Elimination

The round elimination technique works as follows. Assume we are given a problem  $\Pi_1$ , for which we want to prove a lower bound. We show that, given  $\Pi_i$ , we can construct a problem  $\Pi_{i+1}$  that is at least one round easier than  $\Pi_i$ , unless  $\Pi_i$  is already 0 rounds solvable. If we are able to perform this operation for  $T$  times, and prove that for all  $i \leq T$ , the problem  $\Pi_i$  is not 0-rounds solvable, then this implies that  $\Pi_1$  is at least  $T - 1$  rounds harder than  $\Pi_T$ . Since  $\Pi_T$  is not 0-rounds solvable, this implies that  $\Pi_1$  requires at least  $T$  rounds.

In order to define this sequence of problems, we exploit [Bra19, Theorem 4.3], that gives a mechanical way to define  $\Pi_{i+1}$  as a function of  $\Pi_i$ , such that if  $\Pi_i$  requires exactly  $T$  rounds, then  $\Pi_{i+1}$  requires exactly  $\max\{T - 1, 0\}$  rounds. In its general form, this theorem works for any locally checkable problem defined in the black-white formalism, and given a problem  $\Pi$  with complexity  $T$ , it first constructs an intermediate problem  $\Pi'$ , and then it constructs a problem  $\Pi''$  that can be proved to have complexity  $T - 1$ .

We denote by  $\mathcal{R}(\cdot)$  the procedure that can be applied on a problem  $\Pi$  to obtain the intermediate problem  $\Pi'$ , and by  $\overline{\mathcal{R}}(\cdot)$  the procedure that can be applied on  $\Pi'$  to obtain the problem  $\Pi''$ . Note that  $\Pi'' = \overline{\mathcal{R}}(\mathcal{R}(\Pi))$ . In [Bra19],  $\mathcal{R}(\Pi = (\Sigma_{\Pi}, \mathcal{N}_{\Pi}, \mathcal{N}_{\Pi})) = \Pi' = (\Sigma_{\Pi'}, \mathcal{N}_{\Pi'}, \mathcal{E}_{\Pi'})$  is defined as follows.

- $\mathcal{E}_{\Pi'}$  is defined as follows. Let  $C$  be the maximal set such that for all  $(L_1, \dots, L_{\delta}) \in C$  it holds that, for all  $i$ ,  $L_i \in 2^{\Sigma_{\Pi}} \setminus \{\emptyset\}$ , and for all  $(\ell_1, \dots, \ell_{\delta}) \in L_1 \times \dots \times L_{\delta}$ , it holds that a permutation of  $(\ell_1, \dots, \ell_{\delta})$  is in  $\mathcal{E}_{\Pi}$ . In this case, we say that  $(L_1, \dots, L_{\delta})$  satisfies the universal quantifier. The set  $\mathcal{E}_{\Pi'}$  is obtained by removing all *non-maximal* configurations from  $C$ , that is, all configurations  $(L_1, \dots, L_{\delta}) \in C$  such that there exists another configuration  $(L'_1, \dots, L'_{\delta}) \in C$

and a permutation  $\phi$ , such that  $L_i \subseteq L'_{\phi(i)}$  for all  $i$ , and there exists at least one  $i$  such that the inclusion is strict.

- $\Sigma_{\Pi'} \subseteq 2^{\Sigma_{\Pi}}$  contains all the sets that appear at least once in  $\mathcal{E}_{\Pi'}$ .
- $\mathcal{N}_{\Pi'}$  is defined as follows. It contains all configurations  $(L_1, \dots, L_{\Delta})$  such that, for all  $i$ ,  $L_i \in \Sigma_{\Pi'}$ , and there exists  $(\ell_1, \dots, \ell_{\Delta}) \in L_1 \times \dots \times L_{\Delta}$ , such that that a permutation of  $(\ell_1, \dots, \ell_{\Delta})$  is in  $\mathcal{N}_{\Pi}$ . In this case, we say that  $(L_1, \dots, L_{\Delta})$  *satisfies the existential quantifier*.

We may refer to the computation of  $\mathcal{E}_{\Pi'}$  as *applying the universal quantifier*, and to the computation of  $\mathcal{N}_{\Pi'}$  as *applying the existential quantifier*. When applying the universal quantifier, we may use the term *by maximality* to argue that, in order for a configuration to be maximal, some label must appear in a set of labels. The hard part in computing  $\Pi'$  is applying the universal quantifier. In fact, there is an easy way to compute  $\mathcal{N}_{\Pi'}$ , that is the following. Start from all the configurations allowed by  $\mathcal{N}_{\Pi}$ , and for each configuration add to  $\mathcal{N}_{\Pi'}$  the condensed configuration obtained by replacing each label  $y$  by the disjunction of all label sets in  $\Sigma_{\Pi'}$  that contain  $y$ .

The operator  $\overline{\mathcal{R}}(\cdot)$  is defined in the same way, but the roles of  $\mathcal{N}$  and  $\mathcal{E}$  are reversed, that is, the universal quantifier is applied on  $\mathcal{N}_{\Pi'}$ , and the existential quantifier is applied on  $\mathcal{E}_{\Pi'}$ .

In [Bra19, Theorem 4.3], the following is proved.

**Theorem 3.1** ([Bra19], rephrased). *Let  $T > 0$ . Consider a class  $\mathcal{G}$  of hypergraphs with girth at least  $2T + 2$ , and some locally checkable problem  $\Pi$ . Then, there exists an algorithm that solves problem  $\Pi$  in  $\mathcal{G}$  in  $T$  rounds if and only if there exists an algorithm that solves  $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$  in  $T - 1$  rounds.*

Note that, in order to prove lower bounds, when computing  $\mathcal{E}_{\Pi'}$ , it is not necessary to show that any possible choice over the sets allowed by  $\mathcal{E}_{\Pi'}$  is a configuration allowed by  $\mathcal{E}_{\Pi}$ , but it is sufficient to show that there are no maximal configurations that satisfy the universal quantifier that are not in  $\mathcal{E}_{\Pi'}$ . In other words, when proving lower bounds, we can add allowed configurations to  $\mathcal{E}_{\Pi'}$ , as this is only making  $\Pi'$  easier. Hence, we introduce the notion of a relaxation of a configuration (already introduced in [BBO20]).

**Definition 3.2.** Consider two node configurations  $\mathcal{A} = A_1 \dots A_{\Delta}$  and  $\mathcal{B} = B_1 \dots B_{\Delta}$ , where the  $A_i$  and  $B_i$  are sets of labels from some label space. We say that  $\mathcal{A}$  *can be relaxed to  $\mathcal{B}$*  if there exists a permutation  $\rho: \{1, \dots, \Delta\} \rightarrow \{1, \dots, \Delta\}$  such that  $A_i \subseteq B_{\rho(i)}$  for all  $1 \leq i \leq \Delta$ . For simplicity, if not indicated otherwise, we will assume that the sets in the two configurations are ordered such that  $\rho$  is the identity function. We define relaxations of edge configurations analogously.

Given two constraints  $C$  and  $C'$ , if all the configurations allowed by  $C$  can be relaxed to configurations in  $C'$ , we say that  $C'$  is a *relaxation* of  $C$ . Also, given two problems  $\Pi_1 = (\Sigma_1, \mathcal{N}_1, \mathcal{E}_1)$  and  $\Pi_2 = (\Sigma_2, \mathcal{N}_2, \mathcal{E}_2)$ , if  $\mathcal{N}_2$  is a relaxation of  $\mathcal{N}_1$  and  $\mathcal{E}_2$  is a relaxation of  $\mathcal{E}_1$ , then we say that  $\Pi_2$  is a relaxation of  $\Pi_1$ .

**Example: Sinkless Orientation in 2-Colored Graphs.** Recall that this problem can be expressed as follows.

$$\begin{aligned}\mathcal{N}_{\Pi} &= \mathsf{B}[\mathsf{BW}]^{\Delta-1} \\ \mathcal{E}_{\Pi} &= \mathsf{W}[\mathsf{BW}]^{\Delta-1}\end{aligned}$$

We start by computing  $\mathcal{E}_{\Pi'}$ . We need to satisfy the universal quantifier, and this implies that it must not be possible to pick  $\mathsf{B}$  from  $\Delta$  different sets, since we would obtain  $\mathsf{B}^{\Delta}$ , that is not allowed by  $\mathcal{E}_{\Pi}$ . But we can also notice that, as long as we are forced to pick  $\mathsf{W}$  from at least one set, then

everything else is allowed. Hence, the universal quantifier is satisfied for all configurations of the form  $\{W\}[\{\{W\}\{B\}\{W, B\}]^{\Delta-1}$ . If we discard non-maximal configurations, we obtain  $\{W\}\{W, B\}^{\Delta-1}$ .

We can now rename the sets to make it easier to read the constraint, by using the following mapping:

$$\begin{aligned}\{W\} &\rightarrow W \\ \{W, B\} &\rightarrow B\end{aligned}$$

We obtain  $\mathcal{E}_{\Pi'} = WB^{\Delta-1}$ . In order to apply the existential quantifier on  $\mathcal{N}_{\Pi}$ , it is enough to replace each label with the disjunction of all the sets that contain that label. Note that  $B$  is only contained in the new label  $\{W, B\}$ , that is renamed to  $B$ , while  $W$  is contained in both  $\{W\}$  and  $\{W, B\}$ , that are renamed to  $W$  and  $B$ , respectively. Hence,  $\mathcal{N}_{\Pi'} = \{W, B\}[\{\{W\}\{W, B\}]^{\Delta-1} = B[BW]^{\Delta-1}$ . Summarizing, we obtain the following:

$$\begin{aligned}\Sigma_{\Pi'} &= \{B, W\} \\ \mathcal{E}_{\Pi'} &= WB^{\Delta-1} \\ \mathcal{N}_{\Pi'} &= B[BW]^{\Delta-1}\end{aligned}$$

**Relation Between Labels.** Consider a problem  $\Pi$  and let  $C$  be one of its constraints (either node or hyperedge) containing tuples of size  $d$ . It can be useful to relate the labels of  $\Pi$ , according to how easier it is to use a label instead of another, according to  $C$ .

Consider two arbitrary labels  $X, Y$  of  $\Pi$ . We say that  $X$  is *at least as strong as*  $Y$  according to  $C$ , if the following holds. For all configurations  $(\ell_1, \dots, \ell_i, \dots, \ell_d) \in C$  such that there exists an  $i$  such that  $\ell_i = Y$ , a permutation of the configuration  $(\ell_1, \dots, X, \dots, \ell_d)$  is in  $C$ . That is, for all configurations in  $C$  containing  $Y$ , if we replace an arbitrary amount of  $Y$  with  $X$ , then we obtain a configuration that is still in  $C$ . If the constraint is clear from the context, we may omit it and just say that  $X$  is at least as strong as  $Y$ . If  $X$  is at least as strong as  $Y$ , we also say that  $Y$  is *at least as weak as*  $X$ , or  $Y \leq X$ . If  $X$  is at least as strong as  $Y$ , but  $Y$  is not at least as strong as  $X$ , then  $X$  is *stronger* than  $Y$ , and  $Y$  is *weaker* than  $X$ , that we may also denote as  $Y < X$ . Intuitively, stronger labels are easier to use, if we only consider the requirements of  $C$ , because we can always replace a weaker label by a stronger one.

It can be useful to illustrate the strength of the labels by using a directed graph. The *diagram* of  $\Pi$  according to  $C$  is a directed graph where nodes are labels of  $\Pi$ , and there is an edge  $(Y, X)$  if and only if:

- $X \neq Y$ ;
- $X$  is at least as strong as  $Y$ ;
- there is no label  $Z$  different from  $X$  and  $Y$  that is stronger than  $Y$  and weaker than  $X$ .

In other words, we put edges connecting weaker labels to stronger labels, but we omit relations that can be obtained by transitivity. Note that the obtained graph is not necessarily connected, but it is directed acyclic, unless  $\Pi$  contains two labels of equal strength. We may refer to the diagram according to  $\mathcal{N}$  as the *node diagram*, and to the diagram according to  $\mathcal{E}$  as the *hyperedge diagram* (or *edge diagram*, in the case of graphs). The *successors* of some label  $\ell$  w.r.t. some constraint  $C$  are defined to be all labels that are at least as strong as  $\ell$  according to  $C$ , or in other words all labels reachable from  $\ell$  in the diagram according to  $C$ .

When applying the round elimination technique, and in particular when performing the universal quantifier, it will be useful to consider the diagram of the constraint on which the universal quantifier

is applied to. As we will see later, this diagram limits the possible labels appearing in the new constraint. As an example, [Figure 3](#) depicts the edge diagram of the MIS problem, that is, the diagram according to  $\mathcal{E}_\Pi = \text{M}[\text{UP}] \mid \text{UU}$ . There is an edge from P to U because the only edge configuration containing P is MP, and if we replace P with U, we obtain MU, which is also allowed by  $\mathcal{E}_\Pi$ . Conversely, there is no edge from U to P, because UU is allowed, but by replacing an arbitrary amount of U with P we obtain configurations not in  $\mathcal{E}_\Pi$ .




Figure 3: The edge diagram of the MIS problem: U is stronger than P, and there is no relation between labels M and P, and between M and U.

**Additional Notation.** We now introduce some additional notation that can be useful to concisely describe a set of labels. Given a set  $\{\ell_1, \dots, \ell_p\} \subseteq \Sigma_\Pi$ , we denote by  $\langle \ell_1, \dots, \ell_p \rangle$  the set that contains all labels  $\ell$  in  $\Sigma_\Pi$  satisfying that there exists an  $i$  such that  $\ell$  is at least as strong as  $\ell_i$ . In other words,  $\langle \ell_1, \dots, \ell_p \rangle$  contains the labels  $\ell_1, \dots, \ell_p$ , plus all labels that are reachable from at least one of them in the diagram of  $\Pi$  (i.e., all the successors of  $\ell_1, \dots, \ell_p$ ).

In order to properly define  $\langle \cdot \rangle$ , we need to specify to which constraint the strength relation refers to. If the argument of  $\langle \cdot \rangle$  contains labels of some problem  $\Pi$  on which we are applying the operator  $\mathcal{R}(\cdot)$ , then  $\langle \cdot \rangle$  refers to its hyperedge constraint, while if the argument of  $\langle \cdot \rangle$  contains labels of an intermediate problem  $\Pi'$ , that is, a problem on which we are applying the operator  $\overline{\mathcal{R}}(\cdot)$ , then  $\langle \cdot \rangle$  refers to its node constraint. In other words, we always consider the constraint on which we apply the universal quantifier. That is, if we start from  $\Pi$  and compute  $\Pi' = \mathcal{R}(\Pi)$ , then  $\langle \cdot \rangle$  refers to  $\mathcal{E}_\Pi$ , while when computing  $\Pi'' = \overline{\mathcal{R}}(\Pi')$ , then  $\langle \cdot \rangle$  refers to  $\mathcal{N}_{\Pi'}$ .

When computing  $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$  it will be useful to use expressions such as  $\langle\langle \ell \rangle\rangle$ , where  $\ell$  is a label of  $\Pi$ . This expression represents a set of sets of labels from  $\Sigma_\Pi$ , where  $\langle \ell \rangle$  is computed according to  $\mathcal{E}_\Pi$ , and  $\langle\langle \ell \rangle\rangle$  is then taken according to  $\mathcal{N}_{\Pi'}$ .

A set  $S = \{\ell_1, \dots, \ell_p\} \subseteq \Sigma_\Pi$  is called *right-closed* if  $S = \langle \ell_1, \dots, \ell_p \rangle$ . That is,  $S$  is right-closed if and only if for each label  $\ell_i$  contained in  $S$  also all successors of  $\ell_i$  in the diagram (taken w.r.t. the constraint on which the universal quantifier is applied to) are contained in  $S$ . In [\[BBO20\]](#), the following statements have been proved.

**Observation 3.3** ([\[BBO20\]](#)). *Consider an arbitrary collection of labels  $A_1, \dots, A_p \in \Sigma_\Pi$ . If  $\{A_1, \dots, A_p\} \in \Sigma_{\mathcal{R}(\Pi)}$ , then the set  $\{A_1, \dots, A_p\}$  is right-closed (w.r.t.  $\mathcal{E}_\Pi$ ). If  $\{A_1, \dots, A_p\} \in \Sigma_{\overline{\mathcal{R}}(\Pi')}$ , then the set  $\{A_1, \dots, A_p\}$  is right-closed (w.r.t.  $\mathcal{N}_{\Pi'}$ ).*

**Observation 3.4** ([\[BBO20\]](#)). *Let  $U, W \in \Sigma_{\mathcal{R}(\Pi)}$  be two sets satisfying  $U \subseteq W$ . Then  $W$  is at least as strong as  $U$  according to  $\mathcal{N}_{\mathcal{R}(\Pi)}$ . In particular, for any label  $A \in \Sigma_\Pi$  such that  $\langle A \rangle \in \Sigma_{\mathcal{R}(\Pi)}$ , every set  $X \in \Sigma_{\mathcal{R}(\Pi)}$  containing  $A$  is contained in  $\langle\langle A \rangle\rangle$ .*

*Analogous statements hold for  $\overline{\mathcal{R}}(\cdot)$  instead of  $\mathcal{R}(\cdot)$ , by considering the label strength according to  $\mathcal{E}_{\overline{\mathcal{R}}(\cdot)}$ .*

For a set  $S = \{\ell_1, \dots, \ell_p\} \subseteq \Sigma_\Pi$  of labels, we denote by  $\text{disj}(S)$  the disjunction  $[\ell_1 \dots \ell_p]$ . For instance,  $\text{disj}(\langle \ell \rangle)$  is the disjunction of all labels that are at least as strong as  $\ell$ .

## 4 $\Delta$ -Coloring Fixed Point and Its Applications

In this section, we define a problem  $\Pi_\Delta$  that is at least as easy as the  $\Delta$ -coloring problem, and show that  $\Pi_\Delta$  is a non-trivial fixed point under the round elimination framework. It is known that a non-trivial fixed point directly implies an  $\Omega(\log_\Delta n)$  lower bound for deterministic algorithms, and an  $\Omega(\log_\Delta \log n)$  lower bound for randomized algorithms, in the LOCAL model, but we show this formally in Section 7. Then, we also describe the relation of  $\Pi_\Delta$  with the maximal independent set problem.

The problem  $\Pi_\Delta$  is defined as follows. Consider the set of colors  $\mathfrak{C} = \{1, \dots, \Delta\}$ . The label set  $\Sigma_{\Pi_\Delta}$  of  $\Pi_\Delta$  is defined as  $\Sigma_{\Pi_\Delta} = \{\ell(\mathcal{C}) \mid \mathcal{C} \in 2^\mathfrak{C}\}$ . The node constraint  $\mathcal{N}_{\Pi_\Delta}$  of  $\Pi_\Delta$  contains the following allowed configurations:

$$\ell(\mathcal{C})^{\Delta-x} \ell(\emptyset)^x, \text{ for all } \mathcal{C} \in 2^\mathfrak{C} \setminus \{\emptyset\}, \text{ where } x = |\mathcal{C}| - 1.$$

The edge constraint  $\mathcal{E}_{\Pi_\Delta}$  of  $\Pi_\Delta$  contains the following allowed configurations:

$$\ell(\mathcal{C}_1)\ell(\mathcal{C}_2), \text{ for all } \mathcal{C}_1, \mathcal{C}_2 \in 2^\mathfrak{C} \text{ such that } \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset.$$

In Section 4.1, we will prove that  $\Pi_\Delta$  is a fixed point. Note that  $\Pi_\Delta$  is not 0-round solvable, since all configurations allowed by the node constraint contain at least one label  $L = \ell(\mathcal{C}) \neq \ell(\emptyset)$ , and  $LL \notin \mathcal{E}_{\Pi_\Delta}$ . We now provide an explicit example of  $\Pi_\Delta$  where  $\Delta = 3$ , and we give some intuition behind its definition.

**Example:  $\Pi_3$ .** Consider the case where  $\Delta = 3$ , and let us use letters to denote colors, that is,  $\mathfrak{C} = \{A, B, C\}$ . To simplify the notation, we rename each label  $\ell(\{x_1, \dots, x_k\})$  corresponding to a non-empty set of colors as  $\overline{x_1 \dots x_k}$ . Also, we rename  $\ell(\emptyset)$  as  $X$ . For example, the label  $\ell(\{A, B\})$  becomes  $\overline{AB}$ . The problem  $\Pi_3$  is defined as follows:

$\mathcal{N}_{\Pi_3}$ :	$\mathcal{E}_{\Pi_3}$ :
$\overline{A}$ $\overline{A}$ $\overline{A}$	$X$ $[\overline{X A B C} \overline{AB} \overline{BC} \overline{AC} \overline{ABC}]$
$\overline{B}$ $\overline{B}$ $\overline{B}$	$\overline{A}$ $[\overline{X B C} \overline{BC}]$
$\overline{C}$ $\overline{C}$ $\overline{C}$	$\overline{B}$ $[\overline{X A C} \overline{AC}]$
$\overline{AB}$ $\overline{AB}$ $X$	$\overline{C}$ $[\overline{X A B} \overline{AB}]$
$\overline{AC}$ $\overline{AC}$ $X$	$\overline{AB}$ $[\overline{X C}]$
$\overline{BC}$ $\overline{BC}$ $X$	$\overline{AC}$ $[\overline{X B}]$
$\overline{ABC}$ $X$ $X$	$\overline{BC}$ $[\overline{X A}]$
	$\overline{ABC}$ $X$

The intuition behind this problem is the following. We start from the standard definition of 3-coloring, and we add allowed configurations that “reward” nodes that have many neighbors of the same color. For example, if a node  $v$  outputs the configuration  $\overline{AB} \overline{AB} X$ , it means that  $v$  has at least 2 neighbors that do not output neither  $AAA$  nor  $BBB$ . We can think of  $v$  as being colored *both A and B*. Node  $v$  is rewarded for this, and it is allowed to output  $X$  on one of its ports. Label  $X$  can be seen as a wildcard: a label that is compatible with everything. In particular, this means that  $\Pi_3$  can be seen as a variant of arbdefective coloring, where nodes that control many colors are allowed to have a larger arbdefect.

**Relation with MIS.** It is known that, if a problem can be relaxed to a problem that is a fixed point under the round elimination framework, then it requires  $\Omega(\log_{\Delta} n)$  rounds in the LOCAL model for deterministic algorithms, and  $\Omega(\log_{\Delta} \log n)$  rounds for randomized algorithms. Hence, it is perhaps surprising that a fixed point can be helpful to prove lower bounds for problems that can be solved in  $O(f(\Delta) + \log^* n)$  for some function  $f$ , that is, problems that are much easier. We now provide the intuition behind this phenomenon, by considering the case of MIS (similar observations can be made for ruling sets and for arbdefective colored ruling sets).

Already in [BBO20], it has been shown that, when applying the round elimination technique on the MIS problem for many times, we get a problem that looks like the following:

- There is a part of the problem that can be naturally described, that is, after doing  $k$  steps of round elimination, the problem becomes a relaxation of the  $k$ -coloring problem.
- There is a part of the problem that cannot be naturally described, that seems to be just an artifact of the round elimination technique. In particular, after  $k$  steps, the obtained problem also contains what one would obtain by applying the round elimination technique on the  $(k - 1)$ -coloring problem.

Unfortunately, by applying the round elimination technique on coloring problems, we usually obtain problems whose description grows exponentially *at each step* of round elimination. This makes it unfeasible to apply the technique for more than just a few steps. Hence, we need to find a way to avoid this exponential growth, that is, find a relaxation of  $k$ -coloring such that if we apply the technique, we don't get a problem that is much larger. This is exactly what a fixed point for  $k$ -coloring (for  $k \leq \Delta$ ) gives us, since, by applying round elimination on this problem, we obtain the problem itself. Hence, the idea that we use to prove a lower bound for MIS (and other problems) is the following:

- By applying the round elimination technique on MIS, at each step, the number of colors grows by 1, and intuitively this means that the MIS problem requires  $\Omega(\Delta)$ , since  $k$ -coloring is a hard problem, for  $k \leq \Delta$ . To prove this, we need to show that the “unnatural” part of the problem does not help to solve the problem faster.
- The unnatural part would grow exponentially at each step, but we relax it in the same way as we do in the  $k$ -coloring fixed point. In this way, at step  $k + 1$ , we obtain a  $(k + 1)$ -coloring, plus something that can be relaxed to the configurations allowed in the fixed point of the  $(k + 1)$ -coloring problem. This allows us to *concisely* describe the problem obtained at each step.

Hence, the problem obtained by applying the round elimination technique on MIS for  $k$  times can be described as follows: there is a  $k$ -coloring component, plus something that makes the colors grow by 1 at each new step, plus something related to applying the round elimination technique on  $k'$ -coloring for  $k' < k$ . In order to obtain our linear-in- $\Delta$  lower bound for MIS, we essentially embed, into the description of the intermediate problems, a proof that  $k$ -coloring is hard.

#### 4.1 A Fixed Point For $\Delta$ -Coloring

In order to simplify the proofs, instead of showing that  $\Pi_{\Delta}$  is a non-trivial fixed point, we do the following. We define a problem  $\Pi$ , that is essentially equivalent to  $\Pi_{\Delta}$ , where the only difference is that the node constraint of  $\Pi$  is defined in a slightly different way. Then, we show that  $\Pi$  is a non-trivial fixed point. Hence, problem  $\Pi$  is not 0-round solvable, and it is a relaxation of the  $\Delta$ -coloring problem, i.e., a solution for  $\Delta$ -coloring is also a solution for  $\Pi$ , but the reverse does not necessarily hold. Then, we show that, by applying two round elimination steps on  $\Pi$ , we obtain the same problem  $\Pi$ , implying that  $\Pi$  is a non-trivial fixed point.

**Problem  $\Pi$ .** As in the case of  $\Pi_\Delta$ , the label set is defined as  $\Sigma_\Pi = \{\ell(\mathcal{C}) \mid \mathcal{C} \in 2^\mathfrak{C}\}$ , where  $\mathfrak{C} = \{1, \dots, \Delta\}$ . The edge constraint  $\mathcal{E}_\Pi$  is also defined in the same way. It contains all the configurations of the following form:

$$\ell(\mathcal{C}_1)\ell(\mathcal{C}_2), \text{ for all } \mathcal{C}_1, \mathcal{C}_2 \in 2^\mathfrak{C} \text{ such that } \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset.$$

The node constraint  $\mathcal{N}_\Pi$  is defined slightly differently. It contains all the configurations of the following form:

$$\begin{aligned} & \ell(\mathcal{C}_1)\ell(\mathcal{C}_2) \dots \ell(\mathcal{C}_\Delta), \text{ such that} \\ & \exists k \in \{1, \dots, \Delta\} \text{ and } 1 \leq i_1 < i_2 < \dots < i_k \leq \Delta, \text{ such that} \\ & |\mathcal{C}_{i_1} \cap \mathcal{C}_{i_2} \cap \dots \cap \mathcal{C}_{i_k}| \geq \Delta - k + 1. \end{aligned}$$

On the one hand, notice that  $\mathcal{N}_\Pi$  allows strictly more configurations than the node constraint of  $\Pi_\Delta$ . On the other hand, given a solution of  $\Pi_\Delta$ , nodes can solve  $\Pi$  in 0 rounds by outputting the intersection of the  $k$  sets on  $k$  ports, and the empty set on the remaining ones. Hence,  $\Pi$  and  $\Pi_\Delta$  are equivalent. We start by observing some properties of the configurations of  $\Pi$ .

**Observation 4.1.** *If  $\ell(\mathcal{C}_1)\ell(\mathcal{C}_2) \dots \ell(\mathcal{C}_\Delta)$  is a configuration in  $\mathcal{N}_\Pi$ , then also  $\ell(\mathcal{C}'_1)\ell(\mathcal{C}'_2) \dots \ell(\mathcal{C}'_\Delta)$ , where, for all  $1 \leq j \leq \Delta$ ,  $\mathcal{C}_j \subseteq \mathcal{C}'_j$ , is a configuration in  $\mathcal{N}_\Pi$ .*

**Observation 4.2.** *If  $\ell(\mathcal{C}_1)\ell(\mathcal{C}_2)$  is a configuration in  $\mathcal{E}_\Pi$ , then also  $\ell(\mathcal{C}'_1)\ell(\mathcal{C}'_2)$ , where, for all  $1 \leq j \leq 2$ ,  $\mathcal{C}'_j \subseteq \mathcal{C}_j$ , is a configuration in  $\mathcal{E}_\Pi$ .*

**Labels of  $\Pi'$ .** Let  $\Pi' = \mathcal{R}(\Pi)$ . By definition,  $\Sigma_{\Pi'} \subseteq 2^{\Sigma_\Pi}$ , that is,  $\Sigma_{\Pi'}$  consists of a set of sets of labels in  $\Sigma_\Pi$ . Before characterizing the node and edge constraints of  $\Pi'$ , we perform a renaming of the labels of problem  $\Pi'$ , in the following way. Each label set  $\{\ell(\mathcal{C}_1), \dots, \ell(\mathcal{C}_h)\}$  is renamed to label  $\ell(\mathcal{C}) = \ell(\bigcup_{1 \leq j \leq h} \mathcal{C}_j)$ , that is, a label that consists of sets of colors is mapped to a label that consists of the union of the sets of colors. This renaming gives a label-set that contains labels of the form  $\ell(\mathcal{C})$ , where  $\mathcal{C} \in 2^\mathfrak{C}$  and  $\mathfrak{C} = \{1, \dots, \Delta\}$ . Hence, after this renaming,  $\Sigma_{\Pi'} = \Sigma_\Pi$ . Note that different original labels, such that the union of their sets gives the same set, are mapped to the same label.

**Edge Constraint of  $\Pi'$ .** We now characterize the edge constraint  $\mathcal{E}_{\Pi'}$  of  $\Pi'$ . First of all, by definition of the round elimination framework,  $\mathcal{E}_{\Pi'}$  contains all configurations  $L_1 L_2$  such that, for any choice in  $L_1$ , and for any choice in  $L_2$ , we get a configuration that is in  $\mathcal{E}_\Pi$ . Hence, under the renaming that we performed,  $\mathcal{E}_{\Pi'}$  must contain at least all the configurations in  $\mathcal{E}_\Pi$ . We claim that  $\mathcal{E}_{\Pi'}$  does not contain any additional allowed configuration.

Suppose, for a contradiction, that there exists a configuration  $\ell(\mathcal{C}_1)\ell(\mathcal{C}_2)$  in  $\mathcal{E}_{\Pi'}$  such that  $\mathcal{C}_1 \cap \mathcal{C}_2 \neq \emptyset$ . This means that, before the renaming, there exists a configuration  $L_1 L_2$  in  $\mathcal{E}_{\Pi'}$  such that  $L_1$  contains a set  $S_1$  containing  $\ell(\mathcal{C}'_1)$ , and  $L_2$  contains a set  $S_2$  containing  $\ell(\mathcal{C}'_2)$ , such that  $\mathcal{C}'_1 \cap \mathcal{C}'_2 \neq \emptyset$ . Hence, we get that there exists a choice resulting in the configuration  $\ell(\mathcal{C}'_1)\ell(\mathcal{C}'_2)$ , where  $\mathcal{C}'_1 \cap \mathcal{C}'_2 \neq \emptyset$ , which does not satisfy any configuration in  $\mathcal{E}_\Pi$ , which is a contradiction. Hence,  $\mathcal{E}_{\Pi'} = \mathcal{E}_\Pi$ .

**Node Constraint of  $\Pi'$ .** We now characterize the node constraint  $\mathcal{N}_{\Pi'}$  of  $\Pi'$ . As shown in [Section 3](#), the node constraint  $\mathcal{N}_{\Pi'}$  can be constructed in the following way: take all the configurations allowed by the node constraint  $\mathcal{N}_\Pi$ , and replace each label  $\ell(\mathcal{C})$  with the disjunction of all the labels of  $\Pi'$  that, before the renaming, contain  $\ell(\mathcal{C})$ . Note that, after the renaming, the labels that contain  $\ell(\mathcal{C})$  are exactly all the labels  $\ell(\mathcal{C}')$  such that  $\mathcal{C} \subseteq \mathcal{C}'$ . Hence, from each configuration  $\ell(\mathcal{C}_1) \dots \ell(\mathcal{C}_\Delta)$  in  $\mathcal{N}_\Pi$ , we obtain the configuration itself, plus all the configurations  $\ell(\mathcal{C}'_1) \dots \ell(\mathcal{C}'_\Delta)$  such that for all  $1 \leq j \leq \Delta$ ,  $\mathcal{C}_j \subseteq \mathcal{C}'_j$ . By [Observation 4.1](#), these configurations are also in  $\mathcal{N}_\Pi$ . Hence  $\mathcal{N}_{\Pi'} = \mathcal{N}_\Pi$ .

**Labels of  $\Pi''$ .** In order to prove that  $\Pi$  is a fixed point, we show that problem  $\Pi''$ , defined as  $\Pi'' = \mathcal{R}(\Pi') = \overline{\mathcal{R}}(\mathcal{R}(\Pi))$ , after a specific renaming, is the same as problem  $\Pi$ . By definition,  $\Sigma_{\Pi''} \subseteq 2^{\Sigma_{\Pi''}}$ . We rename the labels in  $\Sigma_{\Pi''}$  as follows. Each label set  $\{\ell(\mathcal{C}_1), \dots, \ell(\mathcal{C}_h)\}$  is renamed to label  $\ell(\mathcal{C}) = \ell(\bigcap_{1 \leq j \leq h} \mathcal{C}_j)$ , that is, a label that consists of sets of colors is mapped to a label that consists of the intersection of the sets of colors. This renaming gives a label-set that contains labels of the form  $\ell(\mathcal{C})$ , where  $\mathcal{C} \in 2^{\mathfrak{C}}$  and  $\mathfrak{C} = \{1, \dots, \Delta\}$ . Hence, after this renaming,  $\Sigma_{\Pi''} = \Sigma_{\Pi}$ . Note that different original labels, such that the intersection of their sets gives the same set, are mapped to the same label. We prove that, under this renaming,  $\Pi'' = \Pi$ .

**Node Constraint of  $\Pi''$**  We now characterize the node constraint  $\mathcal{N}_{\Pi''}$  of  $\Pi''$ , and show that it is equal to  $\mathcal{N}_{\Pi}$ . First of all, by definition of the round elimination framework, before renaming,  $\mathcal{N}_{\Pi''}$  contains all and only configurations  $L_1 \dots L_{\Delta}$ , such that, for any choice  $(\ell(\mathcal{C}_1), \dots, \ell(\mathcal{C}_{\Delta})) \in L_1 \times \dots \times L_{\Delta}$ , we get a configuration that is in  $\mathcal{N}_{\Pi'} = \mathcal{N}_{\Pi}$ . We claim that, after the proposed renaming,  $\mathcal{N}_{\Pi''} = \mathcal{N}_{\Pi}$ .

We prove our claim by performing multiple partial renamings. Let  $L_1 \dots L_{\Delta}$  be an arbitrary configuration in  $\mathcal{N}_{\Pi''}$ . Consider an arbitrary label  $L_i = \{\ell(\mathcal{C}_1), \dots, \ell(\mathcal{C}_k)\}$  in the configuration. Create the label  $L'_i = \{\ell(\mathcal{C}_1 \cap \mathcal{C}_2), \ell(\mathcal{C}_3), \dots, \ell(\mathcal{C}_k)\}$ , that is, replace  $\ell(\mathcal{C}_1), \ell(\mathcal{C}_2)$  by their intersection. In the following lemma, we prove that also the configuration obtained by replacing  $L_i$  with  $L'_i$  satisfies that, for any choice over this configuration, we get a configuration that is in  $\mathcal{N}_{\Pi'} = \mathcal{N}_{\Pi}$ . If we repeat this process many times, until every set  $L_i$  contains a single label, then we get the exact same result that we would have obtained by directly replacing each set of labels by their intersection. But this implies that, under the proposed renaming, we get exactly the same configurations that appear in  $\mathcal{N}_{\Pi'}$ , and hence our claim follows.

**Lemma 4.3.** *Let  $L_1 L_2 \dots L_{\Delta}$  be a configuration in  $\mathcal{N}_{\Pi''}$ . Let  $L_t = \{\ell(\mathcal{C}_t^1), \dots, \ell(\mathcal{C}_t^h)\}$  be an arbitrary label in the configuration, and let  $\ell(\mathcal{C}') = \ell(\mathcal{C}_t^1 \cap \mathcal{C}_t^2)$ . If we replace the label set  $L_t$  with  $L'_t = \{\ell(\mathcal{C}'), \ell(\mathcal{C}_t^3), \dots, \ell(\mathcal{C}_t^h)\}$ , we get a configuration that is in  $\mathcal{N}_{\Pi''}$ .*

*Proof.* W.l.o.g., suppose  $t = \Delta$ . Let  $L_{\Delta}^{1 \cap 2} = \{\ell(\mathcal{C}')\}$ , and let  $L_{\Delta}^{3+} = \{\ell(\mathcal{C}_{\Delta}^3), \dots, \ell(\mathcal{C}_{\Delta}^h)\}$ . Since by assumption  $L_1 \dots L_{\Delta-1} L_{\Delta}^{3+}$  already satisfies the universal quantifier of the round elimination framework, we only need to prove that  $L_1 \dots L_{\Delta-1} L_{\Delta}^{1 \cap 2}$  is a configuration in  $\mathcal{N}_{\Pi''}$ , that is, for any choice  $(\ell(\mathcal{C}_1), \dots, \ell(\mathcal{C}_{\Delta-1})) \in L_1 \times \dots \times L_{\Delta-1}$ , we get that  $\ell(\mathcal{C}_1) \dots \ell(\mathcal{C}_{\Delta-1}) \ell(\mathcal{C}')$  is a configuration in  $\mathcal{N}_{\Pi}$ . We prove our claim by contradiction. Hence, assume that there exists a choice  $C' = \ell(\mathcal{C}_1) \dots \ell(\mathcal{C}_{\Delta-1}) \ell(\mathcal{C}')$  that is not a configuration in  $\mathcal{N}_{\Pi}$ .

Let  $\mathcal{C}_{\Delta} = \mathcal{C}_{\Delta}^1$ . By assumption, for any choice in  $L_1 L_2 \dots L_{\Delta}$ , we get a configuration in  $\mathcal{N}_{\Pi}$ . Hence, by the definition of  $\mathcal{N}_{\Pi}$ , we get the following:

$$\begin{aligned} \exists k \in \{1, \dots, \Delta\} \text{ and } 1 \leq i_1 < \dots < i_k \leq \Delta, \text{ such that} \\ |S_1 = \mathcal{C}_{i_1} \cap \dots \cap \mathcal{C}_{i_k}| \geq \Delta - k + 1. \end{aligned}$$

Similarly, if we set  $\mathcal{C}_{\Delta} = \mathcal{C}_{\Delta}^2$ , we get the following:

$$\begin{aligned} \exists k' \in \{1, \dots, \Delta\} \text{ and } 1 \leq i'_1 < \dots < i'_k \leq \Delta, \text{ such that} \\ |S_2 = \mathcal{C}_{i'_1} \cap \dots \cap \mathcal{C}_{i'_k}| \geq \Delta - k' + 1. \end{aligned}$$

Let  $s_1$  be the number of sets  $\mathcal{C}_i$ ,  $i \in \{1, \dots, \Delta - 1\}$ , such that  $S_1 \subseteq \mathcal{C}_i$  and  $S_2 \not\subseteq \mathcal{C}_i$ . Similarly, let  $s_2$  be the number of sets  $\mathcal{C}_{i'}$ ,  $i' \in \{1, \dots, \Delta - 1\}$ , such that  $S_2 \subseteq \mathcal{C}_{i'}$  and  $S_1 \not\subseteq \mathcal{C}_{i'}$ . Also, let  $s$  be the number of labels  $\mathcal{C}_j$ ,  $j \in \{1, \dots, \Delta - 1\}$ , such that  $S_1 \cup S_2 \subseteq \mathcal{C}_j$ .

Consider the configuration  $C = \ell(\mathcal{C}_1) \dots \ell(\mathcal{C}_{\Delta-1}) \ell(\mathcal{C}_{\Delta}^1)$ . Depending on  $\mathcal{C}_{\Delta}^1$ , the number of sets that are supersets of  $S_1$  is either  $s + s_1$ , or  $s + s_1 + 1$ . Hence,  $s + s_1 + 1$  is an upper bound on the



number of sets that are supersets of  $S_1$ . Since, by assumption,  $C$  is a valid configuration of  $\mathcal{N}_\Pi$ , then  $|S_1| \geq \Delta - (s + s_1 + 1) + 1$ . By applying the same reasoning on  $S_2$ , we obtain the following:

$$|S_1| \geq \Delta - (s + s_1), \text{ and} \quad (1)$$

$$|S_2| \geq \Delta - (s + s_2). \quad (2)$$

Also, consider the configuration  $C'$ . If  $|S_1 \cup S_2| \geq \Delta - s + 1$ , then, since we can pick  $s$  sets that are supersets of  $S_1 \cup S_2$ , we get that  $C'$  satisfies the requirements for being a configuration in  $\mathcal{N}_\Pi$ , reaching a contradiction. Hence, we obtain the following:

$$|S_1 \cup S_2| \leq \Delta - s. \quad (3)$$

By combining Equation (1), Equation (2), and Equation (3), we get that:

$$\begin{aligned} |S_1 \cap S_2| &= |S_1| + |S_2| - |S_1 \cup S_2| \\ &\geq (\Delta - s - s_1) + (\Delta - s - s_2) - (\Delta - s) \\ &= \Delta - s - s_1 - s_2. \end{aligned} \quad (4)$$

If  $S_1 \not\subseteq \mathcal{C}_\Delta^1$ , then all the sets of  $C$  that are supersets of  $S_1$  are part of  $\mathcal{C}_1, \dots, \mathcal{C}_{\Delta-1}$ , but this implies that also in  $C'$  there are still  $k$  sets that are supersets of  $S_1$ , and hence that  $C'$  satisfies the requirements for being a configuration in  $\mathcal{N}_\Pi$ , reaching a contradiction. Hence,  $S_1 \subseteq \mathcal{C}_\Delta^1$ . For the same reason,  $S_2 \subseteq \mathcal{C}_\Delta^2$ . We thus get that  $S_1 \cap S_2 \subseteq C'$ . If  $|S_1 \cap S_2| \geq \Delta - (s + s_1 + s_2 + 1) + 1$ , then, since from  $C'$  we can pick  $s + s_1 + s_2 + 1$  sets that are supersets of  $S_1 \cap S_2$ , we get that  $C'$  satisfies the requirements for being a configuration in  $\mathcal{N}_\Pi$ , reaching a contradiction. Hence, we obtain the following:

$$|S_1 \cap S_2| < \Delta - s - s_1 - s_2. \quad (5)$$

But Equation (4) is in contradiction with Equation (5). Hence,  $C'$  is a valid configuration.  $\square$

**Edge Constraint of  $\Pi''$ .** We now characterize the edge constraint  $\mathcal{E}_{\Pi''}$  of  $\Pi''$ . As shown in Section 3, the edge constraint  $\mathcal{E}_{\Pi''}$  can be constructed in the following way: take all the configurations allowed by the edge constraint  $\mathcal{E}_{\Pi'} = \mathcal{E}_\Pi$ , and replace each label  $\ell(C)$  with the disjunction of all the labels of  $\Pi''$  that, before the renaming, contain  $\ell(C)$ . Note that, after the renaming, the labels that contain  $\ell(C)$  are exactly all the labels  $\ell(C')$  such that  $C' \subseteq C$ . Hence, from each configuration  $\ell(C_1)\ell(C_2)$  in  $\mathcal{E}_\Pi$ , we obtain the configuration itself, plus all the configurations  $\ell(C'_1)\ell(C'_2)$  such that for all  $1 \leq j \leq 2$ ,  $C'_j \subseteq C_j$ . By Observation 4.2, these configurations are also in  $\mathcal{E}_\Pi$ . Hence  $\mathcal{E}_{\Pi''} = \mathcal{E}_\Pi$ .

## 5 The Problem Family

In this section, we describe the family of problems that we use to prove our lower bounds. We first start from the formal definition, and then we provide some intuition. Finally, we explain the relation between these problems and some other natural problems.

### 5.1 Problem Definition

Consider a vector  $z$  of  $\beta + 1$  non-negative integers  $z_0, \dots, z_\beta$  and let  $\text{len}(z) = \beta$ . We define  $|z| = z_0 + \dots + z_\beta = |\mathfrak{C}|$ , where  $\mathfrak{C} = \{C_{i,j} \mid 0 \leq i \leq \beta, 1 \leq j \leq z_i\}$  is a set of colors. Given a color  $C_{i,j} \in \mathfrak{C}$ , we define  $\text{level}(C_{i,j}) = i$ , and given a set of colors  $\mathcal{C} \in 2^\mathfrak{C} \setminus \{\emptyset\}$ , we define  $\text{level}(\mathcal{C}) = \max\{i \mid C_{i,j} \in \mathcal{C}\}$ , where  $\text{level}(\emptyset) = -1$ . Finally, let  $\text{prefix}(z)$  be the inclusive prefix sum of  $z$ , that is,  $\text{prefix}(z)_i = \sum_{j \leq i} z_j$ . Assume that  $|z| \leq \Delta$ , then problem  $\Pi_\Delta(z)$  is defined as follows.

**Labels.** The label set of problem  $\Pi_\Delta(z)$  is defined as  $\Sigma_\Delta(z) = \mathcal{P} \cup \mathcal{S} \cup \{\mathbf{X}\}$ , where:

- $\mathcal{P} = \{\mathbf{P}_i, \mathbf{U}_i \mid 1 \leq i \leq \beta\}$ ;
- $\mathcal{S} = \{\ell(\mathcal{C}) \mid \mathcal{C} \in 2^{\mathcal{C}} \setminus \{\emptyset\}\}$ .

**Node Constraint.** The node constraint  $\mathcal{N}_\Delta(z)$  of problem  $\Pi_\Delta(z)$  contains the following allowed configurations for the nodes:

- $\ell(\mathcal{C})^{\Delta-x} \mathbf{X}^x$ , for each  $\mathcal{C} \in 2^{\mathcal{C}} \setminus \{\emptyset\}$ , where  $x = |\mathcal{C}| - 1$ ;
- $\mathbf{P}_i \mathbf{U}_i^{\Delta-1}$ , for each  $1 \leq i \leq \beta$ .

**Edge Constraint.** The edge constraint  $\mathcal{E}_\Delta(z)$  of problem  $\Pi_\Delta(z)$  contains the following allowed configurations for the edges:

- $\ell(\mathcal{C}_1)\ell(\mathcal{C}_2)$ , for all  $\mathcal{C}_1, \mathcal{C}_2 \in 2^{\mathcal{C}} \setminus \{\emptyset\}$  such that  $\mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$ ;
- $\mathbf{U}_i \mathbf{U}_j$ , for each  $1 \leq i, j \leq \beta$ ;
- $\mathbf{U}_i \mathbf{P}_j$ , for each  $1 \leq i < j \leq \beta$ ;
- $\mathbf{U}_i \ell(\mathcal{C})$ , for each  $\mathcal{C} \in 2^{\mathcal{C}} \setminus \{\emptyset\}$  and  $1 \leq i \leq \beta$ ;
- $\mathbf{P}_i \ell(\mathcal{C})$ , for all  $\mathcal{C} \in 2^{\mathcal{C}} \setminus \{\emptyset\}$  such that  $\text{level}(\mathcal{C}) < i$ ;
- $\mathbf{X}\mathbf{L}$ , for each  $\mathbf{L} \in \Sigma_\Delta(z)$ .

**Edge Diagram.** In order to compute  $\mathcal{R}(\Pi_\Delta(z))$ , it will be helpful to know the relation between the labels of  $\Pi_\Delta(z)$  according to the edge constraint  $\mathcal{E}_\Delta(z)$ . The following relation derives directly from the definition of  $\mathcal{E}_\Delta(z)$ . [Figure 4](#) shows an example, for  $\Pi_\Delta([1, 2, 1])$ .

**Observation 5.1.** *The labels of  $\Pi_\Delta(z)$  satisfy all and only the following strength relations w.r.t.  $\mathcal{E}_\Delta(z)$ .*

- $\mathbf{U}_i < \mathbf{U}_j$ , if  $j < i$ .
- $\mathbf{P}_i < \mathbf{P}_j$ , if  $i < j$ .
- $\mathbf{P}_i < \mathbf{U}_j$ , for all  $i, j$ .
- $\mathbf{P}_i < \ell(\mathcal{C})$ , if, for all  $\mathcal{C} \in \mathcal{C}$ ,  $\text{level}(\mathcal{C}) \geq i$ .
- $\ell(\mathcal{C}) < \ell(\mathcal{C}')$ , if  $\mathcal{C}' \subseteq \mathcal{C}$ .
- $\ell(\mathcal{C}) < \mathbf{U}_i$ , if  $\text{level}(\mathcal{C}) \geq i$ .
- $\mathbf{L} < \mathbf{X}$ , for all  $\mathbf{L} \in \Sigma_\Delta(z)$ .

## 5.2 Intuition Behind These Problems

The problem definition contains two main components, a part related to coloring, and a part related to pointers. We start by describing how to interpret the coloring part. Then, we will discuss how pointers can interact between themselves, and with the colors.




Figure 4: The edge diagram of  $\Pi_{\Delta}([1, 2, 1])$ . Color 0 is in level 0, colors 1 and  $1'$  are in level 1, and color 2 is in level 2. For the sake of readability, we omit  $\ell(\cdot)$ .

**Coloring Part.** Let  $\mathfrak{C}$  be a color space and assume that it is split into  $\beta + 1$  levels. We can think of color  $C_{i,j}$  to be the  $j$ -th color of level  $i$ . Nodes can either be colored (nodes outputting a configuration containing  $\ell(\mathcal{C})$ , for some  $\mathcal{C}$ ), or uncolored (nodes outputting a configuration containing  $P_i$ , for some  $i$ ). Colored nodes may even have multiple colors assigned ( $\mathcal{C}$  is a non-empty *set* of colors). Colored nodes do not need to satisfy the requirements of a proper coloring, but it is enough to satisfy the requirements of an arbdefective coloring, that is, they can mark  $x$  edges as outgoing (using label X), and the nodes reachable through those edges are allowed to have the same color as them. If a node is colored with multiple colors, then, ignoring edges that are marked with X by at least one endpoint, the coloring constraints must be satisfied for *all* colors assigned to that node: if a node is red *and* blue, then no neighbor must be red, and no neighbor must be blue (in fact, the edge constraint requires that if two nodes output  $\ell(\mathcal{C}_1)\ell(\mathcal{C}_2)$ , then  $\mathcal{C}_1$  and  $\mathcal{C}_2$  must be disjoint). The arbdefect  $x$  allowed for a node depends on the number of colors assigned to that node: if a node has a single color assigned, then  $x = 0$ , while if a node has two colors assigned, then  $x = 1$ , and so on. In general, the node constraint allows  $x = |\mathcal{C}| - 1$ . The idea here is that if a node manages to get more than 1 color assigned, then it is rewarded by being allowed to have a larger arbdefect. Note that the node constraint requires to have *exactly*  $x$  ports marked X, but this is just a technicality: if a node wants to output X on *at most*  $x$  ports, then it can always change the output on some arbitrary ports to X and obtain *exactly*  $x$  ports marked X, since X is compatible with any other label.

**Pointers Part.** Uncolored nodes must output pointers, and there are  $\beta$  types of pointers ( $P_1, \dots, P_{\beta}$ ). We can think of pointer  $P_i$  to be of level  $i$ . The edge constraint allows pointers to point to either:

- nodes outputting a pointer of strictly lower level ( $P_j$  is compatible with  $U_i$  only if  $i < j$ ), or
- nodes outputting a color set  $\mathcal{C}$  such that all the colors of  $\mathcal{C}$  are of lower levels ( $P_i\ell(\mathcal{C})$  is allowed

only if  $\text{level}(\mathcal{C}) < i$ , that is, the maximum level of the colors in  $\mathcal{C}$  is strictly less than  $i$ ).

Except for these constraints, all the other ports of nodes outputting pointers are unconstrained:  $U_i$  can be neighbor with all other  $U$  labels, with all color sets, and with  $X$ .

Hence, pointers can be interpreted as follows. Pointers must form chains of length at most  $\beta$ , and these chains must reach colored nodes. Not all types of chains are allowed: if the chain terminates by pointing to a color of level  $i$ , then the chain must have length at most  $\beta - i$ .

### 5.3 Relation with Natural Problems.

We now provide some informal examples of what problems our family can capture. We make this more formal in [Section 8](#).

**$\Delta$ -coloring.** Consider the case where  $\beta = 0$  (that is, we do not have pointers at all), and  $\mathfrak{C} = \{1, \dots, \Delta\}$ . We can see that the obtained problem is a relaxation of the  $\Delta$ -coloring problem.

**Arbdefective Colorings.** When  $\beta = 0$ , some variants of arbdefective coloring are captured as well. For example, if the color space contains  $\Delta$  colors, then there exist  $\lfloor \Delta/2 \rfloor$  disjoint sets of colors of size 2. If nodes output sets of size 2, then they are allowed to have arbdefect 1. Hence, by letting each node of color  $0 \leq i < \lfloor \Delta/2 \rfloor$  output the configuration  $\ell(\{C_{0,2i+1}, C_{0,2i+2}\})^{\Delta-1} X$ , we get a problem that is a relaxation of the 1-arbdefective  $\lfloor \Delta/2 \rfloor$ -coloring problem.

**Ruling Sets.** Consider the problem  $\Pi$  in the family where the color space contains a single color, in level 0, and there are  $\beta$  possible pointers. In a solution for the  $(2, \beta)$ -ruling set problem, nodes in the set form an independent set, and nodes not in the set are at distance at most  $\beta$  from nodes in the set. Hence, nodes can spend  $\beta$  rounds to solve  $\Pi$ : nodes in the set output the single color of level 0, and node not in the set output pointer chains based on their distance from the nodes in the set.

**$\alpha$ -Arbdefective  $c$ -Colored  $\beta$ -Ruling Sets.** This problem is a variant of ruling sets, where the nodes in the set do not need to satisfy the independence requirement, but they are only required to be  $c$ -colored with arbdefect at most  $\alpha$ . Note that, if we set  $\alpha = 0$  and  $c = 1$ , we get the  $(2, \beta)$ -ruling set problem. Hence, we can consider a problem of the family where the color space contains  $c(1 + \alpha)$  colors, all in level 0, that allows us to have  $c$  sets of colors with arbdefect  $\alpha$ , and a solution for  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling sets can be converted in  $\beta$  rounds into a solution for the problem in the family.

**Other Problems.** In all the examples made until now, all colors of the color space were in level 0, but our family of problems is more general. For example, our family contains a relaxation of the following problem. Nodes can either be colored or uncolored. There are two groups of colors. One of the two groups is allowed to have some arbdefect, while the other group is allowed some different arbdefect. Uncolored nodes must either be at distance at most 3 from a node colored colored with a color from the first group, or at distance at most 2 from a node colored with a color from the second group.

## 6 Lower Bounds

In this section, we prove some properties about the problems in our family. In particular, we show that there exists a long sequence of problems, all contained in the family, such that the first problem

of the sequence is a relaxation of a problem for which we want to prove a lower bound, and then for each problem in the sequence it holds that:

- it is not 0-round solvable;
- it is at least 1 round harder than the next problem in the sequence;
- it can be described with a small amount of labels.

By combining the first two properties, we obtain that a sequence of  $T$  problems implies a lower bound of  $T$  rounds for the first problem in the port numbering model. As we will see later, in order to lift this lower bound to the LOCAL model, we also need the third property: the number of labels used to describe the problems in the family should stay small.

In order to construct this sequence, we start in [Section 6.1](#) by computing  $\Pi' = \mathcal{R}(\Pi_\Delta(z))$ , and prove some properties about it. While it is easy to check that, if  $|z| \leq \Delta$ , then  $\Pi_\Delta(z)$  contains at most  $2^\Delta + 2\text{len}(z) \leq 2^\Delta(1 + \text{len}(z))$  labels, we also make sure that the intermediate problem  $\Pi'$  can also be described with that amount of labels.

We continue, in [Section 6.2](#), where we show the core part of our proof, that is, we apply the universal quantifier on  $\mathcal{R}(\Pi_\Delta(z))$ .

Then, in [Section 6.3](#), we show that  $\overline{\mathcal{R}(\Pi_\Delta(z))}$  can be relaxed to  $\Pi_\Delta(\text{prefix}(z))$ , implying that if  $\Pi_\Delta(z)$  can be solved in  $T$  rounds, then  $\Pi_\Delta(\text{prefix}(z))$  can be solved in at most  $T - 1$  rounds.

By combining [Section 6.1](#), [Section 6.2](#), and [Section 6.3](#), we obtain the following lemma.

**Lemma 6.1.** *Let  $z' = \text{prefix}(z)$ , and assume that  $|z'| \leq \Delta$  if  $\text{len}(z') = 0$ , and  $|z'| \leq \Delta - 1$  otherwise. Let  $\Pi' = \mathcal{R}(\Pi_\Delta(z))$ . Then,  $\Pi_\Delta(z')$  is a relaxation of  $\overline{\mathcal{R}(\Pi')}$ . The number of labels of  $\Pi_\Delta(z)$ ,  $\Pi'$ , and  $\Pi_\Delta(z')$  is upper bounded by  $2^\Delta(1 + \text{len}(z))$ .*

We continue, in [Section 6.4](#), by showing that, unless  $|z|$  is too large,  $\Pi_\Delta(z)$  cannot be solved in 0 rounds in the port numbering model. In this way, we obtain all the requirements mentioned before.

We conclude, in [Section 6.5](#), by combining all the ingredients and give a lower bound on the length of the problem sequence.

## 6.1 Computing $\mathcal{R}(\Pi_\Delta(z))$

Recall the definition of  $\Pi_\Delta(z)$  (see [Section 5](#)), and define  $U_0 := \ell(\emptyset) := X$ . The following lemma computes  $\mathcal{R}(\Pi_\Delta(z))$ .

**Lemma 6.2.** *The label set of problem  $\mathcal{R}(\Pi_\Delta(z))$  is*

$$\Sigma_{\mathcal{R}} := \{\langle P_i, \ell(\mathcal{C}) \rangle \mid 1 \leq i \leq \beta, \mathcal{C} \in 2^{\mathfrak{C}}, \text{level}(\mathfrak{C} \setminus \mathcal{C}) \leq i - 1\} \cup \{\langle U_i, \ell(\mathcal{C}) \rangle \mid \mathcal{C} \in 2^{\mathfrak{C}}, \text{level}(\mathcal{C}) \leq i \leq \beta\}.$$

*Its edge constraint  $\mathcal{E}_{\mathcal{R}}$  consists of the following configurations:*

$$\begin{aligned} \langle P_i, \ell(\mathcal{C}) \rangle \langle U_{i-1}, \ell(\mathfrak{C} \setminus \mathcal{C}) \rangle & \quad \text{for all } 1 \leq i \leq \beta \text{ and } \mathcal{C} \in 2^{\mathfrak{C}} \text{ satisfying } \text{level}(\mathfrak{C} \setminus \mathcal{C}) \leq i - 1, \text{ and} \\ \langle U_\beta, \ell(\mathcal{C}) \rangle \langle U_\beta, \ell(\mathfrak{C} \setminus \mathcal{C}) \rangle & \quad \text{for all } \mathcal{C} \in 2^{\mathfrak{C}}. \end{aligned}$$

*Its node constraint  $\mathcal{N}_{\mathcal{R}}$  consists of all configurations  $B_1 \dots B_\Delta$  with labels  $B_i \in \Sigma_{\mathcal{R}}$  such that there exists some configuration  $A_1 \dots A_\Delta \in \mathcal{N}_\Delta(z)$  with  $A_i \in B_i$  for all  $1 \leq i \leq \Delta$ .*

*Proof.* Consider an arbitrary configuration  $B_1 B_2$  in the edge constraint  $\mathcal{E}_{\mathcal{R}}$  of  $\mathcal{R}(\Pi_\Delta(z))$ . We will show that  $B_1 B_2$  is of the form stated in the lemma. We consider two cases.

Consider first the case that at least one of  $B_1$  and  $B_2$  (which, by the definition of  $\mathcal{R}(\cdot)$  are sets of labels from  $\Sigma_\Delta(z)$ ) contains a label from  $\{P_1, \dots, P_\Delta\}$ . Let  $i$  be the smallest index for which  $P_i$  is contained in  $B_1$  or  $B_2$ , and assume w.l.o.g. that  $P_i \in B_1$ . Let  $\mathcal{C} \in 2^\mathfrak{C}$  be the set of all colors  $C \in \mathfrak{C}$  such that there exists a set  $\mathcal{C}'$  containing  $C$  for which  $\ell(\mathcal{C}') \in B_1$  (possibly  $\mathcal{C} = \emptyset$ ). In other words,  $\mathcal{C}$  contains all the colors that appear in  $B_1$ . Observe that  $B_2$  cannot contain any label  $\ell(\mathcal{C}'')$  such that  $\mathcal{C}'' \cap \mathcal{C} \neq \emptyset$  since otherwise there exist  $\mathcal{C}', \mathcal{C}'' \in 2^\mathfrak{C}$  such that  $\ell(\mathcal{C}') \in B_1$ ,  $\ell(\mathcal{C}'') \in B_2$ , and  $\mathcal{C}' \cap \mathcal{C}'' \neq \emptyset$ , which is a contradiction to the definition of  $\mathcal{R}(\Pi_\Delta(z))$  (as  $\mathcal{C}' \cap \mathcal{C}'' \neq \emptyset$  implies that  $\ell(\mathcal{C}')\ell(\mathcal{C}'') \notin \mathcal{E}_\Delta(z)$ ). Similarly, for any  $1 \leq j \leq \beta$ , set  $B_2$  does not contain  $P_j$  since  $B_1$  contains  $P_i$ , and  $P_i P_j \notin \mathcal{E}_\Delta(z)$ . Since, for any  $A \in \Sigma_\Delta(z) \setminus (\{\ell(\mathcal{C}'') \mid \mathcal{C}'' \cap \mathcal{C} \neq \emptyset\} \cup \{P_1, \dots, P_\Delta\})$ , we have  $\ell(\mathcal{C})A \in \mathcal{E}_\Delta(z)$ , it follows that  $\ell(\mathcal{C}) \in B_1$ , by maximality. Hence,  $B_1$  contains  $\ell(\mathcal{C})$  and  $P_i$ , but not any  $P_j$  with  $j < i$  (by the definition of  $i$ ) and not any  $\ell(\mathcal{C}'')$  with  $\mathcal{C}'' \not\subseteq \mathcal{C}$  (by the definition of  $\mathcal{C}$ ). By [Observation 5.1](#) and the fact that by [Observation 3.3](#)  $B_1$  is right-closed w.r.t.  $\mathcal{E}_\Delta(z)$ , it follows that the labels in  $B_1$  are exactly the labels from  $\Sigma_\Delta(z)$  that are successors of  $\ell(\mathcal{C})$  (w.r.t.  $\mathcal{E}_\Delta(z)$ ) or  $P_i$ , i.e., we have  $B_1 = \langle P_i, \ell(\mathcal{C}) \rangle$ . Note further that the right-closedness of  $B_1$ , together with [Observation 5.1](#) and the fact that  $P_i \in B_1$ , implies that  $B_1$  contains  $\ell(\{C\})$  for all  $C \in \mathfrak{C}$  satisfying  $\text{level}(C) \geq i$ . Hence,  $\text{level}(\mathfrak{C} \setminus \mathcal{C}) \leq i - 1$ , by the definition of  $\mathcal{C}$ .

By the definition of  $\mathcal{R}(\cdot)$  (in particular by maximality), the set  $B_2$  consists of precisely those labels  $A \in \Sigma_\Delta(z)$  that satisfy  $AA' \in \mathcal{E}_\mathcal{R}$  for all  $A' \in B_1$ . By going through the labels in  $\Sigma_\Delta(z)$  and checking whether this property holds (by going through the list of edge configurations of  $\Pi_\Delta(z)$ ), we see that  $B_2 = \{\ell(\mathcal{C}') \mid \mathcal{C}' \in 2^\mathfrak{C} \setminus \{\emptyset\}, \mathcal{C}' \subseteq \mathfrak{C} \setminus \mathcal{C}\} \cup \{U_1, \dots, U_{i-1}\} \cup \{X\}$ . Note that, for any nonempty  $\mathcal{C}' \subseteq \mathfrak{C} \setminus \mathcal{C}$  and any  $i \leq j \leq \beta$ , we have  $\ell(\mathcal{C}')P_j \in \mathcal{E}_\mathcal{R}$  since  $\text{level}(\mathfrak{C} \setminus \mathcal{C}) \leq i - 1$ , as observed above. Now, by [Observation 5.1](#), the fact that by [Observation 3.3](#)  $B_2$  is right-closed w.r.t.  $\mathcal{E}_\Delta(z)$ , and the fact that  $\text{level}(\mathfrak{C} \setminus \mathcal{C}) \leq i - 1$ , we obtain that  $B_2 = \langle U_{i-1}, \ell(\mathfrak{C} \setminus \mathcal{C}) \rangle$ . It follows that  $B_1 B_2 = \langle P_i, \ell(\mathcal{C}) \rangle \langle U_{i-1}, \ell(\mathfrak{C} \setminus \mathcal{C}) \rangle$ , where  $\text{level}(\mathfrak{C} \setminus \mathcal{C}) \leq i - 1$ , as desired.

Consider now the second case, namely that both  $B_1$  and  $B_2$  do not contain any label from  $\{P_1, \dots, P_\Delta\}$ . Again, let  $\mathcal{C} \in 2^\mathfrak{C}$  be the set of all colors  $C \in \mathfrak{C}$  such that there exists a set  $\mathcal{C}'$  containing  $C$  for which  $\ell(\mathcal{C}') \in B_1$ . Similarly to above, we obtain that  $B_2$  does not contain any label  $\ell(\mathcal{C}'')$  such that  $\mathcal{C}'' \cap \mathcal{C} \neq \emptyset$ , and that  $\ell(\mathcal{C}) \in B_1$ . Also, observe that, by maximality, both  $B_1$  and  $B_2$  contain  $U_\beta$  since  $U_\beta A \in \mathcal{E}_\mathcal{R}$  for any  $A \in \Sigma_\Delta(z) \setminus \{P_1, \dots, P_\Delta\}$ , and  $B_1$  and  $B_2$  do not contain any label from  $\{P_1, \dots, P_\Delta\}$ .

Since  $B_1$  contains  $\ell(\mathcal{C})$  and  $U_\beta$ , but not any  $P_j$  and not any  $\ell(\mathcal{C}'')$  with  $\mathcal{C}'' \not\subseteq \mathcal{C}$ , we obtain  $B_1 = \langle U_\beta, \ell(\mathcal{C}) \rangle$ , by [Observation 5.1](#) and the fact that by [Observation 3.3](#)  $B_1$  is right-closed w.r.t.  $\mathcal{E}_\Delta(z)$ . Moreover, with an analogous argumentation for computing  $B_2$  as in the first case, we obtain  $B_2 = \langle U_\beta, \ell(\mathfrak{C} \setminus \mathcal{C}) \rangle$ . Hence, also in this second case  $B_1 B_2$  is of a form given in the lemma statement.

It is straightforward to verify that the edge configurations given in the lemma statement are indeed contained in  $\mathcal{E}_\mathcal{R}$ , by checking that any two labels picked from the two sets in the considered configuration indeed form an edge configuration of  $\Pi_\Delta(z)$ , and that for both sets it holds that any label that is not in the set cannot be added to the set without violating this property. Hence, the edge constraint  $\mathcal{E}_\mathcal{R}$  is precisely as given in the lemma.

Moreover, by collecting all labels (i.e., sets of labels from  $\Pi_\Delta(z)$ ) that appear in  $\mathcal{E}_\mathcal{R}$ , we obtain precisely the label set  $\Sigma_\mathcal{R}$  given in the lemma. Finally, the statement about the node constraint  $\mathcal{N}_\mathcal{R}$  follows from the definition of  $\mathcal{R}(\cdot)$ .  $\square$

We now prove a property about the labels of  $\mathcal{R}(\Pi_\Delta(z))$ .

**Lemma 6.3.** *Assume  $|z| \leq \Delta$  if  $\text{len}(z) = 0$ , and  $|z| \leq \Delta - 1$  otherwise. For all  $A, B \in \Sigma_\mathcal{R}$  it holds that, if  $A \leq B$  according to  $\mathcal{N}_\mathcal{R}$ , then  $A \subseteq B$ .*

*Proof.* Assume that  $A \not\subseteq B$ , and let  $a \in \Sigma_\Delta(z)$  be a label contained in  $A \setminus B$ . We prove that  $A \not\leq B$ . For a contradiction, assume that  $A \leq B$ .

Consider some configuration  $a\ell_2 \dots \ell_\Delta \in \mathcal{N}_\Delta(z)$ . Note that, by [Lemma 6.2](#) and by [Observation 5.1](#), for each label  $\ell \in \Sigma_\Delta(z)$ , there is the label  $\langle \ell \rangle$  in  $\Sigma_{\mathcal{R}}$ . Hence, by [Lemma 6.2](#), the configuration  $A\langle \ell_2 \rangle \dots \langle \ell_\Delta \rangle$  is in  $\mathcal{N}_{\mathcal{R}}$ . Since  $A \leq B$ , then the configuration  $B\langle \ell_2 \rangle \dots \langle \ell_\Delta \rangle$  must also be in  $\mathcal{N}_{\mathcal{R}}$ .

By the definition of  $\mathcal{N}_{\mathcal{R}}$ , there must exist a configuration  $b\ell'_2 \dots \ell'_\Delta \in \mathcal{N}_\Delta(z)$  that is in  $B \times \langle \ell_2 \rangle \times \dots \times \langle \ell_\Delta \rangle$ . Since, by [Observation 3.3](#),  $B$  is right-closed, then the fact that  $a \notin B$  implies that also all labels that are at least as weak as  $a$  (according to  $\mathcal{E}_\Delta(z)$ ) are not in  $B$ . This implies that  $b$  is not as least as weak as  $a$  according to  $\mathcal{E}_\Delta(z)$ , since otherwise  $a$  would be a successor of  $b$ , and  $b$  being in  $B$  would imply  $a$  being in  $B$ . Also, since  $\ell'_i \in \langle \ell_i \rangle$  for all  $2 \leq i \leq \Delta$ , then we also have that  $\ell_i \leq \ell'_i$ .

Summarizing, there must be two configurations,  $a\ell_2 \dots \ell_\Delta$  and  $b\ell'_2 \dots \ell'_\Delta$ , that are both contained in  $\mathcal{N}_\Delta(z)$ , satisfying that, according to  $\mathcal{E}_\Delta(z)$ ,  $b \not\leq a$  (implying that either  $a < b$ , or that  $a$  and  $b$  are incomparable), and  $\ell_i \leq \ell'_i$  for all  $2 \leq i \leq \Delta$ . We show that all configurations of  $\mathcal{N}_\Delta(z)$  do not satisfy these requirements, reaching a contradiction.

In  $\mathcal{N}_\Delta(z)$  there are two types of configurations:  $\ell(C)^{\Delta-x}X^x$  (color configurations), and  $P_iU_i^{\Delta-1}$  (pointer configurations). We consider three cases separately.

- Both configurations are color configurations,  $C_1 = \ell(C_1)^{\Delta-x_1}X^{x_1}$  and  $C_2 = \ell(C_2)^{\Delta-x_2}X^{x_2}$ . Assume w.l.o.g. that  $x_1 \leq x_2$ . By assumption on  $z$ ,  $x_2 \leq \Delta - 1$ . Also by the definition of  $\mathcal{N}_\Delta(z)$ ,  $|C_1| \leq |C_2|$ . In this case,  $a$  cannot be  $X$  (since  $b$  must be a label strictly stronger than  $a$  or incomparable with  $a$ , but  $X$  is at least as strong as any other label). Hence,  $a$  must be either  $\ell(C_1)$  or  $\ell(C_2)$ .

If  $a = \ell(C_1)$  and  $b = \ell(C_2)$ , then by [Observation 5.1](#), since  $b \not\leq a$ , then  $C_1 \not\subseteq C_2$ , that, together with the fact that  $|C_1| \leq |C_2|$  implies that both  $C_1$  and  $C_2$  are not the set containing all colors. Hence,  $x_1 \leq x_2 \leq \Delta - 2$ , implying that there exists some  $i$  such that  $\ell'_i = \ell(C_2)$ . Then,  $\ell_i$  must be either  $X$  or  $\ell(C_1)$ , but both of them are not as weak as  $\ell(C_2)$  because of [Observation 5.1](#) and the fact that  $|C_1| \leq |C_2|$ .

If  $a = \ell(C_1)$  and  $b = X$ , then there must exist some index  $2 \leq i \leq \Delta$  such that  $\ell'_i = \ell(C_2)$ , but  $\ell(C_2)$  is not a successor of either  $\ell(C_1)$  or  $X$ , by [Observation 5.1](#) and the fact that  $|C_1| \leq |C_2|$ .

If  $a = \ell(C_2)$  and  $b = X$ , then since  $x_2 \geq x_1$ , then there must exist some index  $2 \leq i \leq \Delta$  such that  $\ell_i = X$  and  $\ell'_i = \ell(C_1)$ , that is now allowed.

If  $a = \ell(C_2)$  and  $b = \ell(C_1)$ , then, if  $a < b$ , then by [Observation 5.1](#), we obtain that  $x_1 < x_2$ . This implies that there exists some  $2 \leq i \leq \Delta$  such that  $\ell_i = X$ , and  $\ell'_i = \ell(C_1)$ , which is not allowed. If  $a$  and  $b$  are incomparable, then there must exist some  $2 \leq i \leq \Delta$  such that  $\ell_i = \ell(C_2)$  or  $\ell_i = X$ , and  $\ell'_i = \ell(C_1)$ . In both cases, we have  $\ell_i \not\leq \ell'_i$ .

- Both configurations are pointer configurations,  $C_1 = P_sU_s^{\Delta-1}$  and  $C_2 = P_tU_t^{\Delta-1}$ . Assume w.l.o.g. that  $s \leq t$ . Note that, according to [Observation 5.1](#),  $P_s \leq P_t < U_t \leq U_s$ .

If  $b = P_s$ , then  $a = P_t$  or  $a = U_t$ . We obtain that  $b \leq a$ , which is not allowed.

If  $b = U_s$ , then there exists some  $2 \leq i \leq \Delta$  such that  $\ell'_i = P_s$ . Then,  $\ell_i$  is either  $P_t$  or  $U_t$ . If  $\ell_i = U_t$ , or  $\ell_i = P_t$  and  $s \neq t$ , then  $\ell'_i < \ell_i$ , which is not allowed. If  $\ell_i = P_t$  and  $s = t$ , then  $a = U_s$ , implying  $a = b$ , which is not allowed.

If  $b = P_t$ , then  $a$  cannot be  $U_s$ , since  $P_t < U_s$ . Hence,  $a = P_s$ , implying  $s \neq t$ , since otherwise we would get that  $a = b$ . In this case,  $\ell'_2 = U_t$  and  $\ell_2 = U_s$ , and we obtain  $\ell'_2 < \ell_2$ , which is not allowed.

If  $b = U_t$ , then  $a$  cannot be  $U_s$ , since  $U_t \leq U_s$ . Hence,  $a = P_s$ . In this case, there must exist some  $2 \leq i \leq \Delta$  such that  $\ell'_i = P_t$ . Then,  $\ell_i$  is  $U_s$ , implying that  $\ell'_i < \ell_i$ , which is not allowed.

- One configuration is a color configuration  $C_1 = \ell(\mathcal{C})^{\Delta-x} \mathbf{X}^x$ , and the other is a pointer configuration  $C_2 = \mathbf{P}_s \mathbf{U}_s^{\Delta-1}$ . It is not possible that  $a = \mathbf{X}$ , since we need that  $b \not\leq a$ , but  $\mathbf{X}$  is at least as strong as any other label.

If  $a = \ell(\mathcal{C})$  and  $x \geq 1$ , then there must be some  $i$  such that  $\ell_i = \mathbf{X}$ , and since  $\ell_i \leq \ell'_i$ , then  $\ell'_i = \mathbf{X}$ , which is a contradiction.

If  $a = \ell(\mathcal{C})$  and  $x = 0$ , then for all  $2 \leq i \leq \Delta$ ,  $\ell_i = \ell(\mathcal{C})$ . In order to have  $\ell_i \leq \ell'_i$ , by **Observation 5.1** it must hold that  $\text{level}(\mathcal{C}) \geq s$ , since for some  $i$ , it holds that  $\ell'_i = \mathbf{U}_s$ . For the same reason, it must hold that for all  $2 \leq i \leq \Delta$ ,  $\ell'_i \neq \mathbf{P}_s$ . Hence,  $b = \mathbf{P}_s$ . Since  $x = 0$ , then  $|\mathcal{C}| = 1$ , implying that for all  $C \in \mathcal{C}$ ,  $\text{level}(C) \geq s$ , and hence that  $\mathbf{P}_s < \ell(\mathcal{C})$ , and hence that  $b < a$ , which is not allowed.

If  $a = \mathbf{P}_s$ , then, for all  $2 \leq i \leq \Delta$ ,  $\ell_i = \mathbf{U}_s$ . By **Observation 5.1**, for all  $\mathcal{C}$  it holds that  $\mathbf{U}_s \not\leq \ell(\mathcal{C})$ . Hence, for all  $2 \leq i \leq \Delta$ ,  $\ell'_i$  must be  $\mathbf{X}$ , but this is not possible, since the existence of the configuration  $C_2$  implies that  $\text{len}(z) > 0$ , and hence that  $|z| \leq \Delta - 1$ , and thus that  $x \leq \Delta - 2$ .

If  $a = \mathbf{U}_s$ , by **Observation 5.1** we either have that  $b = \mathbf{X}$ , or that  $b = \ell(\mathcal{C})$  and  $\text{level}(\mathcal{C}) < s$ . If  $b = \mathbf{X}$  then there must exist some  $i$  such that  $\ell'_i = \ell(\mathcal{C})$  such that  $\ell_i = \mathbf{U}_s$  (since, for the same reason as above, it must hold that  $x \leq \Delta - 2$ ), but this is not allowed. If  $b = \ell(\mathcal{C})$  and  $\text{level}(\mathcal{C}) < s$ , then there must be some  $i$  such that  $\ell'_i = \ell(\mathcal{C})$  and  $\ell_i$  is either  $\mathbf{P}_s$  or  $\mathbf{U}_s$ . Since  $\text{level}(\mathcal{C}) < s$ , then both are not as weak as  $\ell(\mathcal{C})$ .

□

Recall the definition of a relaxation of a node configuration in **Definition 3.2**. The second and much more complicated step we have to perform is to show that  $\overline{\mathcal{R}}(\mathcal{R}(\Pi_\Delta(z)))$  can be relaxed to  $\Pi_\Delta(z')$ , where  $z' = \text{prefix}(z)$ . Our key technical lemma is given in the following. Recall that we set  $\mathbf{U}_0 := \mathbf{X}$ .

**Lemma 6.4.** *Assume  $|z| \leq \Delta$  if  $\text{len}(z) = 0$ , and  $|z| \leq \Delta - 1$  otherwise. Then each node configuration of  $\overline{\mathcal{R}}(\mathcal{R}(\Pi_\Delta(z)))$  can be relaxed to*

$$\begin{aligned} \langle \langle \mathbf{P}_i \rangle \rangle \langle \langle \mathbf{U}_i \rangle \rangle^{\Delta-1} & \quad \text{for some } 1 \leq i \leq \beta, \text{ or} \\ \langle \langle \mathbf{U}_i, \ell(\mathcal{C}) \rangle \rangle, \langle \langle \mathbf{P}_j \rangle \rangle^{\Delta-|\mathcal{C}|-i+j+1} \langle \langle \mathbf{X} \rangle \rangle^{|\mathcal{C}|+i-j-1} & \quad \text{for some } \mathcal{C} \in 2^{\mathcal{E}} \setminus \{\emptyset\}, 0 \leq i, j \leq \beta \\ & \quad \text{satisfying } \text{level}(\mathcal{C}) \leq j \text{ and } i \in \{j, j+1\}, \end{aligned}$$

where we set  $\langle \langle \mathbf{U}_i, \ell(\mathcal{C}) \rangle \rangle, \langle \langle \mathbf{P}_j \rangle \rangle := \langle \langle \mathbf{U}_i, \ell(\mathcal{C}) \rangle \rangle$  if  $j = 0$ .

## 6.2 Proof of Lemma 6.4

Throughout this section, assume  $|z| \leq \Delta$  if  $\text{len}(z) = 0$ , and  $|z| \leq \Delta - 1$  otherwise. Let  $\mathcal{N}'$  denote the set of configurations given in **Lemma 6.4**, and  $\mathcal{N}_{\overline{\mathcal{R}}}$  the node constraint of  $\overline{\mathcal{R}}(\mathcal{R}(\Pi_\Delta(z)))$ . We prove **Lemma 6.4** by showing that any configuration with labels from  $2^{\Sigma_{\mathcal{R}}} \setminus \{\emptyset\}$  that cannot be relaxed to any configuration from  $\mathcal{N}'$  is not contained in  $\mathcal{N}_{\overline{\mathcal{R}}}$ .

Let  $\mathcal{S} = \mathbf{S}_1 \dots \mathbf{S}_\Delta$  be an arbitrary configuration with labels from  $2^{\Sigma_{\mathcal{R}}} \setminus \{\emptyset\}$  that cannot be relaxed to any configuration from  $\mathcal{N}'$ . If at least one of the  $\mathbf{S}_k$  is not right-closed (w.r.t.  $\mathcal{N}_{\mathcal{R}}$ ), then, by **Observation 3.3**,  $\mathcal{S}$  is not contained in  $\mathcal{N}_{\overline{\mathcal{R}}}$ , and we are done. Hence, assume throughout the remainder of the section that all  $\mathbf{S}_k$  are right-closed.

In the following, we collect some properties that  $\mathcal{S}$  satisfies due to the fact that it cannot be relaxed to any configuration from  $\mathcal{N}'$ . Observe that, by definition, the  $\mathbf{S}_k$  are sets of sets of labels from  $\Sigma_\Delta(\Pi)$ .



**Lemma 6.5.** *The configuration  $\mathcal{S}$  satisfies the following two properties.*

1. For any  $\mathcal{C} \in 2^{\mathcal{C}} \setminus \{\emptyset\}$ ,  $0 \leq i, j \leq \beta$  satisfying  $\text{level}(\mathcal{C}) \leq j$  and  $i \in \{j, j+1\}$ , there exist at least  $|\mathcal{C}| + i - j$  indices  $k \in \{1, \dots, \Delta\}$  such that there exists some set  $\mathbf{Q}_k \in \mathbf{S}_k$  satisfying
  - (a)  $U_i \notin \mathbf{Q}_k$  or  $\ell(\mathcal{C}) \notin \mathbf{Q}_k$ , **and**
  - (b)  $P_j \notin \mathbf{Q}_k$  if  $j > 0$ .
2. For any  $1 \leq i \leq \beta$ ,
  - (a) there exist sets  $\mathbf{Q}_1 \in \mathbf{S}_1, \dots, \mathbf{Q}_\Delta \in \mathbf{S}_\Delta$  such that  $P_i \notin \mathbf{Q}_k$  for all  $1 \leq k \leq \Delta$ , **or**
  - (b) there exists some index  $1 \leq k \leq \Delta$  and a set  $\mathbf{Q}_k \in \mathbf{S}_k$  such that  $U_i \notin \mathbf{Q}_k$ .

*Proof.* Suppose for a contradiction that the lemma does not hold. Then  $\mathcal{S}$  violates property 1 or property 2.

Consider first the case that  $\mathcal{S}$  violates property 1, and let  $\mathcal{C} \in 2^{\mathcal{C}} \setminus \{\emptyset\}$ ,  $0 \leq i, j \leq \beta$  be a choice for which property 1 is violated, i.e., we have  $\text{level}(\mathcal{C}) \leq j$  and  $i \in \{j, j+1\}$ , and there exist at most  $|\mathcal{C}| + i - j - 1$  indices  $k \in \{1, \dots, \Delta\}$  such that there is some set  $\mathbf{Q}_k \in \mathbf{S}_k$  satisfying properties 1a and 1b. W.l.o.g., let these indices be  $\Delta - |\mathcal{C}| - i + j + 2, \dots, \Delta$ . Then, for any  $1 \leq k \leq \Delta - |\mathcal{C}| - i + j + 1$ , and any set  $\mathbf{Q}_k \in \mathbf{S}_k$ , property 1a is violated or property 1b is violated, i.e., 1)  $U_i \in \mathbf{Q}_k$  and  $\ell(\mathcal{C}) \in \mathbf{Q}_k$ , or 2)  $P_j \in \mathbf{Q}_k$  and  $j > 0$ . If the former holds, then, since by **Observation 3.3**  $\mathbf{Q}_k$  is right-closed w.r.t.  $\mathcal{E}_\Delta(z)$ , we have  $\langle U_i, \ell(\mathcal{C}) \rangle \subseteq \mathbf{Q}_k$ , and we obtain, by **Observation 3.4**, that  $\mathbf{Q}_k$  is at least as strong as  $\langle U_i, \ell(\mathcal{C}) \rangle$  according to  $\mathcal{N}_{\mathcal{R}}$ , which implies  $\mathbf{Q}_k \in \langle \langle U_i, \ell(\mathcal{C}) \rangle \rangle$ . If the latter holds, then  $\langle P_j \rangle \subseteq \mathbf{Q}_k$  and  $j > 0$ , and we analogously obtain that  $\mathbf{Q}_k \in \langle \langle P_j \rangle \rangle$  (and  $j > 0$ ). Since, for any  $1 \leq j \leq \beta$ , we have  $\langle \langle U_i, \ell(\mathcal{C}) \rangle \rangle \subseteq \langle \langle U_i, \ell(\mathcal{C}) \rangle, \langle P_j \rangle \rangle$  and  $\langle \langle P_j \rangle \rangle \subseteq \langle \langle U_i, \ell(\mathcal{C}) \rangle, \langle P_j \rangle \rangle$  (by the definition of  $\langle \cdot \rangle$ ), it follows that  $\mathbf{Q}_k \in \langle \langle U_i, \ell(\mathcal{C}) \rangle, \langle P_j \rangle \rangle$  (where we set  $\langle \langle U_i, \ell(\mathcal{C}) \rangle, \langle P_j \rangle \rangle := \langle \langle U_i, \ell(\mathcal{C}) \rangle \rangle$  if  $j = 0$ ). Note that this property holds for any  $\mathbf{Q}_k \in \mathbf{S}_k$ . Hence, for any  $1 \leq k \leq \Delta - |\mathcal{C}| - i + j + 1$ , we have  $\mathbf{S}_k \subseteq \langle \langle U_i, \ell(\mathcal{C}) \rangle, \langle P_j \rangle \rangle$ , which implies that  $\mathcal{S}$  can be relaxed to the configuration  $\langle \langle U_i, \ell(\mathcal{C}) \rangle, \langle P_j \rangle \rangle^{\Delta - |\mathcal{C}| - i + j + 1} \langle \langle X \rangle \rangle^{|\mathcal{C}| + i - j - 1}$ , yielding a contradiction to the fact that  $\mathcal{S}$  cannot be relaxed to any configuration from  $\mathcal{N}'$ . Here we used that any set of labels from  $\Sigma_{\mathcal{R}}$  is a subset of  $\langle \langle X \rangle \rangle$ , which follows from **Observation 3.4**, the fact that by **Observation 3.3** the sets in  $\Sigma_{\mathcal{R}}$  are right-closed w.r.t.  $\mathcal{E}_\Delta(z)$ , and the fact that  $X$  is at least as strong as any label in  $\Sigma_\Delta(z)$  (which in turn follows from **Observation 5.1**).

Now consider the second case, namely that  $\mathcal{S}$  violates property 2, and let  $1 \leq i \leq \beta$  be a choice for which property 2 is violated, i.e., we have that 1) there exists an index  $1 \leq k' \leq \Delta$  such that any set  $\mathbf{Q}_{k'} \in \mathbf{S}_{k'}$  satisfies  $P_i \in \mathbf{Q}_{k'}$ , and 2) for any  $1 \leq k \leq \Delta$  and any  $\mathbf{Q}_k \in \mathbf{S}_k$ , we have  $U_i \in \mathbf{Q}_k$ . W.l.o.g., assume that  $k' = 1$ . By an analogous argumentation to the one from the first case, we obtain  $\mathbf{Q}_1 \in \langle \langle P_i \rangle \rangle$  for any  $\mathbf{Q}_1 \in \mathbf{S}_1$ , and  $\mathbf{Q}_k \in \langle \langle U_i \rangle \rangle$  for any  $k \in \{2, \dots, \Delta\}$  and any  $\mathbf{Q}_k \in \mathbf{S}_k$ . Hence,  $\mathbf{S}_1 \subseteq \langle \langle P_i \rangle \rangle$ , and  $\mathbf{S}_k \subseteq \langle \langle U_i \rangle \rangle$  for any  $2 \leq k \leq \Delta$ . It follows that  $\mathcal{S}$  can be relaxed to the configuration  $\langle \langle P_i \rangle \rangle \langle \langle U_i \rangle \rangle^{\Delta - 1}$ , again obtaining a contradiction to the fact that  $\mathcal{S}$  cannot be relaxed to any configuration from  $\mathcal{N}'$ .  $\square$

Recall that in order to prove Lemma 6.4, we want to show that certain configurations with labels from  $2^{\Sigma_{\mathcal{R}}} \setminus \{\emptyset\}$  are not contained in  $\mathcal{N}_{\overline{\mathcal{R}}}$ . To this end, the following lemma gives a useful sufficient condition for the property that a configuration is not contained in  $\mathcal{N}_{\overline{\mathcal{R}}}$ .

**Lemma 6.6.** *Let  $D_1 \dots D_\Delta$  be a configuration with labels from  $2^{\Sigma_{\mathcal{R}}} \setminus \{\emptyset\}$ . Assume that there exists some  $(B_1, \dots, B_\Delta) \in D_1 \times \dots \times D_\Delta$  satisfying that*

1. for each  $\mathcal{C} \in 2^{\mathcal{C}} \setminus \{\emptyset\}$ , there are at least  $|\mathcal{C}|$  indices  $k \in \{1, \dots, \Delta\}$  such that  $\ell(\mathcal{C}) \notin B_k$ , **and**
2. there exists some  $0 \leq i \leq \beta$  such that

- (a) if  $i \geq 1$ , then  $P_i \notin B_k$  for all  $1 \leq k \leq \Delta$ , **and**  
(b) if  $i \leq \beta - 1$ , then there exists some index  $1 \leq k \leq \Delta$  such that  $U_{i+1} \notin B_k$ .

Then  $D_1 \dots D_\Delta \notin \mathcal{N}_{\overline{\mathcal{R}}}$ .

*Proof.* Let  $\mathcal{B} = B_1 \dots B_\Delta$  be a configuration with labels from  $\Sigma_{\mathcal{R}}$  satisfying properties 1 and 2. By the definition of  $\overline{\mathcal{R}}(\cdot)$ , in order to prove the lemma, it suffices to show that  $\mathcal{B} \notin \mathcal{N}_{\overline{\mathcal{R}}}$ .

Assume for a contradiction that  $\mathcal{B} \in \mathcal{N}_{\overline{\mathcal{R}}}$ , and recall that the labels in  $\mathcal{B}$  are sets of labels from  $\Sigma_\Delta(z)$ . By the definition of  $\mathcal{R}(\cdot)$ , it follows that there exists some  $(A_1, \dots, A_\Delta) \in B_1 \times \dots \times B_\Delta$  such that  $\mathcal{A} := A_1 \dots A_\Delta \in \mathcal{N}_\Delta(z)$ . If  $\mathcal{A} = \ell(\mathcal{C})^{\Delta-|\mathcal{C}|+1} \chi^{|\mathcal{C}|-1}$  for some  $\mathcal{C} \in 2^{\mathfrak{C}} \setminus \{\emptyset\}$ , then  $\mathcal{B}$  violates property 1. Hence,  $\mathcal{A}$  is not of this form, which implies that  $\mathcal{A} = P_j U_j^{\Delta-1}$  for some  $1 \leq j \leq \beta$ . W.l.o.g., assume that  $A_1 = P_j$  and  $A_k = U_j$  for any  $2 \leq k \leq \Delta$ . By **Observation 3.3** and **Observation 5.1**, it follows that  $P_i \in B_1$  for all  $j \leq i \leq \beta$ , and  $U_i \in B_k$  for all  $1 \leq i \leq j$  and  $1 \leq k \leq \Delta$ . Thus, for each  $0 \leq i \leq j - 1$ , configuration  $\mathcal{B}$  violates property 2b, and, for each  $j \leq i \leq \beta$ , configuration  $\mathcal{B}$  violates property 2a. Hence,  $\mathcal{B}$  violates property 2, which yields the desired contradiction.  $\square$

For a set  $V'$  of nodes of a graph  $G$ , let  $N_G(V')$  denote the set of all nodes of  $G$  that are adjacent to at least one vertex from  $V'$ . In order to prove Lemma 6.4, we will make use of Hall's marriage theorem [**Hal35**, Theorem 1] that can be stated as follows.

**Theorem 6.7** (Hall). *Let  $G = (V \cup W, E)$  be a bipartite graph such that, for each subset  $V' \subseteq V$ , we have  $|N_G(V')| \geq |V'|$ . Then there exists a matching in  $G$  that matches all vertices from  $V$ .*

In the following, we will combine **Lemmas 6.5** and **6.6** and **Theorem 6.7** to prove **Lemma 6.4**. Recall that to prove **Lemma 6.4** it suffices to show that  $\mathcal{S} = S_1 \dots S_\Delta \notin \mathcal{N}_{\overline{\mathcal{R}}}$ . To this end, we will show that there exists some  $(B_1, \dots, B_\Delta) \in S_1 \times \dots \times S_\Delta$  satisfying properties 1 and 2 from **Lemma 6.6**, and then apply **Lemma 6.6**.

Let  $1 \leq i^* \leq \beta$  denote the largest index such that, for each  $S_k$  and each  $Q_k \in S_k$ , we have  $U_{i^*} \in Q_k$ . If no such index exists, i.e., if, for each  $1 \leq i \leq \beta$ , there exists some  $1 \leq k \leq \Delta$  and some  $Q_k \in S_k$  such that  $U_i \notin Q_k$ , then set  $i^* := 0$ .

Consider the bipartite graph  $G = (V \cup W, E)$  where  $V = \mathfrak{C}$ ,  $W = \{S_1, \dots, S_\Delta\}$ , and there is an edge between some color  $C \in V$  and some set  $S_k \in W$  if and only if there exists some  $Q_k \in S_k$  satisfying 1)  $\ell(\{C\}) \notin Q_k$  and 2)  $P_{i^*} \notin Q_k$  if  $i^* > 0$ .

Consider some arbitrary set  $\mathcal{C} \in 2^{\mathfrak{C}} \setminus \{\emptyset\}$ , and set  $i_{\max} := \max(i^*, \text{level}(\mathcal{C}))$ . By applying (property 1 of) **Lemma 6.5** with parameters  $\mathcal{C}$  and  $i = j = i_{\max}$ , we obtain that there exist at least  $|\mathcal{C}|$  indices  $k \in \{1, \dots, \Delta\}$  such that there exists some set  $Q_k \in S_k$  satisfying

1.  $U_{i_{\max}} \notin Q_k$  or  $\ell(\mathcal{C}) \notin Q_k$ , and
2.  $P_{i_{\max}} \notin Q_k$  if  $i_{\max} > 0$ .

Consider an arbitrary such  $Q_k$ . If  $i_{\max} = i^*$ , then, by the definition of  $i^*$ , we have  $U_{i_{\max}} \in Q_k$ , and hence  $\ell(\mathcal{C}) \notin Q_k$  (by property 1). If  $i_{\max} \neq i^*$ , then, by the definition of  $i_{\max}$ , we have  $i_{\max} = \text{level}(\mathcal{C})$ , which implies that  $\ell(\mathcal{C})$  is at least as weak as  $U_{i_{\max}}$  according to  $\mathcal{E}_\Delta(z)$ , by **Observation 5.1**. Since, by **Observation 3.3**,  $Q_k$  is right-closed (w.r.t.  $\mathcal{E}_\Delta(z)$ ), it follows that  $\ell(\mathcal{C}) \in Q_k$  would imply  $U_{i_{\max}} \in Q_k$ . Thus, by property 1, we again obtain  $\ell(\mathcal{C}) \notin Q_k$ . Hence, irrespective of whether  $i_{\max} = i^*$  or  $i_{\max} \neq i^*$ , we have  $\ell(\mathcal{C}) \notin Q_k$ .

Moreover, observe that the characterization of  $\Sigma_{\mathcal{R}}$  given in **Lemma 6.2** implies, by **Observation 5.1**, that for every  $B \in \Sigma_{\mathcal{R}}$  there is a set  $\mathcal{C}' \in 2^{\mathfrak{C}}$  such that, for every  $\mathcal{C}'' \in 2^{\mathfrak{C}} \setminus \{\emptyset\}$ , we have  $\ell(\mathcal{C}'') \in B$  if and only if  $\mathcal{C}'' \subseteq \mathcal{C}'$ : if  $B$  is of the form  $\langle P_{i'}, \ell(\mathcal{C}''') \rangle$ , then  $\mathcal{C}' = \mathcal{C}'''$ , and if  $B$  is of the form  $\langle U_{i'}, \ell(\mathcal{C}''') \rangle$ , then also  $\mathcal{C}' = \mathcal{C}'''$ . We claim that there exists some  $C \in \mathcal{C}$  such that  $\ell(\{C\}) \notin Q_k$ . Assume for a

contradiction that, for all  $C \in \mathcal{C}$ , we have  $\ell(\{C\}) \in \mathbf{Q}_k$ . Then, since  $\mathbf{Q}_k \in \Sigma_{\mathcal{R}}$ , the aforementioned characterization implies that there exists some  $\mathcal{C}' \supseteq \bigcup_{C \in \mathcal{C}} \{C\}$  such that  $\ell(\mathcal{C}') \in \mathbf{Q}_k$  for any  $\mathcal{C}'' \subseteq \mathcal{C}'$ . Since  $\mathcal{C} \subseteq \mathcal{C}'$ , we obtain  $\ell(\mathcal{C}) \in \mathbf{Q}_k$ , yielding the desired contradiction. Hence, there exists some  $C \in \mathcal{C}$  such that  $\ell(\{C\}) \notin \mathbf{Q}_k$ .

Furthermore, if  $i^* > 0$ , then, by [Observation 5.1](#),  $P_{i^*} \leq P_{i_{\max}}$ , which implies  $P_{i^*} \notin \mathbf{Q}_k$ , by [property 2](#) and the right-closedness of  $\mathbf{Q}_k$  (which is due to [Observation 3.3](#)).

Combining the above insights, we obtain that there exist at least  $|\mathcal{C}|$  indices  $k \in \{1, \dots, \Delta\}$  such that there exists some set  $\mathbf{Q}_k \in \mathbf{S}_k$  satisfying that

1. there exists some color  $C \in \mathcal{C}$  with  $\ell(\{C\}) \notin \mathbf{Q}_k$ , and
2.  $P_{i^*} \notin \mathbf{Q}_k$  if  $i^* > 0$ .

As this statement holds for any  $\mathcal{C} \in 2^{\mathcal{C}} \setminus \{\emptyset\}$ , we obtain, by the definition of  $G$ , that, for any subset  $\mathcal{C} \subseteq V$ , we have  $N_G(\mathcal{C}) \geq |\mathcal{C}|$ . By applying [Theorem 6.7](#), i.e., Hall's marriage theorem, we obtain that there exists some matching  $M$  such that each  $C \in \mathcal{C} = V$  is matched. W.l.o.g., assume that the  $\mathbf{S}_k$  that are matched are  $\mathbf{S}_1, \dots, \mathbf{S}_{|\mathcal{C}|}$ . For each  $C \in \mathcal{C}$ , denote the  $\mathbf{S}_k$  color  $C$  is matched to by  $M(C)$ . Similarly, for each  $\mathbf{S}_k$  with  $1 \leq k \leq |\mathcal{C}|$ , denote the color  $C$  set  $\mathbf{S}_k$  is matched to by  $M(\mathbf{S}_k)$ .

Now, let  $\mathcal{C}_{\leq i^*} \in 2^{\mathcal{C}}$  denote the set of colors of level at most  $i^*$ . W.l.o.g., assume that the colors in  $\mathcal{C}_{\leq i^*}$  are matched with the sets from  $\{\mathbf{S}_1, \dots, \mathbf{S}_{|\mathcal{C}_{\leq i^*}|}\}$ . Furthermore, observe that, if  $i^* > 0$ , then [property 2](#) of [Lemma 6.5](#), together with the definition of  $i^*$ , implies that for any  $1 \leq k \leq \Delta$  there exists some set  $\mathbf{Q}_k \in \mathbf{S}_k$  satisfying  $P_{i^*} \notin \mathbf{Q}_k$ . Also, note that it follows from the definition of  $i^*$  that  $i^* = \beta$  or there exist some  $1 \leq k \leq \Delta$  and some  $\mathbf{Q}_k \in \mathbf{S}_k$  such that  $U_{i^*+1} \notin \mathbf{Q}_k$ . If  $i^* < \beta$ , denote the largest index  $k$  such that there exists some  $\mathbf{Q}_k \in \mathbf{S}_k$  satisfying  $U_{i^*+1} \notin \mathbf{Q}_k$  by  $k_{\max}$ . We consider two cases.

**Case I:  $i^* = \beta$ , or  $i^* < \beta$  and  $k_{\max} > |\mathcal{C}_{\leq i^*}|$**

For each  $1 \leq k \leq |\mathcal{C}|$ , set  $\mathbf{B}_k$  to be some  $\mathbf{Q}_k \in \mathbf{S}_k$  such that  $\ell(\{M(\mathbf{S}_k)\}) \notin \mathbf{Q}_k$  and  $P_{i^*} \notin \mathbf{Q}_k$  if  $i^* > 0$ . Such a choice exists by the definition of  $G$ . If  $i^* < \beta$ , set  $\mathbf{B}_{k_{\max}}$  to be some  $\mathbf{Q}_{k_{\max}} \in \mathbf{S}_{k_{\max}}$  such that  $U_{i^*+1} \notin \mathbf{Q}_{k_{\max}}$  (overwriting the previous choice for  $\mathbf{B}_{k_{\max}}$  if  $k_{\max} \leq |\mathcal{C}|$ ). If  $i^* > 0$ , then, for each  $1 \leq k \leq \Delta$  for which  $\mathbf{B}_k$  has not been set yet, set  $\mathbf{B}_k$  to be some  $\mathbf{Q}_k \in \mathbf{S}_k$  satisfying  $P_{i^*} \notin \mathbf{Q}_k$ . Such a choice exists as shown above. If  $i^* = 0$ , then, for each  $1 \leq k \leq \Delta$  for which  $\mathbf{B}_k$  has not been set yet, set  $\mathbf{B}_k$  to be some arbitrary  $\mathbf{Q}_k \in \mathbf{S}_k$ .

We claim that the obtained tuple  $(\mathbf{B}_1, \dots, \mathbf{B}_{\Delta}) \in \mathbf{S}_1 \times \dots \times \mathbf{S}_{\Delta}$  satisfies [properties 1 and 2](#) of [Lemma 6.6](#). Consider first [property 1](#), and consider some arbitrary  $\mathcal{C} \in 2^{\mathcal{C}} \setminus \{\emptyset\}$ . Let  $K$  denote the set of indices  $1 \leq k \leq |\mathcal{C}|$  such that  $M(\mathbf{S}_k) \in \mathcal{C}$ . As  $M$  is a matching, we have  $|K| = |\mathcal{C}|$ .

Consider some arbitrary  $k \in K$ . Since  $1 \leq k \leq |\mathcal{C}|$ , we know, by the definition of  $\mathbf{B}_k$ , that 1)  $\ell(\{M(\mathbf{S}_k)\}) \notin \mathbf{B}_k$ , or 2)  $i^* < \beta$ ,  $k = k_{\max} \leq |\mathcal{C}|$ , and  $U_{i^*+1} \notin \mathbf{B}_k$ . In the latter case, recall that  $k_{\max} > |\mathcal{C}_{\leq i^*}|$ , which implies that  $\text{level}(M(\mathbf{S}_k)) \geq i^* + 1$ . By [Observation 5.1](#), it follows that  $\ell(M(\mathbf{S}_k)) \leq U_{i^*+1}$ , and, by the right-closedness of  $\mathbf{B}_k$  (due to [Observation 3.3](#)), we obtain  $\ell(\{M(\mathbf{S}_k)\}) \notin \mathbf{B}_k$ . Hence, in either case, we have  $\ell(\{M(\mathbf{S}_k)\}) \notin \mathbf{B}_k$ . Now, by [Observation 5.1](#) and the fact that  $M(\mathbf{S}_k) \in \mathcal{C}$ , we obtain  $\ell(\mathcal{C}) \leq \ell(\{M(\mathbf{S}_k)\})$ , which, by the right-closedness of  $\mathbf{B}_k$ , implies  $\ell(\mathcal{C}) \notin \mathbf{B}_k$ . Hence, for each  $k \in K$ , we have  $\ell(\mathcal{C}) \notin \mathbf{B}_k$ . Since  $|K| = |\mathcal{C}|$ , it follows that there are at least  $|\mathcal{C}|$  indices  $k \in \{1, \dots, \Delta\}$  such that  $\ell(\mathcal{C}) \notin \mathbf{B}_k$ , which shows that  $(\mathbf{B}_1, \dots, \mathbf{B}_{\Delta})$  satisfies [property 1](#) of [Lemma 6.6](#).

Now consider [property 2](#) of [Lemma 6.6](#). We show that this property holds for  $(\mathbf{B}_1, \dots, \mathbf{B}_{\Delta})$  by showing that [properties 2a and 2b](#) of [Lemma 6.6](#) are satisfied when setting  $i = i^*$ . Consider first [property 2a](#). If  $i^* = 0$ , then [property 2a](#) is trivially satisfied, hence assume  $i^* > 0$ . By the definition of the  $\mathbf{B}_k$ , we obtain, for each  $1 \leq k \leq \Delta$  that  $P_{i^*} \notin \mathbf{B}_k$  or that  $U_{i^*+1} \notin \mathbf{B}_k$ , which also implies  $P_{i^*} \notin \mathbf{B}_k$ ,

due to the right-closedness of  $B_k$  and the fact that  $P_{i^*} \leq U_{i^*+1}$  (which follows from [Observation 5.1](#)). Therefore, property [2a](#) is satisfied.

Now consider property [2b](#) of [Lemma 6.6](#). If  $i^* = \beta$ , then property [2b](#) is trivially satisfied, hence assume  $i^* < \beta$ . By the definition of the  $B_k$ , we obtain  $U_{i^*+1} \notin B_k$  for  $k = k_{\max}$ , proving that also property [2b](#) is satisfied. Hence,  $(B_1, \dots, B_\Delta)$  satisfies also property [2](#) of [Lemma 6.6](#).

We conclude that  $(B_1, \dots, B_\Delta)$  satisfies the conditions given in [Lemma 6.6](#), and by applying [Lemma 6.6](#), we obtain  $S_1 \dots S_\Delta \notin \mathcal{N}_{\overline{\mathcal{R}}}$ , as desired.

**Case II:  $i^* < \beta$  and  $k_{\max} \leq |\mathcal{C}_{\leq i^*}|$**

The challenge that we have to overcome in Case II is that we cannot simply choose  $B_{k_{\max}}$  to be some  $Q_{k_{\max}} \in S_{k_{\max}}$  such that  $U_{i^*+1} \notin Q_{k_{\max}}$ . Such a choice would make sure that property [2b](#) of [Lemma 6.6](#) is satisfied, but it would not guarantee that the color  $C$  that is matched to  $S_{k_{\max}}$  (or more precisely, the set  $\{C\}$ ) satisfies the condition given in property [1](#) of [Lemma 6.6](#). Our solution is to change the matching  $M$  to another matching with the property that there is some  $S_k$  that is not matched to any color from  $\mathcal{C}_{\leq i^*}$  but contains a set  $Q_k$  satisfying  $U_{i^*+1} \notin Q_k$  (which allows us to choose  $B_k = Q_k$  and thereby satisfy property [2b](#) without violating property [1](#)). To this end, we essentially construct a suitable augmenting path along which we change the matching (while we will not use this exact terminology in the proof).

Let  $\mathcal{E} \subseteq \mathcal{C}$  denote the set of colors  $C$  such that there exists some set  $Q \in M(C)$  satisfying  $U_{i^*+1} \notin Q$ . By the definition of  $k_{\max}$  and the fact that  $k_{\max} \leq |\mathcal{C}_{\leq i^*}|$ , we obtain that  $\emptyset \neq \mathcal{E} \subseteq \mathcal{C}_{\leq i^*}$ . Moreover, recall that  $E$  denotes the set of edges of  $G$ , and recursively define sets  $\mathcal{F}_1, \mathcal{F}_2, \dots$  (which are subsets of  $\mathcal{C}_{\leq i^*}$ ) as follows.

$$\begin{aligned} \mathcal{F}_1 &:= \{C \in \mathcal{C}_{\leq i^*} \mid \text{there exists some } |\mathcal{C}_{\leq i^*}| + 1 \leq k \leq \Delta \text{ such that } \{C, S_k\} \in E\}, \\ \mathcal{F}_y &:= \{C \in \mathcal{C}_{\leq i^*} \setminus (\mathcal{F}_1 \cup \dots \cup \mathcal{F}_{y-1}) \mid \text{there exists some } C' \in \mathcal{F}_{y-1} \text{ such that } \{C, M(C')\} \in E\}, \\ &\text{for all } y \geq 2. \end{aligned}$$

The following lemma shows that in this recursive definition, the  $\mathcal{F}_y$  do not become empty before some  $\mathcal{F}_y$  is reached that has a nonempty intersection with  $\mathcal{E}$ .

**Lemma 6.8.** *Let  $y$  be some positive integer. If  $\mathcal{E} \cap (\mathcal{F}_1 \cup \dots \cup \mathcal{F}_{y-1}) = \emptyset$ , then  $\mathcal{F}_y \neq \emptyset$ .*

*Proof.* Set  $\mathcal{C} := \mathcal{C}_{\leq i^*} \setminus (\mathcal{F}_1 \cup \dots \cup \mathcal{F}_{y-1})$  and observe that  $\text{level}(\mathcal{C}) \leq \text{level}(\mathcal{C}_{\leq i^*}) \leq i^*$  since  $\mathcal{C} \subseteq \mathcal{C}_{\leq i^*}$ . Observe further that  $\mathcal{E} \cap (\mathcal{F}_1 \cup \dots \cup \mathcal{F}_{y-1}) = \emptyset$  implies  $\mathcal{C} \neq \emptyset$  since  $\emptyset \neq \mathcal{E} \subseteq \mathcal{C}_{\leq i^*}$  as observed above. Hence, by property [1](#) of [Lemma 6.5](#) (applied with parameters  $\mathcal{C}$ ,  $j = i^*$ , and  $i = i^* + 1$ ), we know that there exist at least  $|\mathcal{C}| + 1$  indices  $k \in \{1, \dots, \Delta\}$  such that there exists some set  $Q_k \in S_k$  satisfying 1)  $U_{i^*+1} \notin Q_k$  or  $\ell(\mathcal{C}) \notin Q_k$ , and 2)  $P_{i^*} \notin Q_k$  if  $i^* > 0$ . By the pigeonhole principle, it follows that there also exists such a  $k$  additionally satisfying that  $S_k$  is not matched or  $M(S_k) \notin \mathcal{C}$ ; consider such a  $k$ . Since  $\mathcal{E} \subseteq \mathcal{C}$ , we obtain that  $S_k$  is not matched or  $M(S_k) \notin \mathcal{E}$ . Observe that each  $S_{k'}$  such that there exists some  $Q_{k'} \in S_{k'}$  with  $U_{i^*+1} \notin Q_{k'}$  is matched (as, by the definition of  $k_{\max}$ , we have  $k' \leq k_{\max} \leq |\mathcal{C}_{\leq i^*}|$ , and, for each  $k'' \leq |\mathcal{C}_{\leq i^*}|$ , all  $S_{k''}$  are matched (with colors from  $\mathcal{C}_{\leq i^*}$ )) and satisfies  $M(S_{k'}) \in \mathcal{E}$  (by the definition of  $\mathcal{E}$ ). Hence, for all  $Q_k \in S_k$ , we have  $U_{i^*+1} \in Q_k$ .

By the definition of  $k$ , it follows that there exists some set  $Q_k \in S_k$  satisfying 1)  $\ell(\mathcal{C}) \notin Q_k$ , and 2)  $P_{i^*} \notin Q_k$  if  $i^* > 0$ . Recall that, as observed before,  $\ell(\mathcal{C}) \notin Q_k$  implies that there exists some  $C \in \mathcal{C}$  such that  $\ell(\{C\}) \notin Q_k$  (due to the characterization of  $\Sigma_{\mathcal{R}}$  given in [Lemma 6.2](#)). Consider such a  $C$ . By the definition of  $G$ , there is an edge between  $C$  and  $S_k$ .

Recall that  $C \in \mathcal{C}$  and that  $S_k$  is not matched or  $M(S_k) \notin \mathcal{C}$ . If  $y = 1$ , then  $\mathcal{C} = \mathcal{C}_{\leq i^*}$ , and we obtain  $|\mathcal{C}_{\leq i^*}| + 1 \leq k \leq \Delta$ , which implies  $C \in \mathcal{F}_1$ . Thus,  $\mathcal{F}_y \neq \emptyset$ , as desired. Hence, assume that  $y \geq 2$ . Observe that we have  $1 \leq k \leq |\mathcal{C}_{\leq i^*}|$  as otherwise  $C \in \mathcal{F}_1$ , which would contradict

$C \in \mathcal{C} \subseteq \mathcal{C}_{\leq i^*} \setminus \mathcal{F}_1$ . In particular,  $S_k$  is matched, and  $M(S_k) \in \mathcal{C}_{\leq i^*}$ . Since  $M(S_k) \notin \mathcal{C}$ , we have  $M(S_k) \in \mathcal{C}_{\leq i^*} \setminus \mathcal{C} = \mathcal{F}_1 \cup \dots \cup \mathcal{F}_{y-1}$ . As, by definition,  $\mathcal{F}_1, \dots, \mathcal{F}_{y-1}$  are pairwise disjoint, it follows that there is precisely one  $1 \leq y' \leq y-1$  such that  $M(S_k) \in \mathcal{F}_{y'}$ . Again by the definition of  $\mathcal{F}_1, \mathcal{F}_2, \dots$ , we obtain that  $C \in \mathcal{F}_{y'+1}$ . Since  $C \in \mathcal{C} = \mathcal{C}_{\leq i^*} \setminus (\mathcal{F}_1 \cup \dots \cup \mathcal{F}_{y-1})$ , we see that  $y' = y-1$ . Hence,  $C \in \mathcal{F}_y$ , and we obtain  $\mathcal{F}_y \neq \emptyset$ , as desired.  $\square$

Since the  $\mathcal{F}_y$  are subsets of (the finite set)  $\mathcal{C}_{\leq i^*}$  and pairwise disjoint, there is some positive integer  $y'$  with  $\mathcal{F}_{y'} = \emptyset$ . By [Lemma 6.8](#), it follows that  $\mathcal{E} \cap (\mathcal{F}_1 \cup \dots \cup \mathcal{F}_{y'-1}) \neq \emptyset$  (which also implies that  $y' \geq 2$ ). Hence, there is some positive integer  $y^*$  with  $\mathcal{E} \cap \mathcal{F}_{y^*} \neq \emptyset$ . Let  $C'$  be some color contained in  $\mathcal{E} \cap \mathcal{F}_{y^*}$ .

Let  $C^{(1)}, \dots, C^{(y^*)}$  be colors from  $\mathcal{F}_1, \dots, \mathcal{F}_{y^*}$ , respectively, such that  $C^{(y^*)} = C'$  and, for any  $1 \leq y \leq y^* - 1$ , we have  $\{C^{(y+1)}, M(C^{(y)})\} \in E$ . Such  $C^{(y)}$  exist due to the definition of the  $\mathcal{F}_y$ . Moreover, let  $|\mathcal{C}_{\leq i^*}| + 1 \leq k' \leq \Delta$  be some index such that  $\{C^{(1)}, S_{k'}\} \in E$ .

If  $S_{k'}$  is matched in  $M$ , then  $M(S_{k'}) \in \mathcal{C} \setminus \mathcal{C}_{\leq i^*}$ , which implies that  $\text{level}(M(S_{k'})) \geq i^* + 1$ . Hence, in this case,  $\ell(\{M(S_{k'})\}) \leq U_{i^*+1}$ , by [Observation 5.1](#), and the fact that  $M(C^{(y^*)})$  contains some set  $Q$  with  $U_{i^*+1} \notin Q$  (which follows from  $C^{(y^*)} \in \mathcal{E}$ ) implies that there is some  $Q \in M(C^{(y^*)})$  such that  $\ell(\{M(S_{k'})\}) \notin Q$ , by the right-closedness of  $Q$ . If  $i^* > 0$ , then for any such  $Q$ , the fact that  $P_{i^*} \leq U_{i^*+1}$  analogously implies  $P_{i^*} \notin Q$ , and it follows that in  $G$  there is an edge between  $M(S_{k'})$  and  $M(C^{(y^*)})$ .

Let  $M'$  be the matching obtained from  $M$  by

1. removing the edges in  $M$  that are incident to some vertex from  $\{C^{(1)}, \dots, C^{(y^*)}\}$ ,
2. if  $S_{k'}$  is matched in  $M$ , removing the corresponding matching edge,
3. if  $S_{k'}$  was matched in  $M$ , adding edge  $\{M(S_{k'}), M(C^{(y^*)})\}$
4. adding edge  $\{C^{(1)}, S_{k'}\}$ , and
5. adding edge  $\{C^{(y+1)}, M(C^{(y)})\}$ , for each  $1 \leq y \leq y^* - 1$ .

Note that all edges added to  $M'$  are indeed edges in  $G$ , by the above discussion. Note further that for each added edge  $e$ , the two endpoints of  $e$  were unmatched after step 2; hence,  $M'$  is indeed a matching. Finally, observe that the construction of  $M'$  ensures that each color  $C \in \mathcal{C}$  is matched in  $M'$ .

For each vertex  $v$  of  $G$  that is matched in  $M'$ , denote the vertex that  $v$  is matched to by  $M'(v)$ . Moreover, permute the indexing of the  $S_1, \dots, S_\Delta$  such that the colors in  $\mathcal{C}$  are matched (in  $M'$ ) to the sets in  $\{S_1, \dots, S_{|\mathcal{C}|}\}$ , and the colors in  $\mathcal{C}_{\leq i^*}$  are matched to the sets in  $\{S_1, \dots, S_{|\mathcal{C}_{\leq i^*}|}\}$ . Whenever we refer to some  $S_k$  in the remainder of the section, it is w.r.t. this new indexing (unless indicated otherwise). Let  $k^*$  denote the new index of set  $M(C^{(y^*)})$ . By the construction of  $M'$  and the above discussion, we know that  $S_{k^*}$  is unmatched in  $M'$  or matched to some color from  $\mathcal{C} \setminus \mathcal{C}_{\leq i^*}$ , implying that  $k^* \geq |\mathcal{C}_{\leq i^*}| + 1$ . Note that, due to  $S_{k^*} = M(C^{(y^*)})$ ,  $C^{(y^*)} \in \mathcal{E}$ , and by the definition of  $\mathcal{E}$ , there is some set  $Q_{k^*} \in S_{k^*}$  such that  $U_{i^*+1} \notin Q_{k^*}$ . Hence, according to the new indexing,  $k_{\max} \geq k^* \geq |\mathcal{C}_{\leq i^*}| + 1$ .

Hence, we have essentially transformed the setting of Case II to the setting of Case I: by simply setting the  $B_k$  as in Case I (where we use matching  $M'$  (instead of matching  $M$ ), the new indexing, and the new definition of  $k_{\max}$ ) and using the same proof as in Case I, we obtain  $S_1 \dots S_\Delta \notin \mathcal{N}_{\overline{\mathcal{R}}}$ , as before. (The only difference is that we do not have to deal with the case  $i^* = \beta$  (as we cannot have  $i^* = \beta$  in Case II), which only works in our favor.) We conclude that in either case we have  $S_1 \dots S_\Delta \notin \mathcal{N}_{\overline{\mathcal{R}}}$ , which shows that any configuration with labels from  $2^{\Sigma_{\mathcal{R}}} \setminus \{\emptyset\}$  that cannot be relaxed to any configuration from  $\mathcal{N}'$  is not contained in  $\mathcal{N}_{\overline{\mathcal{R}}}$ , and proves [Lemma 6.4](#).

### 6.3 Relaxing $\overline{\mathcal{R}}(\mathcal{R}(\Pi_\Delta(z)))$

Recall that we want to show that, if  $\Pi_\Delta(z)$  can be solved in some number  $T$  of rounds, then  $\Pi_\Delta(z')$ , where  $z' = \text{prefix}(z)$ , can be solved in  $T - 1$  rounds. To this end, we show that  $\overline{\mathcal{R}}(\mathcal{R}(\Pi_\Delta(z)))$  can be relaxed to  $\Pi_\Delta(z')$ , using [Lemma 6.4](#), and then make use of the fact that  $\overline{\mathcal{R}}(\mathcal{R}(\Pi_\Delta(z)))$  can be solved one round faster than  $\Pi_\Delta(z)$ . Throughout this section, we assume that  $|z| \leq \Delta$  if  $\text{len}(z) = 0$ , and  $|z| \leq \Delta - 1$  otherwise (which is a requirement for applying [Lemma 6.4](#)).

Recall that  $\mathcal{N}'$  denotes the set of configurations given in [Lemma 6.4](#), and  $\mathcal{E}_{\mathcal{R}}$  the edge constraint of  $\mathcal{R}(\Pi_\Delta(z))$ . We start by defining a useful “intermediate” problem  $\Pi^*$ . We define the label set  $\Sigma^*$  of  $\Pi^*$  as the set containing all labels that appear in at least one of the node configurations in  $\mathcal{N}'$ , and we set the node constraint of  $\Pi^*$  to be  $\mathcal{N}^* := \mathcal{N}'$ . Moreover, we define the edge constraint  $\mathcal{E}^*$  to be the set of all cardinality-2 multisets  $D_1 D_2$  with  $D_1, D_2 \in \Sigma^*$  such that there exist labels  $B_1 \in D_1$  and  $B_2 \in D_2$  satisfying  $B_1 B_2 \in \mathcal{E}_{\mathcal{R}}$ . Note that, by definition, the labels in  $\Sigma^*$  are sets of labels from  $\Sigma_\Delta(z)$ . We prove that  $\Pi^*$  is at least 1 round easier than  $\Pi_\Delta(z)$ , unless  $\Pi_\Delta(z)$  is already 0-round solvable.

**Lemma 6.9.** *Assume that  $|z| \leq \Delta$  if  $\text{len}(z) = 0$ , and  $|z| \leq \Delta - 1$  otherwise, and let  $T$  be some positive integer. If  $\Pi_\Delta(z)$  can be solved in  $T$  rounds in the deterministic port numbering model, then  $\Pi^*$  can be solved in  $T - 1$  rounds.*

*Proof.* Assume that  $\Pi_\Delta(z)$  can be solved in  $T$  rounds. Then, by [Theorem 3.1](#), there is some  $(T - 1)$ -round algorithm  $\mathcal{A}$  solving  $\overline{\mathcal{R}}(\mathcal{R}(\Pi_\Delta(z)))$ . For each node  $v$ , denote the cardinality- $\Delta$  multiset of labels that  $\mathcal{A}$  outputs at the half-edges incident to  $v$  by  $\mathcal{A}(v) = \{D_1(v), \dots, D_\Delta(v)\}$ . We define a new  $(T - 1)$ -round algorithm  $\mathcal{A}'$  solving  $\Pi^*$  as follows. Each node  $v$  first executes  $\mathcal{A}$  (in  $T - 1$  rounds), and then (without any further communication) chooses a node configuration  $D'_1 \dots D'_\Delta \in \mathcal{N}^*$  satisfying  $D_k \subseteq D'_k$  for each  $1 \leq k \leq \Delta$ , and replaces  $D_k$  by  $D'_k$ , for each  $1 \leq k \leq \Delta$ . Such choices exist due to [Lemma 6.4](#) and the fact that  $\mathcal{N}^* = \mathcal{N}'$ . Observe that, for each half-edge  $h$ , the label (i.e., set)  $\mathcal{A}$  outputs at  $h$  is a subset of the label  $\mathcal{A}'$  outputs at  $h$ . It follows, by the correctness of  $\mathcal{A}$  and the definitions of both  $\mathcal{E}^*$  and the edge constraint of  $\overline{\mathcal{R}}(\mathcal{R}(\Pi_\Delta(z)))$ , that, for each edge  $e$ , the cardinality-2 multiset of labels that  $\mathcal{A}'$  outputs at the half-edges belonging to  $e$  is contained in  $\mathcal{E}^*$ . Hence,  $\mathcal{A}'$  produces a correct output, which concludes the proof.  $\square$

We now explicitly characterize the configurations allowed by the edge constraint of  $\Pi^*$ .

**Lemma 6.10.** *The edge constraint  $\mathcal{E}^*$  of  $\Pi^*$  contains precisely the configurations*

- $\langle\langle U_i, \ell(C) \rangle\rangle, \langle\langle P_j \rangle\rangle \langle\langle U_{i'} \rangle\rangle, \langle\langle P_{j'} \rangle\rangle$  for all  $C, C' \in 2^{\mathcal{C}} \setminus \{\emptyset\}$ ,  $0 \leq i, i', j, j' \leq \beta$  satisfying  $\text{level}(C) \leq j, \text{level}(C') \leq j', i \in \{j, j + 1\}, i' \in \{j', j' + 1\}$ , and
  1.  $C \cap C' = \emptyset$ , or
  2.  $i < j'$ , or
  3.  $i' < j$ ,
- $\langle\langle U_i, \ell(C) \rangle\rangle, \langle\langle P_j \rangle\rangle \langle\langle U_{i'} \rangle\rangle$  for all  $C \in 2^{\mathcal{C}} \setminus \{\emptyset\}$ ,  $0 \leq i, j \leq \beta, 1 \leq i' \leq \beta$  satisfying  $\text{level}(C) \leq j, i \in \{j, j + 1\}$
- $\langle\langle U_i, \ell(C) \rangle\rangle, \langle\langle P_j \rangle\rangle \langle\langle P_{j'} \rangle\rangle$  for all  $C \in 2^{\mathcal{C}} \setminus \{\emptyset\}$ ,  $0 \leq i, j \leq \beta, 1 \leq j' \leq \beta$  satisfying  $\text{level}(C) \leq j, i \in \{j, j + 1\}$ , and  $i < j'$ ,
- $\langle\langle U_i \rangle\rangle \langle\langle U_{i'} \rangle\rangle$  for all  $1 \leq i, i' \leq \beta$
- $\langle\langle U_i \rangle\rangle \langle\langle P_j \rangle\rangle$  for all  $1 \leq i < j \leq \beta$

- $\langle\langle X \rangle\rangle L$ , for each  $L \in \Sigma^*$ ,

where, as before, we set  $U_0 := X$  and  $\langle\langle U_i, \ell(C) \rangle\rangle, \langle P_0 \rangle := \langle\langle U_i, \ell(C) \rangle\rangle$ .

*Proof.* Consider some arbitrary sets  $B_1, B_2, B'_1, B'_2 \in \Sigma_{\mathcal{R}}$  satisfying  $\langle B_1, B_2 \rangle, \langle B'_1, B'_2 \rangle \in \Sigma^*$ . We claim that  $\langle B_1, B_2 \rangle \langle B'_1, B'_2 \rangle \in \mathcal{E}^*$  if and only if at least one of the configurations  $B_1 B'_1, B_1 B'_2, B_2 B'_1, B_2 B'_2$  is contained in  $\mathcal{E}_{\mathcal{R}}$ . The “if” part follows from the definitions of  $\langle \cdot \rangle$  and  $\mathcal{E}^*$ . Now, assume that  $\langle B_1, B_2 \rangle \langle B'_1, B'_2 \rangle \in \mathcal{E}^*$ . Then, by the definition of  $\mathcal{E}^*$ , there is some  $(B_3, B'_3) \in \langle B_1, B_2 \rangle \times \langle B'_1, B'_2 \rangle$  satisfying  $B_3 B'_3 \in \mathcal{E}_{\mathcal{R}}$ . By the definition of  $\langle \cdot \rangle$ , there exist some  $q, r \in \{1, 2\}$  such that  $B_q \leq B_3$  and  $B'_r \leq B'_3$ . By [Lemma 6.3](#), it follows that  $B_q \subseteq B_3$  and  $B'_r \subseteq B'_3$ . Since  $B_3 B'_3 \in \mathcal{E}_{\mathcal{R}}$ , we obtain  $B_q B_r \in \mathcal{E}_{\mathcal{R}}$ , by the definition of  $\mathcal{E}_{\mathcal{R}}$ . This proves the claim.

Now, consider arbitrary labels  $A_1, A_2, A'_1, A'_2 \in \Sigma_{\Delta}(z)$  satisfying  $\langle A_1, A_2 \rangle, \langle A'_1, A'_2 \rangle \in \Sigma_{\mathcal{R}}$ . By the definitions of  $\langle \cdot \rangle$  and  $\mathcal{E}_{\mathcal{R}}$ , we obtain that  $\langle A_1, A_2 \rangle \langle A'_1, A'_2 \rangle \in \mathcal{E}_{\mathcal{R}}$  if and only if all of the configurations  $A_1 A'_1, A_1 A'_2, A_2 A'_1, A_2 A'_2$  are contained in  $\mathcal{E}_{\Delta}(z)$ . Note that all of the above considerations also hold if the two arguments in  $\langle \cdot, \cdot \rangle$  are identical.

The above insights provide a simple method for checking whether a given configuration  $D_1 D_2 = \langle\langle A_{1,1}, A_{1,2} \rangle\rangle, \langle\langle A_{2,1}, A_{2,2} \rangle\rangle \langle\langle A'_{1,1}, A'_{1,2} \rangle\rangle, \langle\langle A'_{2,1}, A'_{2,2} \rangle\rangle$  with labels from  $\Sigma^*$  is contained in  $\mathcal{E}^*$ : it is contained if and only if the following Boolean formula in disjunctive normal form is satisfied:

$$\bigvee_{q,r \in \{1,2\}} \bigwedge_{s,t \in \{1,2\}} A_{q,s} A'_{r,t} \in \mathcal{E}_{\Delta}(z).$$

We illustrate this method for the arguably most complicated situation, namely the one where  $D_1$  and  $D_2$  are of the form  $\langle\langle U_i, \ell(C) \rangle\rangle, \langle P_j \rangle$ . More precisely, consider the configuration  $D_1 D_2 = \langle\langle U_i, \ell(C) \rangle\rangle, \langle P_j \rangle \langle\langle U_{i'}, \ell(C') \rangle\rangle, \langle P_{j'} \rangle$  where  $\mathcal{C}, \mathcal{C}' \in 2^{\mathcal{E}} \setminus \{\emptyset\}$ ,  $0 \leq i, i', j, j' \leq \beta$  and  $\text{level}(\mathcal{C}) \leq j, \text{level}(\mathcal{C}') \leq j', i \in \{j, j+1\}, i' \in \{j', j'+1\}$ . By the above discussion, we obtain that  $D_1 D_2 \in \mathcal{E}^*$  if and only if at least one of  $\langle U_i, \ell(C) \rangle \langle U_{i'}, \ell(C') \rangle, \langle U_i, \ell(C) \rangle \langle P_{j'} \rangle, \langle P_j \rangle \langle U_{i'}, \ell(C') \rangle, \langle P_j \rangle \langle P_{j'} \rangle$  is contained in  $\mathcal{E}_{\mathcal{R}}$ , which in turn happens if we have

1.  $U_i U_{i'} \in \mathcal{E}_{\Delta}(z)$  and  $U_i \ell(C') \in \mathcal{E}_{\Delta}(z)$  and  $\ell(C) U_{i'} \in \mathcal{E}_{\Delta}(z)$  and  $\ell(C) \ell(C') \in \mathcal{E}_{\Delta}(z)$ , or
2.  $U_i P_{j'} \in \mathcal{E}_{\Delta}(z)$  and  $\ell(C) P_{j'} \in \mathcal{E}_{\Delta}(z)$ , or
3.  $P_j U_{i'} \in \mathcal{E}_{\Delta}(z)$  and  $P_j \ell(C') \in \mathcal{E}_{\Delta}(z)$ , or
4.  $P_j P_{j'} \in \mathcal{E}_{\Delta}(z)$ .

If  $j = 0$ , resp.  $j' = 0$ , then the statements about the configurations containing  $P_j$ , resp.  $P_{j'}$ , do not appear. Now, by applying [Observation 5.1](#) to each of the 9 substatements, we obtain that  $D_1 D_2 \in \mathcal{E}^*$  if and only if

1.  $\mathcal{C} \cap \mathcal{C}' = \emptyset$ , or
2.  $i < j'$ , or
3.  $i' < j$ ,

Note that we used here that  $\text{level}(\mathcal{C}) \leq i$  and  $\text{level}(\mathcal{C}') \leq i'$ .

In general, since each label in  $D_1 D_2$  is of one of the four forms  $\langle\langle P_i \rangle\rangle, \langle\langle U_i \rangle\rangle, \langle\langle U_i, \ell(C) \rangle\rangle, \langle P_j \rangle$ ,  $\langle\langle X \rangle\rangle$ , we obtain the lemma by applying an analogous argumentation to each pair of forms (which will yield the restrictions for  $\mathcal{C}, \mathcal{C}', i, i', j, j'$  such that  $D_1 D_2 \in \mathcal{E}^*$  where  $D_1$  and  $D_2$  are of the chosen forms).  $\square$

We now prove that  $\Pi_\Delta(\text{prefix}(z))$  is a relaxation of  $\overline{\mathcal{R}}(\mathcal{R}(\Pi_\Delta(z)))$ .

**Lemma 6.11.** *Assume that  $|z| \leq \Delta$  if  $\text{len}(z) = 0$ , and  $|z| \leq \Delta - 1$  otherwise, and let  $T$  be some positive integer. If  $\Pi_\Delta(z)$  can be solved in  $T$  rounds in the deterministic port numbering model, then  $\Pi(z')$  can be solved in  $T - 1$  rounds, where  $z' = \text{prefix}(z)$ .*

*Proof.* Let  $\mathfrak{C}$  denote the set of colors in  $\Pi_\Delta(z)$ . We identify the set  $\mathfrak{C}'$  of colors in  $\Pi_\Delta(z')$  with the set of pairs  $(C, i)$  satisfying  $C \in \mathfrak{C}$  and  $\text{level}(C) \leq i \leq \beta$ . For some color  $(C, i) \in \mathfrak{C}'$ , set  $\text{level}((C, i)) = i$ . This definition ensures that, for each  $0 \leq j \leq \beta$ , there are precisely  $z'_j = z_0 + \dots + z_j$  many colors of level  $j$  in  $\Pi_\Delta(z')$ , as required.

Assume that  $\Pi_\Delta(z)$  can be solved in  $T$  rounds. Then, by Lemma 6.9, there is some  $(T - 1)$ -round algorithm  $\mathcal{A}$  solving  $\Pi^*$ . We define a new  $(T - 1)$ -round algorithm  $\mathcal{A}'$  as follows. Each node  $v$  first executes algorithm  $\mathcal{A}$  (in  $T - 1$  rounds). Then (without any further communication)  $v$  replaces, for each incident half-edge  $h$ , the label  $\ell$  that  $\mathcal{A}$  produced on  $h$  by a new label  $\ell'$  according to the following rules.

If  $\ell = \langle\langle X \rangle\rangle$ , then  $\ell' := X$ . If  $\ell = \langle\langle U_i \rangle\rangle$  for some  $1 \leq i \leq \beta$ , then  $\ell' := U_i$ . If  $\ell = \langle\langle P_i \rangle\rangle$  for some  $1 \leq i \leq \beta$ , then  $\ell' := P_i$ . If  $\ell = \langle\langle U_i, \ell(C) \rangle\rangle, \langle\langle P_j \rangle\rangle$  for some  $C \in 2^{\mathfrak{C}} \setminus \{\emptyset\}, 0 \leq i, j \leq \beta$  satisfying  $\text{level}(C) \leq j$  and  $i \in \{j, j + 1\}$ , then  $\ell' := \ell(\mathcal{C}^{\text{new}})$  where  $\mathcal{C}^{\text{new}} := \{(C, j) \mid C \in \mathfrak{C}\} \cup \{(C, i) \mid C \in \mathfrak{C}\}$ . (Here, as before, we set  $U_0 := X$  and  $\langle\langle U_i, \ell(C) \rangle\rangle, \langle\langle P_0 \rangle\rangle := \langle\langle U_i, \ell(C) \rangle\rangle$ .)

Finally, consider only those nodes  $v$  such that there is a half-edge incident to  $v$  on which  $\ell' = \ell(\mathcal{C}^{\text{new}})$  for some  $\mathcal{C}^{\text{new}} \in 2^{\mathfrak{C}'} \setminus \{\emptyset\}$ . (We remark that, by the design of the replacement rules and the definition of  $\mathcal{N}^*$ , for any such node  $v$ , each incident half-edge is labeled with either  $\ell(\mathcal{C}^{\text{new}})$  or  $X$  at this point.) Each such node  $v$  selects  $|\mathcal{C}^{\text{new}}| - 1 - x$  half-edges currently labeled with  $\ell(\mathcal{C}^{\text{new}})$  and replaces the labels on those half-edges with the label  $X$ . Here  $x$  denotes the number of half-edges incident to  $v$  that were already labeled with  $X$ . (We will see a bit later that  $|\mathcal{C}^{\text{new}}| - 1 - x$  is always nonnegative.) This concludes the description of  $\mathcal{A}'$ .

We claim that  $\mathcal{A}'$  solves  $\Pi_\Delta(z')$ . We first consider the produced node configurations. Consider some node  $v$ , and let  $N = \ell_1 \dots \ell_\Delta$  denote the node configuration on the half-edges incident to  $v$  produced by  $\mathcal{A}'$ . The replacement rules, together with the definition of  $\mathcal{N}^*$  and the correctness of  $\mathcal{A}$ , imply that  $N \in \mathcal{N}_\Delta(z')$ , provided that  $\mathcal{A}'$  is well-defined. Hence, the only thing that we have to show is that “ $|\mathcal{C}^{\text{new}}| - 1 - x$  is always nonnegative” as claimed in the algorithm description.

Thus, assume that there is a half-edge  $h$  incident to  $v$  that, before the final step of the algorithm, was labeled with  $\ell' = \ell(\mathcal{C}^{\text{new}})$  for some  $\mathcal{C}^{\text{new}} \in 2^{\mathfrak{C}'} \setminus \{\emptyset\}$ . Moreover, let  $x$  be as defined above. By the design of the replacement rules, we see that there are some  $C \in 2^{\mathfrak{C}} \setminus \{\emptyset\}, 0 \leq i, j \leq \beta$  such that the output of  $\mathcal{A}$  on  $h$  was  $\ell = \langle\langle U_i, \ell(C) \rangle\rangle, \langle\langle P_j \rangle\rangle$ . Furthermore, by the definition of  $\mathcal{N}^*$ , we have  $x = |C| + i - j - 1$ , i.e.,  $x = |C| - 1$  if  $i = j$ , and  $x = |C|$  if  $i = j + 1$ . Observe that the definition of  $\mathcal{C}^{\text{new}}$  (in the replacement rule) ensures that  $|\mathcal{C}^{\text{new}}| \geq |C|$  if  $i = j$ , and  $|\mathcal{C}^{\text{new}}| \geq |C| + 1$  if  $i = j + 1$ . Hence, in either case  $|\mathcal{C}^{\text{new}}| - 1 - x \geq 0$ , and we obtain that each node configuration produced by  $\mathcal{A}'$  is contained in  $\mathcal{N}_\Delta(z')$ , as desired.

Now consider the edge configurations produced by  $\mathcal{A}'$ . Consider an arbitrary edge  $e$  and let  $D = D_1 D_2$  denote the edge configuration produced by  $\mathcal{A}'$  on the two half-edges  $h_1, h_2$  belonging to  $e$  (where output  $D_1$  is given on half-edge  $h_1$ ). The replacement rules, together with Lemma 6.10 clearly imply that, if  $D_1, D_2 \in \{U_1, \dots, U_\Delta\} \cup \{P_1, \dots, P_\Delta\}$ ,  $D_1 = X$ , or  $D_2 = X$ , then  $D \in \mathcal{E}_\Delta(z')$ . Hence, assume w.l.o.g., that  $D_1 = \mathcal{C}^{\text{new}}$  for some  $\mathcal{C}^{\text{new}} \in 2^{\mathfrak{C}'} \setminus \{\emptyset\}$ , and that  $D_2$  is some color set from  $2^{\mathfrak{C}'} \setminus \{\emptyset\}$ , some  $U_{i'}$  or some  $P_{j'}$ . Let  $\langle\langle U_i, \ell(C) \rangle\rangle, \langle\langle P_j \rangle\rangle$  be the output of  $\mathcal{A}$  on  $h_1$ , which implies that  $\text{level}(\mathcal{C}^{\text{new}}) = i$ . If  $D_2 = U_{i'}$  for some  $1 \leq i' \leq \beta$ , then  $D \in \mathcal{E}_\Delta(z')$ , simply by the definition of  $\mathcal{E}_\Delta(z')$ .

If  $D_2 = P_{j'}$  for some  $1 \leq j' \leq \beta$ , then we know that the output label of  $\mathcal{A}$  on  $h_2$  is  $\langle\langle P_{j'} \rangle\rangle$ . By Lemma 6.10, we obtain  $i < j'$ , which together with the aforementioned  $\text{level}(\mathcal{C}^{\text{new}}) = i$  (and the definition of  $\mathcal{E}_\Delta(z')$ ) implies  $D \in \mathcal{E}_\Delta(z')$ .



Finally, assume that  $D_2 = \mathcal{C}_2^{\text{new}}$ , and let  $\langle\langle U_{i'}, \ell(\mathcal{C}') \rangle\rangle, \langle P_{j'} \rangle\rangle$  be the output of  $\mathcal{A}$  on  $h_2$ . By [Lemma 6.10](#), we have  $\mathcal{C} \cap \mathcal{C}' = \emptyset$ , or  $i < j'$ , or  $i' < j$ . By the definitions of  $\mathcal{C}^{\text{new}}$  and  $\mathcal{C}_2^{\text{new}}$ , we have that, if  $\mathcal{C} \cap \mathcal{C}' = \emptyset$ , or  $i < j'$ , or  $i' < j$ , then  $\mathcal{C}^{\text{new}} \cap \mathcal{C}_2^{\text{new}} = \emptyset$ . Hence,  $D \in \mathcal{E}_\Delta(z')$ , and we conclude that the output produced by  $\mathcal{A}'$  is indeed a correct solution for  $\Pi_\Delta(z)$ .  $\square$

## 6.4 Zero Round Solvability

We now characterize for which parameters the problems in our family cannot be solved in 0 rounds.

**Lemma 6.12.** *Assume that  $|z| \leq \Delta$ . The problem  $\Pi_\Delta(z)$  is not 0-round solvable in the deterministic port numbering model if the port numbering is unconstrained.*

*Proof.* Since for a 0-round algorithm the view of each node is the same, then all nodes must output the same configuration. We show that, for all the allowed configurations of the node constraints it holds that they are not “self compatible”, meaning that they contain two labels  $\ell_1, \ell_2$ , that are not edge-compatible, that is,  $\ell_1 \ell_2 \notin \mathcal{E}_\Delta(z)$ , implying that for some port numbering assignment the 0 round algorithm must fail.

The allowed node configurations are the following:

- $\ell(\mathcal{C})^{\Delta-x} \mathbf{X}^x$ , for each  $\mathcal{C} \in 2^{\mathcal{E}} \setminus \{\emptyset\}$ , where  $x = |\mathcal{C}| - 1$ .
- $P_i U_i^{\Delta-1}$ , for each  $1 \leq i \leq \beta$ .

Consider the configurations of the first type. Since  $|z| \leq \Delta$ , then  $|\mathcal{C}| \leq \Delta$ , implying that  $x = |\mathcal{C}| - 1 \leq \Delta - 1$ , and hence that  $\ell(\mathcal{C})$  appears at least once in the configuration. Since, for all  $\mathcal{C} \in 2^{\mathcal{E}} \setminus \{\emptyset\}$  it holds that  $\mathcal{C} \cap \mathcal{C} \neq \emptyset$ , then  $\ell(\mathcal{C})$  is not edge-compatible with itself. Consider the configurations of the second type. For all  $i$ ,  $P_i$  is not edge-compatible with itself.  $\square$

## 6.5 A Problem Sequence

We now combine [Lemma 6.1](#) and [Lemma 6.12](#) to show that there exists a problem sequence that satisfies some desirable properties. Intuitively, if a problem in the sequence is  $\Pi_\Delta(z)$ , then the next problem is obtained by taking the prefix sum of  $z$ . If we apply this procedure many times, and we start from a “small” vector it will take some steps before obtaining a “large” vector. The number of steps required to reach a vector  $z'$  such that  $|z'| \geq \Delta$  (or  $> \Delta$  if  $\text{len}(z') = 0$ ) is a lower bound on the length of the sequence.

**Lemma 6.13.** *There exists a problem sequence  $\Pi_0 \rightarrow \Pi_1 \rightarrow \dots \rightarrow \Pi_t$ , where  $\Pi_1 = \Pi_\Delta(z)$ , such that, for all  $0 \leq i < t$ , the following holds:*

- *There exists a problem  $\Pi'_i$  that is a relaxation of  $\mathcal{R}(\Pi_i)$ ;*
- *$\Pi_{i+1}$  is a relaxation of  $\overline{\mathcal{R}}(\Pi'_i)$ ;*
- *The number of labels of  $\Pi_i$ , and the ones of  $\Pi'_i$ , are upper bounded by  $2^\Delta(1 + \text{len}(z))$ .*

*Also,  $\Pi_t$  has at most  $2^\Delta(1 + \text{len}(z))$  labels and is not 0-round solvable in the deterministic port numbering model if the port numbering assignment is unconstrained. Let  $\beta = \text{len}(z)$ . The value of  $t$  is at least the maximum  $i$  such that  $|\text{prefix}^i(z)| < \Delta$  if  $\text{len}(z) > 0$  and such that  $|\text{prefix}^i(z)| \leq \Delta$  if  $\text{len}(z) = 0$ . If there is no maximum, then  $t$  can be made arbitrarily large.*

*Proof.* We set  $\Pi_i = \Pi_\Delta(\text{prefix}^i(z))$ , where  $\text{prefix}^i(z)$  is the result obtained by recursively applying the function  $\text{prefix}$  for  $i$  times, and  $\text{prefix}^0(z) = z$ . By [Lemma 6.1](#) the problem sequence satisfies the requirements. We need to prove for which values of  $t$ ,  $\Pi_t$  is not 0-rounds solvable. By [Lemma 6.12](#) it is sufficient to satisfy  $|\text{prefix}^t(z)| \leq \Delta$ , which is a condition required by the statement.  $\square$

## 7 Lifting Results to the LOCAL Model

In this section we show how to lift the results that we obtained for the deterministic port numbering model to the LOCAL model. For this purpose, we heavily rely on ideas already used in the literature [BBH<sup>+</sup>19, BBO20, BO20, BBE<sup>+</sup>20], by slightly improving some parts to support our use case. While in this work we only consider graphs, we show a general statement that works on hypergraphs as well.

In the following, we assume w.l.o.g. that  $\delta$ , the rank of the hyperedges, is at most  $\Delta$ , the degree of the nodes. If not, we can just switch the roles of nodes and hyperedges, and the same claim still holds (up to an additive constant factor).

**Theorem 7.1.** *Let  $\Pi_0 \rightarrow \Pi_1 \rightarrow \dots \rightarrow \Pi_t$  be a sequence of problems. Assume that, for all  $0 \leq i < t$ , and for some function  $f$ , the following holds:*

- *There exists a problem  $\Pi'_i$  that is a relaxation of  $\mathcal{R}(\Pi_i)$ ;*
- *$\Pi_{i+1}$  is a relaxation of  $\overline{\mathcal{R}}(\Pi'_i)$ ;*
- *The number of labels of  $\Pi_i$ , and the ones of  $\Pi'_i$ , are upper bounded by  $f(\Delta)$ .*

*Also, assume that  $\Pi_t$  has at most  $f(\Delta)$  labels and is not 0-round solvable in the deterministic port numbering model, even if the port numbering assignment satisfies some local constraints  $C^{\text{port}}$ . Then,  $\Pi_0$  requires  $\Omega(\min\{t, \log_\Delta n - \log_\Delta \log f(\Delta)\})$  rounds in the deterministic LOCAL model and  $\Omega(\min\{t, \log_\Delta \log n - \log_\Delta \log f(\Delta)\})$  rounds in the randomized LOCAL model, even if the port numbering satisfies some local constraints  $C^{\text{port}}$ .*

Since all the techniques required to prove [Theorem 7.1](#) are already present in the literature, and we only need to improve the analysis to support an arbitrary large number of labels, we defer the proof of [Theorem 7.1](#) to [Appendix A](#).

Note that, by proving that a problem can be relaxed to a fixed point, then we obtain an arbitrary long sequence of problems that satisfy the requirements. Hence, if the number of labels used in the fixed point problem is not too large, then we get that this problem requires  $\Omega(\log_\Delta n)$  rounds for deterministic algorithms, and  $\Omega(\log_\Delta \log n)$  rounds for randomized ones.

## 8 Lower Bound Results

All the results that we present in this section hold already on  $\Delta$ -regular trees. We start by showing a general result, and then we instantiate it to obtain results for a number of *natural* problems. By combining [Lemma 6.13](#) with [Theorem 7.1](#), we obtain the following theorem.

**Theorem 8.1.** *Let  $z$  be a vector  $[z_0, \dots, z_\beta]$  such that  $\beta \leq 2^\Delta$ . Let  $t$  be the maximum value such that  $|\text{prefix}^t(z)| \leq \Delta$  if  $\text{len}(z) = 0$ , and such that  $|\text{prefix}^t(z)| < \Delta$  if  $\text{len}(z) > 0$ . If there is no maximum, let  $t = \infty$ . Then,  $\Pi_\Delta(z)$  requires  $\Omega(\min\{t, \log_\Delta n\})$  rounds in the deterministic LOCAL model and  $\Omega(\min\{t, \log_\Delta \log n\})$  rounds in the randomized LOCAL model.*

Notice that the assumption on  $\beta$  is only required in order to make the total number of labels small enough so that  $\log_\Delta \log f(\Delta) = O(1)$ , where  $f(\Delta)$  is an upper bound on the number of labels of the problems in the sequence.

**Arbdefective Colorings.** We start by proving that  $\Pi_\Delta([\Delta])$  is a relaxation of some variants of arbdefective coloring, including the  $\Delta$ -coloring problem itself.

**Theorem 8.2.** *Let  $\vec{d} = (d_1, \dots, d_c)$  be an arbdefect vector that is not 1-relaxed. Then, the  $\vec{d}$ -arbdefective  $c$ -coloring problem requires  $\Omega(\log_\Delta n)$  rounds in the deterministic LOCAL model and  $\Omega(\log_\Delta \log n)$  rounds in the randomized LOCAL model.*

*Proof.* First of all, notice that  $\text{prefix}([\Delta]) = [\Delta]$ , implying that for  $\Pi_\Delta([\Delta])$  we can apply [Theorem 8.1](#) with  $t = \infty$ . We show that, if  $\vec{d}$  is not 1-relaxed, then the  $\vec{d}$ -arbdefective  $c$ -coloring problem can be relaxed to  $\Pi_\Delta([\Delta])$ . Since by assumption  $\vec{d}$  is not 1-relaxed, by definition we have that:

$$\text{cap}(\vec{d}) = \sum_{i=1}^c (d_i + 1) \leq \Delta.$$

Consider a color space  $S$  of size  $\sum_{i=1}^c (d_i + 1)$ . We partition the colors of  $S$  into  $c$  sets, where each set  $\mathcal{C}_i$  contains  $d_i + 1$  colors. For each set  $\mathcal{C}_i$ , the problem  $\Pi_\Delta([\Delta])$  contains the configuration  $\ell(\mathcal{C}_i)^{\Delta-x} \mathsf{X}^x$ , where  $x = |\mathcal{C}_i| - 1 = d_i$ . Also, since all sets  $\mathcal{C}_i$  are pairwise disjoint, then  $\ell(\mathcal{C}_i)\ell(\mathcal{C}_j)$  is contained in the edge constraint, for all  $i \neq j$ . Moreover,  $\mathsf{X}$  is compatible with every other label. Hence, given a solution for arbdefective coloring, nodes of color  $i$  can label  $\ell(\mathcal{C}_i)$  their edges, and then mark their arbdefective edges with  $\mathsf{X}$ . The number of arbdefective edges is at most  $d_i$ . Nodes can then mark other arbitrary edges to make the number of  $\mathsf{X}$  exactly  $d_i$ . In this way, they obtain a solution for  $\Pi_\Delta([\Delta])$ .  $\square$

By allowing the same arbdefect to each color, we obtain the following corollary.

**Corollary 8.3.** *The  $d$ -arbdefective  $c$ -coloring problem requires  $\Omega(\log_\Delta n)$  rounds in the deterministic LOCAL model and  $\Omega(\log_\Delta \log n)$  rounds in the randomized LOCAL model, if  $(d + 1)c \leq \Delta$ .*

Also, by setting  $c = \Delta$  and  $d = 0$ , we obtain the following.

**Corollary 8.4.** *The  $\Delta$ -coloring problem requires  $\Omega(\log_\Delta n)$  rounds in the deterministic LOCAL model and  $\Omega(\log_\Delta \log n)$  rounds in the randomized LOCAL model.*

**Arbdefective Colored Ruling Sets.** We now consider  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling sets. Before showing our results, we prove some useful lemmas. We start by showing a lemma proved in [\[BBO20\]](#).

**Lemma 8.5** (Lemma 16 of [\[BBO20\]](#), arXiv version). *Assume  $z = [1, 0, \dots, 0]$  and  $k \leq \beta = \text{len}(z)$ . For all  $j \geq 1$ ,  $\text{prefix}^j(v)_k = \binom{j+k-1}{k}$ .*

We now show a lower bound for  $\Pi_\Delta(z)$  for the case in which only the first element of  $z$  is non-zero.

**Lemma 8.6.** *Assume  $z = [k, 0, \dots, 0]$ , where  $k \geq 1$ , and  $\beta = \text{len}(z) \leq 2^\Delta$ . Let  $t$  be the maximum value such that  $\binom{t+\beta}{\beta} < \Delta/k$ . The problem  $\Pi_\Delta(z)$  requires  $\Omega(\min\{t, \log_\Delta n\})$  rounds in the deterministic LOCAL model and  $\Omega(\min\{t, \log_\Delta \log n\})$  rounds in the randomized LOCAL model.*

*Proof.* Note that  $|\text{prefix}^t(z)| = \text{prefix}^{t+1}(z)_\beta$ , and  $\text{prefix}^{t+1}([k, 0, \dots, 0]) = k \cdot \text{prefix}^{t+1}([1, 0, \dots, 0])$ . Hence,  $\text{prefix}^{t+1}(z)_\beta = \text{prefix}^{t+1}([k, 0, \dots, 0])_\beta = k \cdot \text{prefix}^{t+1}([1, 0, \dots, 0])_\beta = k \binom{t+\beta}{\beta}$ , where the last equality holds by [Lemma 8.5](#). Hence, since by assumption  $t$  is the maximum value such that  $k \binom{t+\beta}{\beta} < \Delta$ , then, by [Theorem 8.1](#), the claim follows.  $\square$

Note that  $\Pi_\Delta([1, 0, \dots, 0])$  is a relaxation of a variant of  $(2, \beta)$ -ruling sets where nodes also need to know an upper bound on the distance, and the direction, towards a ruling set node. For this variant of ruling sets, in [BBO20, Lemma 15 (arXiv version)] it has been shown that it can be solved in  $t$  rounds, where  $t$  is the minimum value such that  $\binom{t+\beta}{\beta} \geq c$ , if nodes are provided a  $c$ -coloring in input. Lemma 8.6 implies that this result is tight, at least in the case in which a  $(\Delta + 1)$ -coloring is given in input to the nodes: it is not possible to solve this variant of ruling sets in  $t$  rounds, if  $\binom{t+\beta}{\beta} < \Delta$ . In fact, from Lemma 6.13 we actually get a lower bound of  $t + 1$  rounds, where  $t$  is the maximum value such that  $\binom{t+\beta}{\beta} < \Delta$ . Hence, our result is tight in the case in which a  $(\Delta + 1)$ -coloring is given to the nodes. Note that, improving this result for the case in which a larger coloring is given in input would imply that a  $(\Delta + 1)$ -vertex coloring cannot be computed fast, but knowing whether  $(\Delta + 1)$ -vertex coloring requires  $\omega(\log^* n)$  is a long standing open question.

We now use the result of Lemma 8.6 to show lower bounds for arbdefective colored ruling sets. In particular, we show that, a lower bound of  $t$  rounds for  $\Pi([c(1 + \alpha), 0, \dots, 0])$  implies a lower bound of  $t - \beta$  rounds for  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling sets.

**Theorem 1.5.** *Let  $t$  be the maximum value such that  $\binom{t+\beta}{\beta} < \frac{\Delta}{c(\alpha+1)}$ . The  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling set problem requires  $\Omega(\min\{t, \log_\Delta n\} - \beta)$  in the deterministic LOCAL model and  $\Omega(\min\{t, \log_\Delta \log n\} - \beta)$  in the randomized LOCAL model.*

*Proof.* Assume  $z = [c(1 + \alpha), 0, \dots, 0]$  and  $\text{len}(z) = \beta$ . By Lemma 8.6, the problem  $\Pi_\Delta(z)$  requires  $\Omega(\min\{t, \log_\Delta n\})$  rounds in the deterministic LOCAL model and  $\Omega(\min\{t, \log_\Delta \log n\})$  rounds in the randomized LOCAL model, where  $t$  is the maximum value such that  $c(1 + \alpha)\binom{t+\beta}{\beta} < \Delta$ .

We show that, given a solution for the  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling set problem, nodes can spend  $\beta$  rounds to solve  $\Pi_\Delta(z)$ . Hence, a lower bound of  $t$  rounds for  $\Pi_\Delta(z)$  implies a lower bound of  $t - \beta$  for the  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling set problem. In  $\beta$  rounds, uncolored nodes can find the nearest colored node  $v$ , and nodes at distance  $i$  from  $v$  can output  $P_i$  on their port towards  $v$  and  $U_i$  on all the other ports. Then, since the color space contains  $c(1 + \alpha)$  colors, then there must exist  $c$  disjoint sets  $\mathcal{C}_i$  of colors of size  $1 + \alpha$ . For each set  $\mathcal{C}_i$ , there must exist an allowed configuration of the form  $\ell(\mathcal{C}_i^{\Delta-x})\mathbf{X}^x$ , for  $x = \alpha$ . Nodes of color  $i$  can label their edges  $\ell(\mathcal{C}_i)$ , and then mark their edges oriented towards other nodes of the same color with  $\mathbf{X}$ . The number of edges marked  $\mathbf{X}$  is at most  $\alpha$ . Nodes can then mark other arbitrary edges to make the number of  $\mathbf{X}$  exactly  $\alpha$ . It is easy to check that this labeling satisfies the constraints of  $\Pi_\Delta(z)$ .  $\square$

We now provide some more human-friendly versions of Theorem 1.5. We start by making an observation.

**Observation 8.7.** *For all  $t \leq \frac{\beta}{2e}x^{1/\beta}$  and  $\beta \leq \frac{1}{2e} \log x$ , it holds that  $\binom{t+\beta}{\beta} \leq x$ .*

*Proof.*

$$\binom{t+\beta}{\beta} \leq \left(\frac{e(t+\beta)}{\beta}\right)^\beta \leq \left(\frac{e(\frac{\beta}{2e}x^{1/\beta} + \frac{1}{2e} \log x)}{\beta}\right)^\beta \leq \left(\frac{e(\frac{\beta}{e}x^{1/\beta})}{\beta}\right)^\beta = x$$

$\square$

We are now ready to provide an explicit lower bound for  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling sets as a function of  $\Delta$  and  $n$ .

**Corollary 8.8.** *There exists a constant  $\varepsilon_0 > 0$  such that the  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling set problem requires  $\Omega(\min\{\beta(\frac{\Delta}{c(1+\alpha)})^{1/\beta}, \log_\Delta n\})$  rounds, for  $\beta \leq \varepsilon_0 \min\{\log \frac{\Delta}{c(1+\alpha)}, \log_\Delta n\}$ , in the deterministic LOCAL model, and it requires  $\Omega(\min\{\beta(\frac{\Delta}{c(1+\alpha)})^{1/\beta}, \log_\Delta \log n\})$  rounds, for  $\beta \leq \varepsilon_0 \min\{\log \frac{\Delta}{c(1+\alpha)}, \log_\Delta \log n\}$ , in the randomized LOCAL model.*

*Proof.* Let  $t$  be the maximum value such that  $\binom{t+\beta}{\beta} < \frac{\Delta}{c(1+\alpha)}$ . By [Observation 8.7](#) and the assumption on  $\beta$ ,  $t = \Omega(\beta(\frac{\Delta}{c(1+\alpha)})^{1/\beta})$ . By the assumption on  $\beta$ ,  $t - \beta = \Omega(t)$ . Also, by the assumption on  $\beta$  in the case of deterministic algorithms,  $\log_{\Delta} n - \beta = \Omega(\log_{\Delta} n)$ , and in the case of randomized algorithms,  $\log_{\Delta} \log n - \beta = \Omega(\log_{\Delta} \log n)$ . By [Theorem 1.5](#), the claim follows.  $\square$

**Corollary 8.9.** *There exists some constant  $\varepsilon > 0$  such that the  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling set problem requires  $\Omega(\frac{\log n}{\log(1+\alpha)+\log c+\beta \log \log n})$  rounds for  $\beta \leq \varepsilon \sqrt{\frac{\log n}{\log(1+\alpha)+\log c+\log \log n}}$  in the deterministic LOCAL model, and  $\Omega(\frac{\log \log n}{\log(1+\alpha)+\log c+\beta \log \log \log n})$  rounds for  $\beta \leq \varepsilon \sqrt{\frac{\log \log n}{\log(1+\alpha)+\log c+\log \log \log n}}$  in the randomized LOCAL model.*

*Proof.* Apply [Corollary 8.8](#) by picking  $\Delta \approx c(1+\alpha)(\frac{\log n}{\log \log n})^{\beta}$  in the deterministic case, and  $\Delta \approx c(1+\alpha)(\frac{\log \log n}{\log \log \log n})^{\beta}$  in the randomized one.  $\square$

By setting  $\alpha = 0$  and  $c = 1$ , from [Theorem 1.5](#) and [Corollaries 8.8](#) and [8.9](#) we obtain lower bounds for  $(2, \beta)$ -ruling sets.

**Corollary 8.10.** *Let  $t$  be the maximum value such that  $\binom{t+\beta}{\beta} < \Delta$ , where  $\beta \leq 2^{\Delta}$ . The  $(2, \beta)$ -ruling set problem requires  $\Omega(\min\{t, \log_{\Delta} n\} - \beta)$  rounds in the deterministic LOCAL model and  $\Omega(\min\{t, \log_{\Delta} \log n\} - \beta)$  rounds in the randomized LOCAL model.*

*In particular, there exists some constant  $\varepsilon > 0$  such that the  $(2, \beta)$ -ruling set problem requires  $\Omega(\min\{\beta \Delta^{1/\beta}, \log_{\Delta} n\})$  rounds for  $\beta \leq \varepsilon \min\{\log \Delta, \log_{\Delta} n\}$  in the deterministic LOCAL model, and  $\Omega(\min\{\beta \Delta^{1/\beta}, \log_{\Delta} \log n\})$  rounds for  $\beta \leq \varepsilon \min\{\log \Delta, \log_{\Delta} \log n\}$  in the randomized one.*

*If we express the complexity of the problem solely as a function of  $n$ , then it requires  $\Omega(\frac{\log n}{\beta \log \log n})$  rounds for  $\beta \leq \varepsilon \sqrt{\frac{\log n}{\log \log n}}$  for deterministic algorithms, and  $\Omega(\frac{\log \log n}{\beta \log \log \log n})$  rounds for  $\beta \leq \varepsilon \sqrt{\frac{\log \log n}{\log \log \log n}}$  for randomized ones.*

Then, by setting  $\beta = 1$ , we obtain a lower bound for MIS.

**Theorem 8.11.** *The maximal independent set problem requires  $\Omega(\min\{\Delta, \log_{\Delta} n\})$  rounds in the deterministic LOCAL model and  $\Omega(\min\{\Delta, \log_{\Delta} \log n\})$  rounds in the randomized LOCAL model. If we express the complexity of the problem solely as a function of  $n$ , then it requires  $\Omega(\frac{\log n}{\log \log n})$  rounds for deterministic algorithms and  $\Omega(\frac{\log \log n}{\log \log \log n})$  rounds for randomized ones.*

In particular, we get that the maximal independent set problem requires  $\Omega(\log n / \log \log n)$  deterministic rounds on trees, which is tight, as in [\[BE10\]](#) it has been shown that MIS on trees can be solved in  $O(\log n / \log \log n)$  deterministic rounds.

**More Results.** Until now, we only considered cases where the vector  $z$  has non-zero values only in the first position. We conclude the section by showing that [Theorem 8.1](#) can be used to prove lower bounds for even more variants of ruling sets and arbdefective colorings.

Consider the following problem. We have two disjoint color spaces  $\mathcal{C}^1$  and  $\mathcal{C}^2$ , one containing 4 colors and the other containing 5 colors. Nodes can either be colored, or uncolored. Uncolored nodes must have either at least one node colored from the space  $\mathcal{C}^1$  at distance at most 3, or at least one node colored from the space  $\mathcal{C}^2$  at distance at most 2, or both. Nodes colored from the space  $\mathcal{C}^1$  must have arbdefect at most 1, while nodes colored from the space  $\mathcal{C}^2$  must have arbdefect at most 2. We show how to use [Theorem 8.1](#) to prove that this problem requires  $\Omega(\Delta^{1/3})$  rounds.

**Corollary 8.12.** *The example problem requires  $\Omega(\min\{\Delta^{1/3}, \log_{\Delta} n\})$  rounds in the deterministic LOCAL model and  $\Omega(\min\{\Delta^{1/3}, \log_{\Delta} \log n\})$  rounds in the randomized LOCAL model.*

*Proof.* Consider the problem  $\Pi_\Delta(z)$  where  $z = [8, 15, 0, 0]$ . We show that, given a solution for the example problem, nodes can solve in  $O(1)$  rounds the problem  $\Pi_\Delta(z)$ .

Uncolored nodes can spend 3 rounds to see the color of the nearest colored node. Nodes at distance  $i$ , can output  $P_i$  on the port towards the colored node if the colored node has a color from  $\mathcal{C}^1$ , or  $P_{i+1}$  if the colored node has a color from  $\mathcal{C}^2$ , and if both cases apply, nodes from  $\mathcal{C}^1$  are prioritized. Note that the assumptions guarantee that only labels from  $\{P_1, P_2, P_3\}$  are required. Uncolored nodes that output  $P_i$  on one port, then output  $U_i$  on all the other ports.

Colored nodes proceed as follows. There are 8 colors in level 0, and 15 colors in level 1. We partition the 8 colors in 4 groups of 2 colors,  $\{\mathcal{C}_i^1 \mid 1 \leq i \leq 4\}$ , and we partition the 15 colors in 5 groups of 3 colors,  $\{\mathcal{C}_i^2 \mid 1 \leq i \leq 5\}$ . In  $\Pi_\Delta(z)$  there must exist the following configurations:  $\ell(\mathcal{C}_i^1)^{\Delta-1}\mathbf{X}$  such that  $\text{level}(\mathcal{C}_i^1) = 0$  for  $i \in \{1, 2, 3, 4\}$ , and  $\ell(\mathcal{C}_i^2)^{\Delta-2}\mathbf{X}^2$  such that  $\text{level}(\mathcal{C}_i^2) = 1$  for  $i \in \{1, 2, 3, 4, 5\}$ , such that the intersection between each distinct pair of color sets is disjoint. As in the arbdefective colored ruling set case, nodes can map their colors and put  $\mathbf{X}$  on the right edges, such that the problem is solved correctly.

For the time lower bound, note that:

$$\text{prefix}^t([8, 15, 0, 0]) \leq \text{prefix}^t([15, 15, 15, 15]) \leq 15 \cdot \text{prefix}^t([1, 1, 1, 1]) \leq 15 \cdot \text{prefix}^{t+1}([1, 0, 0, 0]).$$

Hence, by [Lemma 8.5](#), [Observation 8.7](#), and [Theorem 8.1](#), the claim follows.  $\square$

## 9 Upper Bounds

We first discuss the problem of computing generalized arbdefective colorings as defined in [Definition 1.1](#). We start with a simple lemma, which states that if we are given a proper vertex  $C$  coloring of a graph, as long as a given arbdefective coloring instance is 1-relaxed, we can compute a valid solution by iterating over the  $C$  color classes.

**Lemma 9.1.** *Let  $G = (V, E)$  be a properly  $m$ -colored graph. Any 1-relaxed generalized arbdefective coloring problem on  $G$  can be solved in  $O(m)$  rounds deterministically.*

*Proof.* Let  $C$  be the number of colors of the given generalized arbdefective coloring instance and let  $\vec{d} = (d_1, \dots, d_C)$  be the arbdefect vector. The coloring of the nodes  $V$  and the orientation of the edges  $E$  is computed greedily in  $m$  phases by iterating over the  $m$  colors of the given proper  $m$ -coloring of  $G$ . Assume that the proper initial coloring assigns color  $\gamma(v) \in [m]$  to some node  $v \in V$ . The edge orientation is defined as follows. For every edge  $\{u, v\} \in E$ , the edge is oriented from  $u$  to  $v$  if  $\gamma(u) > \gamma(v)$  and it is oriented from  $v$  to  $u$  otherwise. In phase  $p \in \{1, \dots, m\}$ , the nodes  $v$  of initial color  $\gamma(v) = p$  are colored. Note that the given edge orientation implies that when coloring a node  $v$ , all its outneighbors have already been colored in earlier phases and because the initial  $m$ -coloring is proper, we also never color two adjacent nodes in the same phase. Each node  $v \in V$  that picks a color can now just pick a color  $x \in [C]$  for which the number of outneighbors of color  $x$  is at most  $d_x$ . Note that such a color must exist because  $v$  has at most  $\Delta$  outneighbors and because the arbdefect vector  $\vec{d}$  is 1-relaxed and therefore  $\sum_{i=1}^C (d_i + 1) > \Delta$ .  $\square$

[Theorem 9.2](#) now follows directly from the above lemma.

**Theorem 9.2.** *Let  $G = (V, E)$  be an  $n$ -node graph. Any 1-relaxed generalized arbdefective coloring problem on  $G$  can be solved in  $O(\Delta + \log^* n)$  rounds deterministically.*

*Proof.* We can first compute a proper  $(\Delta + 1)$ -coloring of  $G$  in time  $O(\Delta + \log^* n)$  (e.g., by using the algorithm of [\[BEK14\]](#)) and then apply [Lemma 9.1](#).  $\square$

The following lemma provides a simple generalization the  $O(\Delta/d + \log^* n)$ -round  $d$ -arbdefective  $O(\Delta/d)$ -coloring algorithm of [BEG18].

**Lemma 9.3.** *For every  $\alpha \geq 0$  and every  $\eta \geq 1$ , an  $\alpha$ -arbdefective  $O(\eta\Delta/\alpha)$ -coloring of an  $n$ -node graph  $G$  with maximum degree  $\Delta$  can be computed deterministically in  $O(\frac{\Delta}{\eta\alpha} + \log^* n)$  rounds.*

*Proof.* We first compute  $\Delta/\eta$ -defective  $O(\eta^2)$ -coloring of  $G$  in time  $O(\log^* n)$  by using the distributed defective coloring algorithm of [Kuh09]. Each of the  $O(\eta^2)$  color classes now has a maximum degree of  $\leq \Delta/\eta$  and we can therefore compute an  $\alpha$ -arbdefective  $O(\Delta/(\eta\alpha))$ -coloring in time  $O(\Delta/(\eta\alpha) + \log^* n)$ . In combination with the  $O(\eta^2)$  colors of the initial  $\Delta/\eta$ -defective coloring, we obtain the desired  $\alpha$ -arbdefective coloring of  $G$  with  $O(\eta\Delta/\alpha)$  colors.  $\square$

## 9.1 Algorithms for Arbdefective Colored Ruling Sets

We first show that if we are already given an appropriate arbdefective coloring, an  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling set can be computed in time matching the lower bound given in Theorem 1.5.

**Theorem 1.6.** *Let  $G = (V, E)$  be a graph and let  $\alpha \geq 0$ ,  $c \geq 1$ ,  $\beta \geq 1$ , and  $C \geq c$ . If an  $\alpha$ -arbdefective  $C$ -coloring of  $G$  is given, then an  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling set of  $G$  can be computed deterministically in the minimum time  $t$  such that  $\binom{t+\beta}{\beta} \geq \lceil \frac{C}{c} \rceil$  in the LOCAL model.*

*Proof.* We partition the  $C$  colors of the given  $\alpha$ -arbdefective  $C$ -coloring into  $K := \lceil C/c \rceil$  groups of size at most  $c$ . Within each of the  $K$  groups, we rename the at most  $c$  colors from 1 to  $c$ . For every  $i \in [K]$  let  $V_i$  be the set of nodes in color group  $i$ . Note that the given coloring provided an  $\alpha$ -arbdefective  $c$ -coloring for each of the induced subgraphs graphs  $G[V_i]$ . Let  $E' \subseteq E$  be the set of edges between two nodes of the same color group  $i \in [K]$  and consider the subgraph  $H = (V, E \setminus E')$  of  $G$ . Note that each set  $V_i$ ,  $i \in [K]$  is an independent set of  $H$  and therefore the partitioning into  $K$  color groups provides a proper  $K$ -coloring of  $H$ . We now compute a  $(2, \beta)$ -ruling set  $S \subseteq V$  of  $H$ . By using the algorithm of [BBO20] (Lemma 15 in the full version of [BBO20]), a  $(2, \beta)$ -ruling set  $S$  of a  $K$ -colored graph can be computed in the minimum time  $t$  such that  $\binom{t+\beta}{\beta} \geq K$ . The set  $S$  is an independent set of  $H$  and therefore the only edges of the subgraph  $G[S]$  of  $G$  induced by the nodes in  $S$  are edges between two nodes  $u, v \in V_i$  for some  $i$  (i.e., between nodes of the same color group). We therefore know that the initial  $\alpha$ -arbdefective coloring provides an  $\alpha$ -arbdefective  $c$ -coloring of  $G[S]$  and  $S$  therefore is an  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling set of  $G$ .  $\square$

By using the fact that  $\binom{a}{b} \geq (a/b)^b$ , we directly obtain the following corollary.

**Corollary 9.4.** *Let  $G = (V, E)$  be a graph and let  $\alpha \geq 0$ ,  $c \geq 1$ ,  $\beta \geq 1$ , and  $C \geq c$ . If one is given an  $\alpha$ -arbdefective  $C$ -coloring of  $G$ , then an  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling set of  $G$  can deterministically be computed  $O(\beta \cdot (C/c)^{1/\beta})$  rounds in the LOCAL model.*

In order to compute the initial arbdefective coloring that is required to apply Corollary 9.4, we can use Lemma 9.3. This gives the following complexity for computing an  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling set.

**Corollary 9.5.** *Let  $G = (V, E)$  be a graph and let  $\alpha \geq 0$ ,  $c \geq 1$ ,  $\beta \geq 1$ , and  $C \geq c$ . If one is given an  $\alpha$ -arbdefective  $C$ -coloring of  $G$ , then an  $\alpha$ -arbdefective  $c$ -colored  $\beta$ -ruling set of  $G$  can deterministically be computed  $O(\beta(\frac{\Delta}{(\alpha+1)\sqrt{c}})^{2/(\beta+1)})$  rounds in the LOCAL model.*

*Proof.* We first compute an  $\alpha$ -arbdefective  $C := O(\eta\Delta/\alpha)$ -coloring. By Lemma 9.3, this can be done in  $O(\Delta/(\eta\alpha) + \log^* n)$  rounds. We choose  $\eta$  such that  $\Delta/(\eta\alpha) = (C/c)^{1/\beta} = (\eta\Delta/(\alpha c))^{1/\beta}$ . We obtain  $\eta = (\Delta/\alpha)^{\frac{\beta-1}{\beta+1}} \cdot c^{\frac{1}{\beta+1}}$ , which together with Corollary 9.4 gives the required round complexity.  $\square$

## 10 Open Problems

There is a natural distinction between locally checkable problems: those who are solvable in  $O(f(\Delta) + \log^* n)$ , that is, *easy* problems, and those who require  $\Omega(\log_\Delta n)$  for deterministic algorithms and  $\Omega(\log_\Delta \log n)$  for randomized ones, that is, *hard* problems. We also know from prior work that, for bounded-degree graphs, there cannot be problems in between [CKP19, CP19]. In this work, we proved lower bounds for many easy natural problems, by exploiting the hardness of hard problems, namely,  $c$ -coloring for  $c \leq \Delta$ . For this purpose, we exploited the fact that there exists a fixed point proof for the hardness of  $\Delta$ -coloring. But what if there were no such proof? It would have been much more difficult to prove a tight lower bound for ruling sets and other problems. Hence, a first question that we ask is whether all hard problems admit a proof based on fixed points.

**Open Problem 1.** *Prove or disprove that, for all locally checkable problems with deterministic complexity  $\Omega(\log_\Delta n)$  in the LOCAL model, there exists a relaxation that is a non-trivial fixed point.*

Another interesting question is how broadly our technique can be applied. By applying the round elimination technique to  $O(f(\Delta) + \log^* n)$ -round solvable problems, we typically obtain problem sequences where the size of the problems grows really fast. Often, it seems that *most* of the configurations allowed in the problems of the sequence do not actually make the problem easier; they just make the description of the problem more complicated, similarly to what happens with ruling sets, where these configurations are related to  $c$ -coloring. It would be very interesting to find other cases where we can get rid of these “useless” allowed configurations by constructing some fixed point problem that we can embed into the description of the problem to make these configurations disappear.

**Open Problem 2.** *Can we find other interesting problems that can be solved in  $O(f(\Delta) + \log^* n)$  rounds for which we can prove lower bounds by exploiting the hardness of problems that require  $\Omega(\log_\Delta n)$ ?*

Also, we now have three different lower bound proofs based on round elimination for the maximal independent set problem on trees. The first one appeared in [BBO20]; in that proof the number of labels of the problems in the lower bound sequence grows linearly and is always upper bounded by  $O(\Delta)$ , and with similar techniques the authors also obtain lower bounds for ruling sets. The second one appeared in [BBKO21], and in that case the number of labels is always  $O(1)$ , but the proof does not extend to ruling sets. In both cases, the obtained lower bounds are exponentially worse than the lower bounds presented in this work, but here we needed to use a number of labels that is exponential in  $\Delta$ . We would like to understand this phenomenon, and better understand the relation between the number of labels and the achievable lower bounds. For example, can we prove the same lower bound presented in our work by using a subexponential number of labels?

**Open Problem 3.** *Understand the tradeoff between the number of labels and the maximum lower bound achievable with a round elimination problem sequence where each problem does not exceed this number of labels.*

In our work we characterized exactly which variants of arbdefective coloring are easy, and which variants of arbdefective coloring are hard. The same question is open for another important subroutine in efficient coloring algorithms—defective coloring.

**Open Problem 4.** *Characterize which variants of the defective coloring problem can be solved in  $O(f(\Delta) + \log^* n)$  rounds, for some function  $f$ , and which of them require  $\Omega(\log_\Delta n)$  rounds.*



In order to apply the round elimination technique, we need to encode problems in the node-edge formalism. Some natural problems, for example  $(2, \beta)$ -ruling sets, cannot be encoded directly, and we need to require nodes to produce a node-edge checkable proof that they solved the problem. We call this variant of ruling sets an *encoded* version of the problem. In the encoded variant of  $(2, \beta)$ -ruling sets, we require nodes to output pointer chains that point towards a ruling set node, such that these chains are of length at most  $\beta$ . Unfortunately, this means that if we obtain a lower bound of  $T$  rounds for the encoded problem, then we only obtain a  $(T - \beta)$ -round lower bound for the natural problem. Is the encoded version really  $\beta$  rounds harder than the natural one? For the encoded version of  $(2, \beta)$ -ruling sets we obtained truly tight bounds, whereas for the natural version we only obtained asymptotically tight bounds, and only for  $\beta$  small enough. While we think that encoded versions are interesting per se, we would also like to know the right complexity of the natural versions. In particular, for the natural versions, our lower bounds only hold for  $\beta \leq \varepsilon \log \Delta$  and some small enough  $\varepsilon$ , and we would like to know if they also hold for larger values of  $\beta$ . We conjecture that  $(2, O(\log \Delta))$ -ruling sets require  $\Omega(\log \Delta)$  rounds.

**Open Problem 5.** *Understand the difference between the complexity of encoded problems and natural problems.*

A limit of the current version of round elimination is that the best achievable lower bound, when  $\Delta \geq 3$ , is  $\Omega(\log_{\Delta} n)$  rounds for deterministic algorithms and  $\Omega(\log_{\Delta} \log n)$  rounds for randomized ones. We would like to understand whether this is really a limit of the technique or just of our current use of the general idea behind it.

**Open Problem 6.** *Can we use the round elimination technique to prove lower bounds of  $\omega(\log_{\Delta} n)$  rounds for deterministic algorithms, or of  $\omega(\log_{\Delta} \log n)$  rounds for randomized ones?*

In particular, in bounded degree graphs, we do not have any example of a problem that is known to be solvable in  $O(\text{poly log log } n)$  rounds with randomized algorithms, but also known to require  $\omega(\log n)$  rounds with deterministic ones.

**Open Problem 7.** *Prove or disprove that there exists a locally checkable problem that, in bounded degree graphs, is solvable in  $O(\text{poly log log } n)$  rounds with randomized algorithms in the LOCAL model, but that requires  $\omega(\log n)$  rounds for deterministic algorithms.*

We would like to mention that a related question has been asked in 2017 by Chang and Pettie [CP19]. There, the authors asked whether the distributed constructive Lovász Local Lemma problem can be solved in  $O(\log \log n)$  rounds on bounded degree graphs. It is known that answering this question in the affirmative would imply that there are no locally checkable problems in bounded degree graphs with randomized complexity between  $\omega(\log \log n)$  and  $o(\log n)$  [CP19], and it is also known that, for locally checkable problems on bounded degree graphs, an  $O(\log \log n)$  randomized algorithm implies an  $O(\log n)$  deterministic one [CKP19].

One of the few natural problems that is currently a candidate for answering **Open Problem 7** is  $\Delta$ -coloring. This is a problem for which randomness helps, and it can in fact be solved in  $O(\log^2 \log n)$  rounds with randomized algorithms in bounded degree graphs [GHKM18]. For deterministic algorithms, a lower bound of  $\Omega(\log n)$  rounds is known, but the best known upper bound is  $O(\log^2 n)$  rounds [GHKM18, PS95]. The currently best known algorithm uses ruling sets as a subroutine, and  $O(\log^2 n)$  is in fact the time that is used to compute a  $(2, \log n)$ -ruling set in  $G^{O(\log n)}$ , the  $O(\log n)$ -th power of  $G$ . While it may be possible to improve the complexity of  $\Delta$ -coloring by exploiting the fact that the ruling set is computed on a power graph (where, for some reason, it may be easier to compute ruling sets), we would like to understand whether computing a ruling set is really necessary for solving

$\Delta$ -coloring fast, or if some other (easier) symmetry breaking primitive is sufficient (which, given our lower bounds for ruling sets, might provide a more promising way of improving the complexity of  $\Delta$ -coloring).

**Open Problem 8.** *The best known algorithm for  $\Delta$ -coloring has a deterministic runtime of  $O(\log^2 n)$  rounds in bounded degree graphs, and uses ruling sets as a subroutine. Can we find a genuinely different algorithm for  $\Delta$ -coloring, not based on ruling sets, with a runtime that is at least as good?*

Proofs based on the round elimination technique have been getting progressively more complicated. We have automatic tools to efficiently check them for specific values of  $\Delta$  [Oli19], but generalizing these proofs for any value of  $\Delta$  usually requires a complex analysis (see e.g. the proof of Lemma 6.1). It would be nice to find ways to prove such statements in some easier way, or to entirely automate the process. Also, it would be interesting to have machine-checked versions of existing proofs.

**Open Problem 9.** *Can we find an automatic way to prove statements that are based on round elimination?*

## References

- [ABI86] Noga Alon, László Babai, and Alon Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *Journal of Algorithms*, 7(4):567–583, 1986.
- [AGLP89] Baruch Awerbuch, Andrew V. Goldberg, Michael Luby, and Serge A. Plotkin. Network decomposition and locality in distributed computation. In *Proc. 30th Symp. on Foundations of Computer Science (FOCS)*, pages 364–369, 1989.
- [Bar16] Leonid Barenboim. Deterministic  $(\Delta+1)$ -Coloring in Sublinear (in  $\Delta$ ) Time in Static, Dynamic, and Faulty Networks. *Journal of ACM*, 63(5):1–22, 2016.
- [BBE<sup>+</sup>20] Alkida Balliu, Sebastian Brandt, Yuval Efron, Juho Hirvonen, Yannic Maus, Dennis Olivetti, and Jukka Suomela. Classification of distributed binary labeling problems. In *Proc. 34th Symp. on Distributed Computing (DISC)*, 2020.
- [BBH<sup>+</sup>19] Alkida Balliu, Sebastian Brandt, Juho Hirvonen, Dennis Olivetti, Mikaël Rabie, and Jukka Suomela. Lower bounds for maximal matchings and maximal independent sets. In *Proc. 60th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 481–497, 2019.
- [BBKO21] Alkida Balliu, Sebastian Brandt, Fabian Kuhn, and Dennis Olivetti. Improved distributed lower bounds for MIS and bounded (out-)degree dominating sets in trees. In *Proc. 40th ACM Symposium on Principles of Distributed Computing (PODC)*, 2021.
- [BBO20] Alkida Balliu, Sebastian Brandt, and Dennis Olivetti. Distributed lower bounds for ruling sets. In *Proc. 61st IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 365–376, 2020.
- [BE10] Leonid Barenboim and Michael Elkin. Sublogarithmic distributed MIS algorithm for sparse graphs using Nash-Williams decomposition. *Distributed Comput.*, 22:363–379, 2010.
- [BE11] Leonid Barenboim and Michael Elkin. Deterministic distributed vertex coloring in polylogarithmic time. *Journal of ACM*, 58:23:1–23:25, 2011.

- [BE13] Leonid Barenboim and Michael Elkin. *Distributed Graph Coloring: Fundamentals and Recent Developments*. Morgan & Claypool Publishers, 2013.
- [BEG18] Leonid Barenboim, Michael Elkin, and Uri Goldenberg. Locally-iterative distributed  $(\Delta + 1)$ -coloring below Szegedy-Vishwanathan barrier, and applications to self-stabilization and to restricted-bandwidth models. In *Proc. 37th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 437–446, 2018.
- [BEK14] Leonid Barenboim, Michael Elkin, and Fabian Kuhn. Distributed  $(\Delta+1)$ -Coloring in Linear (in  $\Delta$ ) Time. *SIAM Journal on Computing*, 43(1):72–95, 2014.
- [BEPS16] Leonid Barenboim, Michael Elkin, Seth Pettie, and Johannes Schneider. The Locality of Distributed Symmetry Breaking. *Journal of the ACM*, 63(3):1–45, 2016.
- [BFH<sup>+</sup>16] Sebastian Brandt, Orr Fischer, Juho Hirvonen, Barbara Keller, Tuomo Lempinen, Joel Rybicki, Jukka Suomela, and Jara Uitto. A lower bound for the distributed Lovász local lemma. In *Proceedings of the 48th ACM Symposium on Theory of Computing (STOC 2016)*, pages 479–488. ACM Press, 2016.
- [BHOS19] Alkida Balliu, Juho Hirvonen, Dennis Olivetti, and Jukka Suomela. Hardness of minimal symmetry breaking in distributed computing. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, (PODC)*, pages 369–378, 2019.
- [BKP14] Tushar Bisht, Kishore Kothapalli, and Sriram V. Pemmaraju. Brief announcement: Super-fast  $t$ -ruling sets. In *Proceedings of the 2014 Symposium on Principles of Distributed Computing (PODC)*, pages 379–381, 2014.
- [BO20] Sebastian Brandt and Dennis Olivetti. Truly tight-in- $\Delta$  bounds for bipartite maximal matching and variants. In *Proc. 39th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 69–78, 2020.
- [Bra19] Sebastian Brandt. An automatic speedup theorem for distributed problems. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, pages 379–388, 2019.
- [Bro41] Rowland Leonard Brooks. On colouring the nodes of a network. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 37, pages 194–197. Cambridge University Press, 1941.
- [CHL<sup>+</sup>20] Yi-Jun Chang, Qizheng He, Wenzheng Li, Seth Pettie, and Jara Uitto. Distributed edge coloring and a special case of the constructive Lovász local lemma. *ACM Transactions on Algorithms*, 16(1):8:1–8:51, 2020.
- [CKP19] Yi-Jun Chang, Tsvi Kopelowitz, and Seth Pettie. An exponential separation between randomized and deterministic complexity in the LOCAL model. *SIAM Journal on Computing*, 48(1):122–143, 2019.
- [CLP18] Yi-Jun Chang, Wenzheng Li, and Seth Pettie. An optimal distributed  $(\Delta+1)$ -coloring algorithm? In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, (STOC 2018)*, pages 445–456, 2018.
- [CP19] Yi-Jun Chang and Seth Pettie. A time hierarchy theorem for the LOCAL model. *SIAM J. Comput.*, 48(1):33–69, 2019.

- [CPS17] K. Censor-Hillel, M. Parter, and G. Schwartzman. Derandomizing local distributed algorithms under bandwidth restrictions. In *Proc. 31st Symp. on Distributed Computing (DISC)*, pages 11:1–11:16, 2017.
- [FG17] Manuela Fischer and Mohsen Ghaffari. Sublogarithmic distributed algorithms for Lovász local lemma, and the complexity hierarchy. In *31st International Symposium on Distributed Computing, DISC 2017, October 16-20, 2017, Vienna, Austria*, volume 91 of *LIPICs*, pages 18:1–18:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.
- [FHK16] Pierre Fraigniaud, Marc Heinrich, and Adrian Kosowski. Local conflict coloring. In *Proc. 57th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 625–634, 2016.
- [FP20] Pierre Fraigniaud and Ami Paz. The topology of local computing in networks. In *Proc. 47th Int. Coll. on Automata, Languages, and Programming (ICALP)*, pages 128:1–128:18, 2020.
- [GGR21] Mohsen Ghaffari, Christoph Grunau, and Václav Rozhon. Improved deterministic network decomposition. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 2904–2923, 2021.
- [Gha16] Mohsen Ghaffari. An improved distributed algorithm for maximal independent set. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2016)*, pages 270–277, 2016.
- [Gha19] Mohsen Ghaffari. Distributed maximal independent set using small messages. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA 2019)*, pages 805–820, 2019.
- [GHK18] Mohsen Ghaffari, David G. Harris, and Fabian Kuhn. On derandomizing local distributed algorithms. In *Proc. 59th Symp. on Foundations of Computer Science (FOCS)*, pages 662–673, 2018.
- [GHKM18] Mohsen Ghaffari, Juho Hirvonen, Fabian Kuhn, and Yannic Maus. Improved distributed  $\Delta$ -coloring. In *Proc. 37th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 427–436, 2018.
- [GK21] Mohsen Ghaffari and Fabian Kuhn. Deterministic distributed vertex coloring: Simpler, faster, and without network decomposition. In *Proc. 62nd IEEE Symp. on Foundations of Computing (FOCS)*, 2021.
- [GP19] Mohsen Ghaffari and Julian Portmann. Improved Network Decompositions Using Small Messages with Applications on MIS, Neighborhood Covers, and Beyond. In *Proceedings of the 33rd International Symposium on Distributed Computing (DISC 2019)*, pages 18:1–18:16, 2019.
- [GV07] Beat Gfeller and Elias Vicari. A randomized distributed algorithm for the maximal independent set problem in growth-bounded graphs. In *Proceedings of the 2007 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 53–60, 2007.
- [Hal35] P. Hall. On representatives of subsets. *Journal of the London Mathematical Society*, s1-10(1):26–30, 1935.

- [HKN16] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. A deterministic almost-tight distributed algorithm for approximating single-source shortest paths. In *Proceedings 48th Symp. on Theory of Computing (STOC 2016)*, pages 489–498, 2016.
- [KMW16] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. Local computation: Lower and upper bounds. *Journal of ACM*, 63(2):17:1–17:44, 2016.
- [KMW18] Fabian Kuhn, Yannic Maus, and Simon Weidner. Deterministic distributed ruling sets of line graphs. In *Structural Information and Communication Complexity - 25th International Colloquium, SIROCCO 2018, Ma'ale HaHamisha, Israel, June 18-21, 2018, Revised Selected Papers*, pages 193–208, 2018.
- [KP12] Kishore Kothapalli and Sriram V. Pemmaraju. Super-fast 3-ruling sets. In *Annual Conference on Foundations of Software Technology and Theoretical Computer Science (IARCS 2012)*, volume 18, pages 136–147, 2012.
- [Kuh09] Fabian Kuhn. Local weak coloring algorithms and implications on deterministic symmetry breaking. In *Proc. 21st ACM Symp. on Parallelism in Algorithms and Architectures (SPAA)*, 2009.
- [Kuh20] Fabian Kuhn. Faster deterministic distributed coloring through recursive list coloring. In *Proc. 32nd ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 1244–1259, 2020.
- [KW85] Richard M. Karp and Avi Wigderson. A fast parallel algorithm for the maximal independent set problem. *Journal of ACM*, 32:762–773, 1985.
- [Lin87] Nathan Linial. Distributive graph algorithms – Global solutions from local data. In *Proc. 28th Symp. on Foundations of Computer Science (FOCS 1987)*, pages 331–335. IEEE, 1987.
- [Lin92] Nathan Linial. Locality in Distributed Graph Algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992.
- [Lub86] Michael Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM Journal on Computing*, 15(4):1036–1053, 1986.
- [LW11] Christoph Lenzen and Roger Wattenhofer. MIS on trees. In *Proceedings of the 2011 Annual ACM Symposium on Principles of Distributed (PODC)*, pages 41–48, 2011.
- [Mau21] Yannic Maus. Distributed graph coloring made easy. In *Proc. 33rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2021.
- [MT20] Yannic Maus and Tigran Tonoyan. Local conflict coloring revisited: Linial for lists. *CoRR*, abs/2007.15251, 2020.
- [Nao91] Moni Naor. A lower bound on probabilistic algorithms for distributive ring coloring. *SIAM Journal on Discrete Mathematics*, 4(3):409–412, 1991.
- [Oli19] Dennis Olivetti. Round Eliminator: a tool for automatic speedup simulation, 2019.
- [Pel00] David Peleg. *Distributed Computing: A Locality-Sensitive Approach*. Society for Industrial and Applied Mathematics, 2000.

- [PPP<sup>+</sup>17] Shreyas Pai, Gopal Pandurangan, Sriram V. Pemmaraju, Talal Riaz, and Peter Robinson. Symmetry breaking in the CONGEST model: Time- and message-efficient algorithms for ruling sets. In *Proceedings of 31st International Symposium on Distributed Computing, (DISC 2017)*, volume 91, pages 38:1–38:16, 2017.
- [PS95] Alessandro Panconesi and Aravind Srinivasan. The local nature of  $\Delta$ -coloring and its algorithmic applications. *Combinatorica*, 15(2):255–280, 1995.
- [PS96] Alessandro Panconesi and Aravind Srinivasan. On the Complexity of Distributed Network Decomposition. *Journal of Algorithms*, 20(2):356–374, 1996.
- [RG20] Václav Rozhoň and Mohsen Ghaffari. Polylogarithmic-time deterministic network decomposition and distributed derandomization. In *Proceedings of 52nd Annual ACM Symposium on Theory of Computing (STOC 2020)*, 2020.
- [SEW13] Johannes Schneider, Michael Elkin, and Roger Wattenhofer. Symmetry breaking depending on the chromatic number or the neighborhood growth. *Theoretical Computer Science*, 509:40–50, 2013.
- [SW10a] Johannes Schneider and Roger Wattenhofer. A new technique for distributed symmetry breaking. In *Proceedings of the 2010 Annual ACM Symposium on Principles of Distributed Computing, (PODC)*, pages 257–266, 2010.
- [SW10b] Johannes Schneider and Roger Wattenhofer. An optimal maximal independent set algorithm for bounded-independence graphs. *Distributed Computing*, 22:349–361, 2010.

## A Proof of Theorem 7.1

All proofs present in this section are already present in the literature [BBH<sup>+</sup>19, BBO20, BO20, BBE<sup>+</sup>20]. We report them here for completeness, but we also make them slightly more general: while the existing proofs only tolerate problem sequences containing  $O(\Delta^2)$  labels, we remove this limitation and we assume that the number of labels is upper bounded by  $f(\Delta)$ , for some arbitrary function  $f$ .

The standard way for lifting lower bounds, obtained through round elimination, from the port numbering model to the LOCAL model is the following:

1. Show that the round elimination technique also works for randomized algorithms for the port numbering model. In particular, show that if  $\Pi$  can be solved in  $T$  rounds with local failure probability at most  $p$  (where the local failure probability of a node  $v$  is the probability that the constraints of some problem  $\Pi$  are not satisfied on  $v$ ), then  $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$  can be solved in  $T - 1$  rounds with some local failure probability at most  $p'$ , for  $p'$  that is not too large compared to  $p$ .
2. Show that if a problem is not 0 rounds solvable in the deterministic port numbering model, then any 0 round algorithm for the problem must fail with large probability in the randomized port numbering model.
3. Combine the two parts above to show that any algorithm that runs “too fast”, that is, faster than the deterministic port numbering model lower bound, must fail with large probability.

4. Show for which values of  $T$ , as a function of  $n$ , the obtained local failure probability is too large. This gives a lower bound on the local failure probability of randomized algorithms that run in infinite  $\Delta$ -regular trees in the port numbering model, as a function of their running time.
5. Show that the last point implies a lower bound in a finite  $\Delta$ -regular tree in the randomized port numbering model. This directly gives a lower bound in the randomized LOCAL model as well.
6. Show that a randomized lower bound implies an even stronger deterministic lower bound.

While this was not emphasized in previous works, the standard techniques used to lift lower bounds from the deterministic port numbering model to the LOCAL model do not work if there are inputs that can be adversarially assigned. In our setting, the inputs of the nodes are

- the size of the graph,
- the random bits,
- the port numbering.

In order to make the lifting work, we need to assume that the port numbering is random among the ones allowed by the port numbering constraints.

We start by discussing [Item 1](#). In [\[BBE<sup>+</sup>20\]](#), the authors show that, for a random port numbering assignment, the round elimination technique also works for randomized algorithms, and they provide a generic analysis, that does not depend on a specific problem, that shows how the local failure probability evolves when applying the round elimination technique. The same result holds also if the random port number assignment comes from a locally constrained port numbering assignment. This result holds also for hypergraphs, and in that case  $\Delta$  should be an upper bound on the degree of the nodes and on the rank of the hyperedges.

**Lemma A.1** (Lemma 41 of [\[BBE<sup>+</sup>20\]](#)). *Let  $A$  be a randomized  $t$ -round algorithm for  $\Pi$  with local failure probability at most  $p$  (where  $t > 0$ ). Let  $\Pi'$  be a relaxation of  $\mathcal{R}(\Pi)$ . Then there exists a randomized  $(t - 1)$ -round algorithm  $A'$  for  $\overline{\mathcal{R}}(\Pi')$  with local failure probability  $p'' \leq 2^{\frac{1}{\Delta+1}} (\Delta|\Sigma'|)^{\frac{\Delta}{\Delta+1}} p'^{\frac{1}{\Delta+1}} + p'$ , where  $p' \leq 2^{\frac{1}{\Delta+1}} (\Delta|\Sigma|)^{\frac{\Delta}{\Delta+1}} p^{\frac{1}{\Delta+1}} + p$  and  $\Sigma'$  is the label set of  $\Pi'$ .*

This lemma essentially states that, if there exists an algorithm  $A$  that solves  $\Pi$  with some local failure probability, then we can make  $A$  running one round faster and solve (a relaxation of)  $\overline{\mathcal{R}}(\mathcal{R}(\Pi))$  with a local failure probability that is larger, but not too large.

We can apply [Lemma A.1](#) multiple times, to show an upper bound on the failure probability of an algorithm that solves (a relaxation of) the problem obtained by applying  $\overline{\mathcal{R}}(\mathcal{R}(\cdot))$  for  $j$  times, as a function of the failure probability of an algorithm solving  $\Pi$  in  $t$  rounds. This lemma (and its proof) already appeared in previous works [\[BBH<sup>+</sup>19, BBO20, BO20, BBE<sup>+</sup>20\]](#), but we now make it parametric on the number of labels  $f(\Delta)$  of the problem family.

**Lemma A.2.** *Let  $\Pi_0 \rightarrow \Pi_1 \rightarrow \dots \rightarrow \Pi_t$  be a sequence of problems satisfying the conditions of [Theorem 7.1](#). Let  $A$  be a randomized  $t$ -round algorithm for  $\Pi_0$  with local failure probability at most  $p$ . Then there exists a randomized  $(t - j)$ -round algorithm  $A'$  for  $\Pi_j$  with local failure probability at most  $(2\Delta f(\Delta))^2 p^{1/(\Delta+1)^{2j}}$ , for all  $0 < j \leq t$ .*

*Proof.* By assumption,  $|\Sigma|$  and  $|\Sigma'|$  are upper bounded by  $f(\Delta)$ . Hence, by [Lemma A.1](#), we obtain:

$$\begin{aligned}
p'' &\leq 2^{\frac{1}{\Delta+1}} (\Delta|\Sigma'|)^{\frac{\Delta}{\Delta+1}} p'^{\frac{1}{\Delta+1}} + p' \\
&\leq (\Delta|\Sigma'|)^{\frac{1}{\Delta+1}} (\Delta|\Sigma'|)^{\frac{\Delta}{\Delta+1}} p'^{\frac{1}{\Delta+1}} + p' \\
&= (\Delta|\Sigma'|) p'^{\frac{1}{\Delta+1}} + p' \\
&\leq (\Delta f(\Delta)) p'^{\frac{1}{\Delta+1}} + p' \\
&\leq (2\Delta f(\Delta)) p'^{\frac{1}{\Delta+1}}.
\end{aligned}$$

Similarly,

$$\begin{aligned}
p' &\leq 2^{\frac{1}{\Delta+1}} (\Delta|\Sigma|)^{\frac{\Delta}{\Delta+1}} p^{\frac{1}{\Delta+1}} + p \\
&\leq (\Delta|\Sigma|) p^{\frac{1}{\Delta+1}} + p \\
&\leq (2\Delta f(\Delta)) p^{\frac{1}{\Delta+1}}.
\end{aligned}$$

Summarizing, we get that:

$$\begin{aligned}
p'' &\leq (2\Delta f(\Delta)) p'^{\frac{1}{\Delta+1}}, \text{ where} \\
p' &\leq (2\Delta f(\Delta)) p^{\frac{1}{\Delta+1}}.
\end{aligned}$$

By recursively applying [Lemma A.1](#), we get the following:

$$p_j \leq (2\Delta f(\Delta)) p_{j-1}^{\frac{1}{\Delta+1}},$$

where  $p_0 = p$  and  $p_{2j}$ , are, respectively, the local failure probability bounds for  $\Pi_0$  and  $\Pi_j$ . We prove by induction that for all  $j > 0$ ,

$$p_j \leq (2\Delta f(\Delta))^2 p^{\frac{1}{(\Delta+1)^j}}.$$

For the base case where  $j = 1$ , we get that  $p_1 \leq (2\Delta f(\Delta))^2 p^{\frac{1}{\Delta+1}}$ , which holds, as we showed above. Let us assume that the claim holds for  $j$ , and let us prove it for  $j + 1$ . We obtain the following, where the second inequality holds by the inductive hypothesis.

$$\begin{aligned}
p_{j+1} &\leq 2\Delta f(\Delta) p_j^{\frac{1}{\Delta+1}} \leq 2\Delta f(\Delta) \left( (2\Delta f(\Delta))^2 p^{\frac{1}{(\Delta+1)^j}} \right)^{\frac{1}{\Delta+1}} \\
&\leq (2\Delta f(\Delta))^{1+\frac{2}{\Delta+1}} p^{\frac{1}{(\Delta+1)^{j+1}}} \leq (2\Delta f(\Delta))^2 p^{\frac{1}{(\Delta+1)^{j+1}}}.
\end{aligned}$$

□

We now handle [Item 2](#). We show that, the mere fact that a problem  $\Pi$  cannot be solved in 0 rounds in the deterministic port numbering model implies a lower bound on the local failure probability of any algorithm that tries to solve  $\Pi$  in the randomized port numbering model. Also in this case, we consider hypergraphs, and we assume  $\Delta$  to be an upper bound on the degree of the nodes and on the rank of the hyperedges.

**Lemma A.3.** *Let  $\Pi$  be a problem that cannot be solved in 0 rounds with deterministic algorithms in the port numbering model given a port numbering assignment that satisfies some constraints  $C^{\text{port}}$ , such that  $|\Sigma| \leq f(\Delta)$ , for some function  $f$ . Then any randomized 0-round algorithm solving  $\Pi$  must fail with probability at least  $\frac{1}{\Delta^3 \Delta^2 f(\Delta) \Delta^2}$ , if the port numbering is chosen uniformly at random among the ones allowed by  $C^{\text{port}}$ .*



*Proof.* Since in 0 rounds of communication all nodes have the same information, we can see any 0-round algorithm as a probability assignment to each node configuration. That is, for each  $c_i \in \mathcal{N}$ , the algorithm outputs  $c_i$  with probability  $p_i$ , such that  $\sum p_i = 1$ . Since the number of labels is at most  $f(\Delta)$ , then the number of possible configurations appearing in  $\mathcal{N}$  is at most  $f(\Delta)^\Delta$ . Hence, by the pigeonhole principle, there exists some configuration  $\bar{c}$  that all nodes output with probability at least  $\frac{1}{f(\Delta)^\Delta}$ .

Also, since there are at most  $\Delta!$  different permutations of  $\bar{c}$ , a node can map the labels of  $\bar{c}$  to its  $\Delta$  ports in at most  $\Delta! \leq \Delta^\Delta$  different ways. Hence, conditioned on a node to output  $\bar{c}$ , there is at least one permutation  $\sigma$  that is chosen with probability at least  $1/\Delta^\Delta$ . Hence, there is at least one configuration  $\bar{c}$  that is output using some permutation  $\sigma$  with probability at least  $\frac{1}{\Delta^\Delta f(\Delta)^\Delta}$ . The probability that all the nodes incident to the same hyperedge output the configuration  $\bar{c}$  using the same permutation  $\sigma$  is at least  $(\frac{1}{\Delta^\Delta f(\Delta)^\Delta})^\Delta = \frac{1}{\Delta^{3\Delta^2} f(\Delta)^{\Delta^2}}$ .

Since the problem is not 0-round solvable in the deterministic setting, we know that there is a port numbering assignment allowed by  $C^{\text{port}}$ , such that, if all nodes around the same hyperedge output  $\bar{c}$  with the same permutation  $\sigma$ , then the constraint on the hyperedge is not satisfied. If we consider a hyperedge  $h$ , and  $\delta$  nodes incident to it,  $v_1, \dots, v_\delta$ , there are  $2\Delta\delta \leq 2\Delta^2$  ports in total: for each  $v_i$  there are  $2\Delta$  ports, the ones connecting them to their incident hyperedges, and the ports connecting the hyperedges to them. Since the number of port assignments allowed by  $C^{\text{port}}$  is upper bounded by the total number of port assignments, that is at most  $\Delta^{2\Delta^2}$ , then the probability that all nodes incident to the same hyperedge output the same configuration using the same permutation such that the port numbering assignment makes this output fail on the hyperedge is at least  $\frac{1}{\Delta^{2\Delta^2} \Delta^{\Delta^2} f(\Delta)^{\Delta^2}} = \frac{1}{\Delta^{3\Delta^2} f(\Delta)^{\Delta^2}}$ .  $\square$

We now handle **Item 3**. We show that any algorithm that runs in strictly less than  $t$  rounds, where  $t$  is a lower bound for the deterministic port numbering model, must fail with large probability. This lemma also appeared in previous works [BBH<sup>+</sup>19, BBO20, BO20, BBE<sup>+</sup>20], and here we just make it parametric on  $f(\Delta)$ .

**Lemma A.4.** *Let  $\Pi_0 \rightarrow \Pi_1 \rightarrow \dots \rightarrow \Pi_t$  be a sequence of problems satisfying the conditions of [Theorem 7.1](#). Any algorithm for  $\Pi_0$  running in strictly less than  $t$  rounds must fail with probability at least  $(\frac{1}{f(\Delta)^{\Delta^{10t}}})$ .*

*Proof.* By applying [Lemma A.2](#) we get that an algorithm solving  $\Pi_0$  in  $t' < t$  rounds with local failure probability at most  $p$  implies an algorithm solving  $\Pi_{t'}$  in 0 rounds with local failure probability at most  $(2\Delta f(\Delta))^2 p^{1/(\Delta+1)^{2t'}}$ . Then, since  $\Pi_{t'}$  is not 0-round solvable in the port numbering model using deterministic algorithms by assumption, by applying [Lemma A.3](#) we get the following:

$$(2\Delta f(\Delta))^2 p^{1/(\Delta+1)^{2t'}} \geq \frac{1}{\Delta^{3\Delta^2} f(\Delta)^{\Delta^2}},$$

that implies the following:

$$\begin{aligned}
p &\geq \left( \frac{1}{\Delta^{3\Delta^2} f(\Delta)^{\Delta^2} (2\Delta f(\Delta))^2} \right)^{(\Delta+1)^{2t'}} \\
&\geq \left( \frac{1}{4\Delta^{3\Delta^2+2} f(\Delta)^{\Delta^2+2}} \right)^{(\Delta+1)^{2t'}} \\
&\geq \left( \frac{1}{f(\Delta)^{\Delta^2+2+(3\Delta^2+4)\log\Delta}} \right)^{(\Delta+1)^{2t'}} \\
&\geq \left( \frac{1}{f(\Delta)^{6\Delta^3}} \right)^{(\Delta+1)^{2t'}} \\
&\geq \left( \frac{1}{f(\Delta)^{6\Delta^3}} \right)^{\Delta^{4t'}} \\
&\geq \frac{1}{f(\Delta)^{6\Delta^3\Delta^{4t'}}} \\
&\geq \frac{1}{f(\Delta)^{\Delta^{10t'}}} \\
&\geq \frac{1}{f(\Delta)^{\Delta^{10t}}}.
\end{aligned}$$

□

We now handle [Item 4](#), and show that if an algorithm runs “too fast” then it must fail with probability strictly larger than  $1/n$ . This lemma essentially says that, if an algorithm, on infinite  $\Delta$ -regular trees, runs in  $o(\min\{t, \log_\Delta \log n\})$ , for some given  $n$ , where  $t$  is a lower bound for the deterministic port numbering model, then it cannot locally succeed with high probability.

**Lemma A.5.** *Let  $\Pi_0 \rightarrow \Pi_1 \rightarrow \dots \rightarrow \Pi_t$  be a sequence of problems satisfying the conditions of [Theorem 7.1](#). Any randomized algorithm running in strictly less than  $\min\{t, \frac{1}{10}(\log_\Delta \log n - \log_\Delta \log f(\Delta))\}$  rounds must fail with probability  $> 1/n$ .*

*Proof.* By [Lemma A.4](#), any algorithm running in strictly less than  $t$  rounds must fail with probability  $p \geq \frac{1}{f(\Delta)^{\Delta^{10t}}}$ . We show that, for all  $t < \frac{1}{10}(\log_\Delta \log n - \log_\Delta \log f(\Delta))$ , it holds that  $p > 1/n$ .

$$\begin{aligned}
t &< \frac{1}{10}(\log_\Delta \log n - \log_\Delta \log f(\Delta)) && \Rightarrow \\
10t \log \Delta &< \log \log n - \log \log f(\Delta) && \Rightarrow \\
10t \log \Delta + \log \log f(\Delta) &< \log \log n && \Rightarrow \\
\Delta^{10t} \log f(\Delta) &< \log n && \Rightarrow \\
f(\Delta)^{\Delta^{10t}} &< n && \Rightarrow \\
\frac{1}{f(\Delta)^{\Delta^{10t}}} &> 1/n
\end{aligned}$$

□

We now handle [Item 5](#). We show that if an algorithm fails with large probability in infinite trees, then it also fails with large probability in finite trees. Again, this lemma also appeared in previous works [[BBH<sup>+</sup>19](#), [BBO20](#), [BO20](#), [BBE<sup>+</sup>20](#)], and here we just make it parametric on  $f(\Delta)$ .

**Lemma A.6.** *Let  $\Pi_0 \rightarrow \Pi_1 \rightarrow \dots \rightarrow \Pi_t$  be a sequence of problems satisfying the conditions of [Theorem 7.1](#). Any randomized algorithm running in the LOCAL model, in  $\Delta$ -regular balanced trees of  $n$  nodes, that fails with probability at most  $1/n$ , requires  $\Omega(\min\{t, \log_\Delta \log n - \log_\Delta \log f(\Delta)\})$  rounds.*

*Proof.* We prove the result by contradiction. Assume that there is an algorithm for the LOCAL model that contradicts our claim. Note that this algorithm can be simulated in the randomized port numbering model with essentially the same failure probability, since an algorithm in the randomized port numbering model can first generate unique IDs with high probability of success, and then simulate the LOCAL algorithm. This algorithm is for finite  $\Delta$ -regular balanced trees of  $n$  nodes. We show that we can use such an algorithm in infinite  $\Delta$ -regular trees and obtain the same local failure probability, contradicting [Lemma A.5](#).

The claim follows by a standard argument: we can run this algorithm on infinite trees, and if it fails with too large local failure probability, then we can consider the finite tree of  $n$  nodes obtained by considering the subgraph induced by the  $O(\log_\Delta n)$  neighborhood of the failing node. By running the algorithm on the obtained tree of  $n$  nodes, the view of the failing node itself, and its neighbors, is the same as in the original instance, since the running time is  $o(\log_\Delta n)$ . Hence, by a standard indistinguishability argument, the algorithm fails with probability strictly larger than  $1/n$ , contradicting the correctness of the given algorithm.  $\square$

We can finally handle [Item 6](#), and show that a strong lower bound for the randomized LOCAL model implies an even stronger lower bound for the deterministic LOCAL model. The proof uses a standard “lie about  $n$ ” argument, and again, it has been used in the context of round elimination in [[BBH<sup>+</sup>19](#), [BBO20](#), [BO20](#), [BBE<sup>+</sup>20](#)], and we report the same proof here modified to handle an arbitrary number of labels  $f(\Delta)$ .

**Lemma A.7.** *Let  $\Pi_0 \rightarrow \Pi_1 \rightarrow \dots \rightarrow \Pi_t$  be a sequence of problems satisfying the conditions of [Theorem 7.1](#). Any deterministic algorithm running in the LOCAL model, in  $\Delta$ -regular balanced trees of  $n$  nodes, requires  $\Omega(\min\{t, \log_\Delta n - \log_\Delta \log f(\Delta)\})$  rounds.*

*Proof.* The proof follows by a standard “lie about  $n$ ” argument. If we run an algorithm  $A$  and lie about the size of the graph, if the algorithm cannot notice the lie, then it must work correctly, with a running time that now depends on the value of  $n$  provided to the algorithm. In particular, in [[CKP19](#), [Theorem 6](#)], it is shown that, given an algorithm that runs in  $T(n)$  rounds, if we lie about the size of the graph and say that it has size  $N$ , then the obtained algorithm is correct and runs in  $T = T(N)$  rounds if the following conditions are satisfied:

- The IDs in each  $(T + 1)$ -radius neighborhood are unique and from  $\{1, 2, \dots, N\}$ ;
- Each  $(T + 1)$ -radius neighborhood contains at most  $N$  nodes.

Assume for a contradiction that we have a deterministic algorithm  $A$  that contradicts the claim. In particular, this means that it terminates in  $o(\log_\Delta n)$  rounds. Let us fix  $N = \log n$ , and let us construct a new algorithm  $A'$  that runs  $A$  by lying about the size of the graph, and claims that it has size  $N$ . In order to satisfy the first requirement, we need to create a new ID assignment from a smaller space. For this purpose, we can compute an  $N$ -coloring of  $G^{2T+2}$ , the  $(2T + 2)$ th power of  $G$ , using three iterations of Linial’s coloring algorithm [[Lin92](#), [Corollary 4.1](#)], and use the obtained coloring as a new ID assignment. This algorithm can be used to color a  $k$ -colored graph of maximum degree  $\bar{\Delta}$  with e.g.

$$O(\bar{\Delta}^2(\log \log \log k + \log \bar{\Delta}))$$

colors in three rounds. The power graph  $G^{2T+2}$  has maximum degree  $\bar{\Delta} \leq \Delta^{2T+2}$  and thus we get a coloring of  $G^{2T+2}$  with

$$O(\Delta^{2(2T+2)}(\log \log N + \log \Delta^{2T+2}))$$

colors. Since  $T(N) \in o(\log_{\Delta} N)$ , then each  $(T+1)$ -radius neighborhood contains  $O(\Delta^{T+1}) < N$  nodes (and hence the second requirement is satisfied) and that the number of colors is less than  $N$  (and hence the first requirement is satisfied). The coloring can be computed in time  $O(T)$  in the original graph  $G$ . Hence, if  $A$  is correct, then  $A'$  is also correct and terminates in  $o(\min\{t, \log_{\Delta} N - \log_{\Delta} \log f(\Delta)\}) = o(\min\{t, \log_{\Delta} \log n - \log_{\Delta} \log f(\Delta)\})$  rounds, but this contradicts the claim of [Lemma A.6](#).  $\square$

By combining [Lemma A.6](#) and [Lemma A.7](#), we obtain [Theorem 7.1](#).